What Makes a Kind an *Art*-kind?

Michel-Antoine Xhignesse

michelxhignesse@capilanou.ca

Abstract: The premise that every work belongs to an art-kind has recently inspired a kind-centred approach to theories of art. Kind-centred analyses posit that we should abandon the project of giving a general theory of art and focus instead on giving theories of the *arts*. The main difficulty, however, is to explain what makes a given kind an *art*-kind in the first place. Kind-centred theorists have passed this buck on to appreciative practices, but this move proves unsatisfactory. I argue that the root of this dissatisfaction stems not from the act of kicking the can down the road, but from not kicking it far enough. The missing ingredient, I argue, is a notion of convention which does the work of marking the difference between art and non-art for a given physical medium.

1. Introduction

Let us begin with a single premise:

(P): Every artwork belongs to an art-kind.

Although we might quibble about the full extent of the claim—e.g. does each work belong to just one art-kind? How many art-kinds are there, what are their boundaries, and how should we class mixed-media works?—it is nevertheless largely uncontroversial. After all, we make daily use of our kind-categorizations: this is a novel, that's a TV show, this is a film, that's a painting, etc. We are surrounded by entities which fit neatly into our classificatory schema. In recent years a great deal of excitement has attended this premise, which has been used to generate kind-centred analyses of 'art' which do not rely on identifying properties common to all artworks. Where ontology is concerned, this premise has led David Clowney to argue that art-making must proceed deliberately—that is to say, making art requires an agent to intend to instantiate an entity belonging to an art-kind.[1] On the definitional front, the premise has inspired both top-down and bottom-up approaches. Top-down theorists like Christy Mag Uidhir have argued that art-kind membership is determined by our theories of art, so that in order for an entity to be an artwork, it must first satisfy the requirements of art-kind membership, and second, the way in which it satisfies those requirements must also satisfy the conditions for being an artwork, whatever those may be.[2] On the other hand, bottom-up theorists such as Dominic McIver Lopes have used it to argue that we ought to abandon the project of giving a theory of art, and focus our attention instead on the more promising project of giving theories of the *arts*, since we have a better understanding of the limits and nature of individual art-kinds.[3]

---

[1] David Clowney, 'Definitions of Art and Fine Art's Historical Origins', *Journal of Aesthetics and Art Criticism* 69.3 (2011), 309-20.

[2] Christy Mag Uidhir, *Art & Art-Attempts* (New York: OUP, 2013).

[3] Dominic McIver Lopes, 'Nobody Needs a Theory of Art', *The Journal of Philosophy* 105.3 (2008), 109-27; and *Beyond Art* (New York: OUP, 2014).

This paper takes up the idea that we can analyse 'art' in terms of its constituent art-kinds. Although the project of focusing philosophical attention on art-kinds is at least *prima facie* promising, I will argue that it does not manage to pass the buck back far enough, and that its deficiencies become evident when we ask what makes a kind an *art*-kind. Ultimately, I argue that buck-passing bottoms out in a buck-*stopping* analysis of art-kinds in terms of networks of conventions. This suggests a proceduralist explanation of how our kind-categorizations take place, and how they are taught and transmitted, which does not rely on agents' direct intentions or propositional attitudes. This framework also offers us a method in principle for determining the factors that inform those kind-categorizations in the first place, although the paucity of historical information makes it incredibly difficult to do so in practice for all but the most recent developments in our artistic practices.

2. Centring kinds on medium

The idea that we should focus our theoretical attention on giving theories of art-kinds rather than of art in general was first proposed systematically by Dominic McIver Lopes, who argues for the following buck-passing principle (R):

> (R): item x is a work of art if and only if x is a work in activity P and P is one of the arts.[4]

'The arts' here just designates the various art-kinds: media such as dance, literature, painting, music, sculpture, etc. On this model, something is a work of art iff it is classified as such by a theory of some art-kind. So, nobody actually needs a theory of art—all we need are theories of art-kinds. A theory of sculpture, for instance, should suffice to tell us why a particular sculpture is art—the notion of 'art' in general has nothing more to contribute. The kind-centred approach thus passes the explanatory buck on to theories of the individual arts.

---

[4] Lopes, 'Nobody', 109.

