A Model of Artifact Creation Reconsidered

David Pearce*

Universidad Politécnica de Madrid, Spain

Abstract. I examine a model of artifact creation developed by (Houkes & Vermaas 2009). I suggest a way to extend the model by regarding the agents involved in the design and making of new artifacts to be cooperating in the 'we-mode' as described by Raimo Tuomela. Contrary to the authors' own claims, I suggest that their model then becomes largely compatible with the 'intentionalist' view of artifacts propounded by Amie Thomasson and others, according to which the nature of a new artifactual kind is essentially dependent upon the intentions of its creator. In particular the model explains how an artifactual kind can be understood via a form of *collective* intentionality.

1 Introduction: Artifacts and Intentionality

We are surrounded by artifacts. PC-tablets, mobile phones, tables, chairs and coffee cups form a staple part of our living space and the physical ontology around us. Yet, artifacts, unlike rocks, lakes, salt and granite, also form part of our social ontology. This seems self-evident in the case of entities like coins of the realm and national flags that have an overtly social connotation; their very existence would be meaningless without a social and even institutional structure of support. But is the same true of bowler hats, light bulbs, churches and football stadia? And how do these items, from an ontological point of view, differ, if at all, from one another? And what of Facebook, Twitter and other artifacts that populate our tablets and smart phones?

We associate with Franz Brentano the idea that there is a realm of intentional phenomena not reducible in kind to the physical realm. In the 20th century a prominent exponent of the idea that human, cultural objects are intentional was Roman Ingarden, a student of two of Brentano's most illustrious students, Kazimierz Twardowski and Edmund Husserl. As illustrative examples, Ingarden took the cases of flags and churches that require social or institutional acts to give them meaning. Ingarden also stressed the dual, physical and intentional, structure of certain forms of art. For example, he referred to the two-layered structure of architectural works, as intentional artworks on the one hand and as 'real' buildings of bricks and mortar, on the other.[1]

This idea has been extended to artifacts more generally. What we might term the *intentionalist* view of artifacts holds that the nature of an artifact or artifactual kind depends on the intentions of its author or creator. The view was developed in particular by Risto Hilpinen who formulated a dependence condition as follows (Hilpinen 1992), p.65, (see also (Hilpinen 2011) for a more recent discussion):

* Partially supported by the project SINTELNET.

[1] See in particular (Ingarden 1961).

(DEP) The existence and some of the properties of an artifact depend on an author's intention to make an object of a certain kind.

Stronger or weaker forms of the intentionalist view are obtained by regarding the dependence condition as applying to few, many or all of the principal properties of an artifact. For example a strong intentionalist view of artifacts has been propounded by Amie Thomasson, especially in her writings (Thomasson 2003, Thomasson 2007).

Thomasson rejects a common view that artifactual kinds are individuated by means of their *function*. Her suggestion is that the creator's intentions form the most relevant guide to the extension of an artifact type (term) and so artifacts must be the intended products of human activities. A consequence of this view is that, since intentions are constitutive of meaning, makers (in a very general sense of maker) cannot be completely wrong about the nature of the artifacts they produce. As with social concepts, beliefs about artifacts may partly be constitutive of their meaning.

Thomasson defends her view by considering a simplified account of a single, individual artisan intentionally creating a prototype artifact of a new kind. Since the object is of a new type, there is no question of her aiming to copy an existing design or conform to a previously available specification of some product, to which she could succeed or fail. Instead the features that are relevant to the new kind K are those features intentionally given by her. She must have a substantive idea of what kind of object a K must be, along with the intention to produce that kind of object. By creating a new artifact and delineating a new kind she has stipulated the features that are relevant for and constitutive of that kind.

As Thomasson remarks, there is an element of 'self-referentiality' in artifactual concepts. It follows that there is a natural comparison but also a contrast with Searle's view of institutional concepts in (Searle 1995).

> Whereas for a certain sort of thing to be money, it is necessary (and sufficient) that it be the sort of thing that is collectively regarded as money, for an individual object to be a chair, it must itself have been intended to be a chair. ((Thomasson 2007), p.58)

In recent years the dual nature of *technical* artifacts has been highlighted and studied by a number of scholars working on the philosophical foundations of technology. At the beginning of this century a research programme entitled *The Dual Nature of Technical Artefacts* ran for several years coordinated by the Delft University of Technology in The Netherlands. According to its manifesto, this programme

> starts from the observation that technical artifacts have a twofold nature: They are (i) designed physical structures which realize (ii) intentionality-bearing functions. [2]

A key aspect of the programme was to investigate the nature of (i) and (ii) and especially their interrelations; this was realised through a substantial output of research papers and books. We will shortly examine one of these works in more detail.

[2] See http://www.dualnature.tudelft.nl/.

In the modern world of artifact creation and manufacturing typically many agents will be involved in the process of creating an artifact or indeed designing and making a new artifactual kind. Both Hilpinen and Thomasson accept that there may be a plurality of agents involved. The question then arises: What becomes of the intentionalist view? Does it become indefensible once individual intentionality is no longer relevant? Or should we rather explore the idea that individual intentionality should be replaced by a form of collective or we-intentionality? Collective intentionality has become a pivotal concept of social intelligence. The concept of we-intention was already introduced and analysed in the works of Wilfred Sellars (eg. in (Sellars 1968)). Later a systematic account of we-intentions was given by Raimo Tuomela (Tuomela 1984) and developed by (Tuomela & Miller 1985) as a central concept in the theory of social action. More recently John Searle has made we-intentionality a key component of his account of social reality, as in (Searle 2010).

The path to collective intentionality is highly inviting, but now it is clear that we will need a more sophisticated model of artifact creation: the single-agent, 'artisan' model is not designed to explain how collective intentionality underlies the nature of an artifactual kind. The alternative model that I will turn to is one that has been proposed by two authors contributing to the 'Dual Nature' research programme just mentioned.

2 An 'engineering' model of artifact creation

(Houkes & Vermaas 2004, Houkes & Vermaas 2009) have developed an elaborate account of artifacts that is intended to reflect engineering practice. It involves a detailed model of agents' actions, beliefs and intentions in terms of action plans relevant for artifact design and production. Like Thomasson they reject the idea that artifacts can be understood in terms of their functions and they aim to replace the function oriented philosophy of artifacts with an account based on intentional actions[3].

There are several reasons to take an interest in the account provided by Houkes and Vermaas. It purports to reflect modern engineering practice by distinguishing between, design, use and manufacture, as well as between the different agents involved in these processes. In particular their account includes a use plan, a design plan, as well as plans for product-designing, making and manufacturing. Each of these involves the intentions, goals and beliefs of the specific agents involved. Not only is this a genuinely multi-agent view that aims to reflect engineering practice, also, since it is action-theoretic, it can be embedded in more general theories of social action.

Neverthless, there is a major obstacle to overcome in applying this model in the present context. Houkes and Vermaas reject the intentionalist stance and claim that their model is incompatible with views such as those of Thomasson. While they accept the idea "that artifacts are intentionally produced by humans", ((Houkes & Vermaas 2009), p.410), with explicit reference to Thomasson they question the idea that artifactual kinds can be identified on the basis of makers' intentions. They also reject her suggestion that makers can in general be said to have substantively correct ideas about properties relevant to an artifactual kind. As they say of their own model:

[3] See especially (Houkes & Vermaas 2004). Although this aspect contrasts somewhat with the quoted manifesto of the research programme, the authors also investigate technical functions.

> It does not afford a clear relation between the intentions of any of the agents involved in producing the artefact and membership of an artefact kind. ((Houkes & Vermaas 2009), p.404).

The main thesis that Houkes and Vermaas attack and Thomasson defends is therefore the claim:

> (A): Artifacts are the intended products of largely successful intentions to create something of that kind.