But, as Beardsley pointed out, the proper appreciation of a dance requires more than just the observation of a series of bodily motions; it also requires the realization that their presentation to an audience is part of a larger enterprise with particular representational or expressive goals (among other possibilities).[5] Our proper interest is thus not in the bodily motions considered in themselves (as the work's vehicular medium), but rather considered as 'movings' and 'posings'.[6] This serves to highlight an important point: art-kinds and physical matter come apart, so that not every entity composed of the same kind of matter is an artwork in the same kind. As Lopes himself observes, the fact that painting is one of the arts does not mean that every painted thing—nor even every painted canvas—is an artwork.[7] The kind-theorist owes us an explanation of what makes some of the entities belonging to a kind artworks, while others—which may even be perceptually indistinguishable from them—are not.

So, for example, a theory of the arts ought to be able to explain why a piece of bizen-yaki is art, while an ordinary coffee mug is not. Lopes calls this the *coffee mug objection*, and its force comes from the fact that in ordinary discourse, talk of an entity's 'medium' carries the tacit assumption that the entity in question is an artwork.[8] We should, of course, be wary of accepting such a contentious assumption: not every painted surface is an artwork, as any bridge, fence, or house will readily testify. Any theory of art-kinds will have to offer us principled reasons to refrain from affixing the label of 'art' to all the entities falling under a vehicular medium.

In an effort to respond to the coffee mug objection, Lopes argues that art-kinds are fundamentally *appreciative* kinds, which is just to say that they are the products of appreciative *practices*.[9] The exact nature of each of the arts thus depends in part on the value a particular

---

[5] Monroe C. Beardsley, 'What Is Going on in a Dance?' *Dance Research Journal* 15.1 (1982), 31-6.
[6] Beardsley 'What Is Going On', 34-5.
[7] Lopes, 'Nobody', 124.
[8] Lopes, *Beyond Art*, 17.
[9] Ibid, 121 and 148.

community attaches to each of the members of the kind in question.[10] A 'medium,' on Lopes's view, is a technical resource: an inert array of variegated matter ('resource'), completed by a set of procedures for its transformation ('technique').[11] Artistic media—what Lopes calls 'appreciative kinds'—are not individuated by a single vehicular medium, but by what he calls a 'medium profile,' a nested set of media joined together by a practice exploiting a core set of technical resources in order to realize artistic properties and values.[12] To say that two entities are works in the same art-kind is just to say that they are both part of the same appreciative practice.[13] The suggestion, then, is that we should look to our social practices for the answer to the coffee mug objection. It is the social practices from which bizen-yaki results that secure its belonging to the art-kind of ceramics and, thus, its art-status, just as it is the social practices from which a coffee mug results which preclude it from being art.[14]

The second major problem facing an analysis of the arts is posed by what Lopes calls 'free agents', entities that are clearly artworks but which do not seem to belong to any particular pre-existing art-kind or appreciative practice.[15] Consider Tracey Emin's *Everyone I Have Ever Slept With 1963–1995* (1995), a tent with the appliquéd names of 102 people with whom the artist literally shared a bed in the designated period. Works such as *Everyone* seem to defy the strategy of explaining 'art' by reference to theories of the arts because they have no obvious associated art-kind. This is not to say that free agents have a wholly bare medium profile—many will use some physical substrate as

---

[10] Ibid, 130.
[11] Ibid, 139.
[12] See Lopes, *Beyond Art*, 140, 144, and 182. The introduction of the notion of a 'medium profile' is meant to forestall the thought that Lopes is re-introducing the doctrine of medium-specificity, and to accommodate the existence of mixed- and multi-media works. A work's medium-profile does not (or, at least, need not) correspond to a single physical medium. See Lopes, *Beyond Art*, 195; for a similar argument against singularity of medium, see Noël Carroll, 'Forget the Medium!", *Engaging the Moving Image* (New Haven: Yale UP, 2003), 1-9, at 5-6. The notion of a medium profile conveniently sidesteps the problem of deciding how fine-grained we need to be about medium (see Ibid, 7).
[13] Lopes, *Beyond Art*, 182.
[14] Ibid, 148.
[15] Lopes, *Beyond Art*, 17 and 196-8.

a vehicle for their artistic significance—but rather that they do not fit into any established appreciative practices.