While for Houkes and Vermaas the ascription of technical functions to artifacts plays a role in their model of artifact design and production, functions do not provide the 'essences' of artifacts. Rather their theory of function ascription is action oriented and highlights the capacities of artifacts relative to a plan, e.g. a use plan. It is therefore a relative and context dependent notion based on artifact actions.[4]

In their 'engineering' model of artifact creation, Houkes and Vermaas consider four agent types: *user*, *designer*, *maker* and *manufacturer*. A key element in the model is a *use plan*, understood as a series of goal-directed actions. Use is then characterised as the carrying out (by the user) of the actions that make up the use plan. The designer's role on the other hand is help users to realise their goals by constructing and communicating use plans.

To characterize a class of artifacts, three stages of *product-designing*, *making* and *manufacturing* are described in more detail. Behind product design is the idea that the designer describes a new item that contributes to the realisation of a use plan by prospective users of the artifact: the designer "d intends to contribute to producing items x_i, x_j, etc. that do not yet exist by *product-designing* them."

In simplified form, the product-design stage is described as follows ((Houkes & Vermaas 2009), p. 410):

- designer d believes that an item x with (physiochemical) capacity φ does not exist
- d contributes to realising a goal g_{dx} by describing item x (with capacity φ)
- d believes that a composite of components $c_1, \ldots, c_i, \ldots$ with capacities $\varphi_1, \ldots, \varphi_i, \ldots$ achieves the desired capacity φ.
- d believes that the various design tasks g_{dci} are simultaneously fulfilled.
- d intends to communicate a description of x (to appropriate agents).

To the extent that some of the components c_i may also need to be product-designed, the first four steps may iterated for each such c_i. This feature is built into the full description of the model.

The process of *making* a new artifactual kind is also characterised as a goal-directed series of actions. In (Houkes & Vermaas 2009) this is described as follows:

- maker m wants to bring about the existence of item x as described by agent a who, with that item x, wants to bring about the goal state g_a.
- m chooses or constructs a suitable make plan mp.
- m intends to carry out mp (and acts accordingly.

[4] See especially (Houkes & Vermaas 2004).

Again, this is a slight simplification that omits steps at which the maker verifies whether the plan mp has succeeded. This part of the model distinguishes the making of x from the agent (a) who with x wants to bring about the goal state g_a. For example, as Houkes and Vermaas note, agent a may be the product-designer, while m could be a different engineer involved in the production process.

A third type of agent is the manufacturer who supports all the processes involve in product creation and delivery. In the model it is supposed that the manufacturer designs the make plans that makers use:

- manufacturer mf wants to contribute to maker's goal of bringing about item x as described by agent a.
- mf constructs a suitable make plan mp involving items $y_1, y_2, \ldots$.
- mf contributes to producing $y_1, y_2, \ldots$.
- mf intends to communicate mp to users.

Again this is only a core part of the model, since (Houkes & Vermaas 2009) consider also processes whereby manufacturers verify and adjust the make plan in order to arrive at a successful version of the item x.

2.1 Summarising the basic features of the model

If we restrict attention to successful artifact creation, then what we have described here represents the core features of the Houkes and Vermaas model. Let us summarise the main components. This is a multi-agent account comprising several agent types with specific roles. There is an overarching goal, to design, produce and use an item x of an artifactual kind, say K, according to certain plans. To reach this goal there are various subgoals, each with corresponding plans. These may involve designing and producing the required subcomponents of x, $c_1, \ldots, c_i, \ldots$ assembled according to suitable make plans. The model involves agent beliefs, about capacities, properties and plans, as well as intentions, to act and carry out plans. There is also agent *communication* involved, since for example designers communicate plans to makers, manufacturers communicate plans to users, and so forth. We now need to examine whether and how these features may support an intentionalist understanding of artifacts.

3 A role for collective intentionality

Can we reconcile the view that artifacts are determined by a collective form of intentionality with the model of artifact creation proposed by Houkes and Vermaas? They themselves are sceptical about this. They take it to be a key feature (and improvement) that their model distinguishes between plan-designers, product designers, makers and manufacturers. However, even when these agents act deliberately and their "intentions are related to the characteristics of the produced item" (p.415), it is claimed that they may have "incomplete or conflicting notions of what is required to create a successful object of the kind." Their reluctance to accept that artifact kinds are collectively determined by all the agents involved in the design and production process seems to rest on the doubt that such kinds could ever be said to be collectively determined (by a group or

team) when members of the group display conflicting conceptions and intentions. This seems to be a primary reason for their rejection of thesis (A).[5]

Rather than constituting an argument against collective intentionality, it seems to me that individual differences in the conceptions held by the key agents in the production process may yield an argument in favour of introducing collective intentionality into the Houkes and Vermaas model. As an example, I will use elements of the approach to collective intentionality taken by Raimo Tuomela (especially (Tuomela 2007)), though other approaches could also be taken up in this context. For our present purposes it should suffice to pick out some of the essential features of each in order to show how the theory of artifacts is compatible with and can be extended by concepts from Tuomela's framework.

In Tuomela's account, the core concepts of collective intentionality are joint intentions and social actions, group attitudes and cooperation. Fundamental to this is his important distinction between the *I-mode* and the *we-mode* in forming intentions and in cooperating to achieve goals. It is the sharing of goals and cooperation in the we-mode that help distinguish a form of collective intentionality that is not reducible to individual intentions as represented in the I-mode. The Houkes/Vermaas model of artifacts incorporates action, agency and a plurality of agents, and so is precisely of the right kind for analysing within Tuomela's framework under the we-mode of intention and cooperation. As we have seen it is a multi-agent model of cooperation and there are explicit and common goals, communication between agents and even shared plans to achieve goals.

Here are some relevant features of Tuomela's model, CWM; the acronym standing for cooperation in the we-mode, given for the simplified case of two agents.

> A_1 and A_2 *(successfully) cooperate with each other in the we-mode* in bringing about a goal G if and only if
>
> 1. G is a collective goal type, namely an "achievement-whole" the achievement of which can be divided – either *ex ante actu* or *ex post actu* – into A_1's and A_2's parts;
> 2. A_1 and A_2 jointly intend to achieve G by acting jointly in the sense of (AT^*), and they achieve G jointly in accordance with and partly because of this joint intention of their to achieve G together. ((Tuomela 2007) p.165)

This is Tuomela's 'conceptually minimal' notion of a we-mode cooperation that is further embellished to a 'weakly rational' form of cooperation by requiring additionally that A_1 and A_2 rationally believe (1) and (2), and that (2) holds in part because of this shared belief.

This model of cooperation depends on an underlying property of *joint action*. This is understood according to a further condition (AT^*) as follows. X is jointly performed (by two agents 'you' and 'I') if

> 1. X is a collective action type, that is, an "achievement-whole" divisible – either *ex ante actu* or *ex post actu* – into your and my parts;

[5] For arguments sake, I will assume that team members might individually have different conceptions about the nature of the kind. In the above quotation, it seems that Houkes and Vermaas's claim is a weaker one, that there may be a conflict not over the kind itself but over the means needed to create it.

2. we jointly intended to perform X jointly;
3. we performed X jointly in accordance with and partly because of our joint intention to perform X (or some "closely related" action) jointly;
4. you and I mutually believed – or at least shared the belief – that 1,2, and 3;
5. 2 in part because of 4. ((Tuomela 2007), p. 112)

In short, that agents acted jointly in the sense of (AT^*) means that there was a joint intention to perform the action of a certain type X and that the performance of X was in accordance with a joint intention to perform X. Moreover there is mutual belief about this among the agents concerned.

Without going into all the subtleties of joint actions and collective goals as analysed by (Tuomela 2007), we can nevertheless make the following suggestions. First, suppose that a collective goal type G is the goal to produce an item x of an artifact type K and let us suppose that this goal succeeds and x is produced. The agents involved in the cooperation to achieve G may include the product-designers, makers and manufacturers, each with roles as described in the account of Houkes and Vermaas. They design, communicate and carry out plans. These agents cooperate in the we-mode according to CWM since they jointly intend to achieve G by acting jointly in the sense of Tuomela's joint action account; each agent a_i performing according to its appropriate role type. This cooperation involves the mutual belief that they so intend. The various subgoals, say $g_1, \ldots g_i, \ldots$ involving components $c_1, \ldots c_i, \ldots$ are similarly achieved via joint actions and cooperation in the we-mode.

Both models refer to what is eventually a successful form of cooperation, since ultimately an artifact of type K is produced. Having G as a collective goal type means in this case that there is a shared commitment to produce an artifact of type K and that the agents involved are thus acting as a group. By saying that 'the production of an artifact of type K' is a shared goal, and referring to 'K' in this statement, we are not begging the question about collective intentionality. We are supposing that an artifact is successfully produced, that this did not come about by accident, but by design, and that a group effort was involved in this. As we saw, the Houkes/Vermaas model refers to 'an item x' in several different places. So in the main goal (and its subgoals) the same item x is being described in each case. [6]

The Houkes and Vermaas model clearly conforms to part 1 of CWM since it divides the achievement of the goal into different agents' parts and moreover there is even a sharing of subgoals. We have just seen this in the examples of the maker's and manufacturer's intentions described earlier. Another example arises in the description of design plans found in (Houkes & Vermaas 2009) that we did not examine in detail. In clause D1 of their design plan one finds the condition: "The designer d wants to

[6] It seems clear, then, that within the Houkes/Vermaas model we can legitimately maintain that there is a common goal G even if there are differences of opinion about K (or item x) within the team of producers. But if necessary we could replace the expression 'artifact of type K' by an expression such as 'an artifact of the type later marketed under the label 'K' ' or some equivalent formulation.

contribute to a user's goal of bringing about a state g_u." ((Houkes & Vermaas 2009), p. 406.) So here, g_u forms part of a use plan and is a goal shared by both agents.[7]

What is the nature of the collective intentionality surrounding the artifactual kind K? The key to understanding this is to see that it rests on shared intentions to carry out joint plans according to an agreed assignment of roles. It does not rest on the assumption that agents have mutual beliefs about K or perhaps shared mental models of item x. This is a very fundamental point. Tuomela's cooperation model and the action performance submodel do not specifically refer to other attitudes like first-order beliefs about K. The mutual beliefs in the CWM model are higher-level beliefs about group cooperation, action types etc, not first-level beliefs about the nature of K. They refer to they fact that agents assume that other agents are acting in a cooperative manner and indeed this seems to be a clear feature of the Houkes and Vermaas model. In our case they do not need to share (first-order) beliefs about K.

This point is crucial for understanding (and obviating) Houkes and Vermaas's objections to the intentionalist view of artifacts. In (Houkes & Vermaas 2009) they do consider the question whether product designers, manufacturers and makers might *collectively* determine artifact kinds. Their main objection here is that agents may disagree on certain matters:

> ...the literature on engineering teamwork shows that conflicting conceptions and intentions among participants often survive intensive teamwork. ((Houkes & Vermaas 2009), p. 417).

Indeed, we might say that without disagreements most innovation would be stymied. In a corporate climate of 'yes-men' there would be no room for innovative ideas challenging current thinking to emerge bottom-up. This would be the antithesis of the approach of corporations such as Honda Motor Company that makes us of a type of meeting, known as *waigaya*, that aims to improve any aspect of the design and manufacturing (or sales and marketing) processes through a thorough and possibly lengthy discussion among different actors. In *waigaya* few or many employees from different departments and responsibilities meet to discuss a given issue or problem. Such meetings may be held regularly for days or even years until an agreed proposal or solution is found. In keeping with the 'flat' hierarchical structure at Honda, participants in a *waigaya* carry equal weight in their opinions. Ideas generated and results of the discussions become corporate property, and at the end of *waigaya* a set of corporate decisions and plans for enacting are generated.[8]

[7] Again, depending on the complexities of the example, presumably g_u could be the overall goal G or merely one of several subgoals.

[8] See for example (Rothfeder 2014) Ch. 3 for a lengthy discussion. According to Jeffrey Rothfeder, *waigaya* proceeds according to four basic rules:

1. Everybody is equal in *waigaya* – there are no bad ideas except those that are not aired.
2. All ideas must be disputed and rejected until they are either proven valid or vanquished.
3. When a person shares an idea, he or she doesn't own it anymore – it belongs to Honda and the group can do with it what it will.

What the example of *waigaya* illustrates is how from differing conceptions an agreed plan may emerge – by whatever discussion process that happened to take place – that the team then has to adhere to. It is instructive to note that ideas, once aired, are no longer 'owned' by the person who raised them.

While in Toumela's model (in the weakly rational form) the shared beliefs are in the form of assumptions about other agents' intentions to perform jointly, in the account of Houkes and Vermaas first-order beliefs do explicitly arise, e.g. at the design and make stages, and these beliefs are communicated between agents.[9] However, the model does not require that these are mutual beliefs held by all agents. They can be individual beliefs that are communicated or (in the important cases) *group* beliefs when they are taken up by a team as a whole. Group beliefs and other related group attitudes, like acceptance, also form a cornerstone of current work in the area of social intelligence. While there are differing accounts of these attitudes, a rather common and well established view is that *group* beliefs should not be understood to be beliefs that are necessarily held by every single member of a given collective, whereas *common* and *shared* beliefs are so regarded. In Tuomela's own approach, for example, group beliefs are understood to involve an intentional joint acceptance and commitment on the part of certain group members, but they need not be mutually held by all of the agents in the group. And this is very likely to be the case within a corporation and its operative teams.[10]

Therefore, in the case where the plans for cooperating and achieving the goal type rely necessarily on specific beliefs (about the nature of an artifactual kind K) i.e. on conceptions of K, it is sufficient if these are group beliefs, the important feature being that they are not necessarily held by each agent individually (though they may be communicated and acted upon). Another point to observe is that not all agents involved in a production process need be cooperating in the we-mode. While in modern forms of production, technicians and assembly-line workers are usually highly 'integrated' into the production process and aware of their role in the manufacturing the final product, we can allow for the case of an assembly-line worker who acts only in the I-mode in Tuomela's sense. We can leave room for this while holding that those agents who are essentially responsible for the final outcome are cooperating in the we-mode.

As quoted earlier, Houkes and Vermaas suggest that even conflicting intentions among participants "often survive intensive teamwork." In such a case team members would be cooperating at best in the I-mode (and at worst not at all). Yet this does not seem to be a plausible account of the type of case that we are dealing with here. The I-mode form of cooperation will hardly explain how a new artifact type could successfully be produced. In the I-mode, while cooperation between agents is possible, typically joint goals, collective actions and group attitudes are missing. This scarcely seems

---

4. At the end of *waigaya*, decisions and responsibilities are generated – a precise list of who is to do what next and by when. ((Rothfeder 2014), pp.69-70).

[9] For instance they may be beliefs about x and its properties, or about components of x, or the correctness of their make plans.

[10] For Tuomela's account see for instance (Tuomela 1992, Tuomela 2007). For a recent overview of approaches to group belief and a reconstruction in modal logic, see (Gaudou et al 2014).

adequate to explain how the outcomes of modern design and production processes can be so effective.

It might seem at first counterintuitive to claim that the nature of an artifactual kind K depends on a form of collective intentionality, while the actors involved may hold differing *conceptions* of K. This might indeed appear strange if one is thinking that the 'essence' of K is collectively determined by some mutual idea, a shared mental model or perhaps an agreed list of technical functions and defining properties. But this is not the case under the multi-agent, action-theoretic model of artifact creation. In this model the 'nature' of K depends on the plans that the agents enact under (I claim) a we-intentional form of cooperation. Those beliefs that support different phases of artifact creation, inasmuch as they are beliefs about the nature of K, should be regarded as group beliefs to which core members of the team are committed *as a group* in order to pursue the overarching goal. If this is correct, then production of an artifact in the manner envisaged by Houkes and Vermaas can illustrate an elaborate form of cooperation in the we-mode, with common goals and group commitments, and therefore exhibit exemplary features of collective intentionality. If those essentially involved in the design and production process act as a team, then their group commitments, attitudes and goals are what help to ensure the successful outcome. Moreover, the notion of we-mode and the appeal to collective intentionality and group attitudes helps to explain how a coherent artifactual type K may be successfully established even if different beliefs and conceptions about K are held by individual team members.

4 Conclusions

To summarise, we have seen how Houkes and Vermaas reject Thomasson's thesis of intentionality and artifact determination largely due to a disagreement about what agents constitute the 'maker' of an artifact and the fact that different agents involved in the process, from invention to production, may have conflicting beliefs and conceptions related to the artifactual kind. They also doubt whether the role of intentionality can be saved by claiming that the actors involved *collectively* determine artifact kinds. This scepticism seems unwarranted if we apply concepts of collective intentionality in the manner discussed by Tuomela and other authors. Despite differences of opinion among agents, it seems that modern production processes rely on a teamwork that involves shared goals and cooperation in the we-mode. These features are what capture collective intentionality and, as I have tried to show, they are fully compatible with the Houkes and Vermaas model even when not explicitly acknowledged there. We should also not forget the potential role of users and consumers in the intentionality of artifacts. This, for Houkes and Vermaas, should not be a point of contention. In their model users are represented not only indirectly by the use plans but also directly by the fact that the agent a referred to in the model may also be a innovative user of the technology or even just a potential user ((Houkes & Vermaas 2004), p.66).

References

Gaudou, B, Herzig, A, Longin, D, & Lorini, E. On modal logics of group belief, forthcoming, 2014.

Hilpinen, R., On Artifacts and Works of Art, *Theoria* 58 (1992), 58–82.
Hilpinen, R, "Artifact", The Stanford Encyclopedia of Philosophy (Winter 2011 Edition), Edward N. Zalta (ed.), URL = http://plato.stanford.edu/archives/win2011/entries/artifact/.
Houkes, W. & Vermaas, P., Actions versus Functions: A Plea for an alternative Metaphysics of Artefacts, *The Monist* 87 (2004), 52–71.
Houkes, W. & Vermaas, P., Contemporary Engineering and the Metaphysics of Artefacts: Beyond the Artisan Model, *The Monist* 92 (2009), 403–419.
Houkes, W., Vermaas, P., Dorst, K., & de Vries, M., Design and use as plans: an action-theoretical account, *Design Studies* Volume 23:3 (2002), 303–320.
Ingarden, R, *Untersuchungen zur Ontologie der Kunst*, Max Niemeyer Verlag, 1961.
Margolis, E. & Laurence. S., *Creations of the Mind. Theories of Artifacts and their Representation*, Oxford University Press, 2007.
Rothfeder, J. *Driving Honda. Inside the world's most innovative car company*, Portfolio Penguin, 2014.
Searle, J., *The Construction of Social Reality*, Allen Lane, The Penguin Press, 1995.
Searle, J., *Making the Social World. The Structure of Human Civilization*, OUP, 2010.
Sellars, W, *Science and Metaphysics*, Routledge and Kegan Paul, 1968.
Thomasson, A, Realism and Human Kinds, *Philosophy and Phenomenological Research* 67 (2003), 580–609.
Thomasson, A, Artifacts and Human Concepts, in (Margolis & Laurence 2007), 52–73.
Tuomela, R, *A Theory of Social Action*, Reidel, 1984.
Tuomela, R, We-intentions and Social Action, *Analyse & Kritik* 7 (1985), 26–43.
Tuomela, R, Group Beliefs, *Synthese* 91 (1992), 285–318.
Tuomela, R, *The Philosophy of Sociality*, Oxford University Press, 2007.