One strategy for dealing with free agents is to relegate the entity in question to an existing art-kind, such as sculpture. *Everyone* makes for an uneasy sculpture, however, since its vehicular medium does less artistic work than the idea conveyed by the title. Similar problems attend its classification under other art-kinds; free agents are difficult to classify precisely because they so often deliberately set out to challenge existing artworld classifications.

The other option, according to Lopes, is to assign the work to a 'new' art-kind such as conceptual art, whose medium profile is characterized by language or sets of ideas (especially about art itself). There is no doubt that 'conceptual art' is a legitimate art-kind; its appreciative practices concretized in the late 1970s, after decades of accumulated work and theory by Isidore Isou, Henry Flynt, Sol LeWitt, and Joseph Kosuth, among many others. But it did not spring from the cultural world fully-formed and all at once. Duchamp's readymades (c. 1917), for instance, are widely credited with paving the way for conceptual art by showing how to separate a work's artistic merit from its vehicular medium. Duchamp's work, in turn, was heavily influenced by the increasing abstractions of cubism and expressionism (especially via Kandinsky), which themselves had their roots in aestheticism, impressionism, and post-impressionism. In fact, we can see early gestures towards separating a work's artistic merit from its vehicular medium in late nineteenth century titling practices, which began to exploit violations of Gricean maxims in order to focus an audience's critical attention on particular features of the work in an unusual way.[16] Just consider Whistler's whimsically titled paintings, especially his *Arrangement in Grey and Black No. 1* (1871), which is more commonly known as the *Portrait of Whistler's Mother*. The work's artist-given title suggests that its proper subject is not the woman depicted (and her relation to the artist), but rather the formal

---

[16] Michel-Antoine Xhignesse, 'Entitled Art: What Makes Titles Names?', *Australasian Journal of Philosophy* 97 (2019), 437–450 at 442.

properties that she embodies. These considerations suggest a different question that we should put to kind-centred analyses: if at least some of these precursors are plausibly interpreted as works of conceptual art, then were they free agents until the 1960s and '70s, when conceptual art was conceived?

Certainly not. Lopes quite rightly observes that the best interpretation of such putatively free agents is as pioneering works: their existence calls for a theory of the art-kind which they pioneer, and that can only develop over time as artists follow in the pioneer's footsteps and audiences begin to accept new such works because they have already accepted the last.[17] In other words, they call for the development of conventions.

Pioneering free agents, when accepted as artworks, are likely to be accepted as ill-fitting members of a pre-existing kind, rather than being recognized as the progenitors of a new art-kind. This was the case with *Fountain* (1917), which is routinely called a sculpture[18]—Duchamp himself referred to it as a sculpture in a letter to his sister in 1917.[19] We have no trouble recognizing these as artworks, given the right contextual cues; what gives us trouble, rather, is their classification under an existing art-kind. And herein lies the trouble for kind-based accounts of 'art': these works usher in the dawn of a *K*-centred appreciative practice but, because that practice is not yet established, they cannot (yet) be artworks *in virtue of* belonging to it. Instead, they must be art in virtue of belonging, however poorly, to some other pre-existing art-kind (e.g., sculpture). And yet, on the kind-centred approach it is their status as members of that same established art-kind which we find so

---

[17] Lopes, *Beyond Art*, 202.

[18] No less illustrious a figure than Thierry de Duve does so in his commentary on *Fountain* where, in the course of arguing for its status as "pure" art, he says that "a urinal is a sculpture only when you accept seeing it as art"— see his *Kant after Duchamp* (Cambridge: MIT Press, 1996), at 153.

[19] In a letter to his sister Suzanne dated 11 April 1917, Duchamp says "Une de mes amies sous un pseudonyme masculin, Richard Mutt, avait envoyé une pissotière en porcelaine comme sculpture." See Marcel Duchamp, *Affectionately, Marcel: The Selected Correspondence of Marcel Duchamp*, Ed. Francis M. Naumann and Hector Obalk (Ghent: Ludion Press, 2000), at 47.

contentious in the first place. In other words, pioneering free agents show us that the kind-centred approach flirts with circularity.

The danger lies in focusing on the entities in question rather than on the actions that generate them, the practices that guide these actions, and the cultural contexts in which the entities find themselves. The object itself does virtually none of the substantive work involved in its classification under an art-kind; that work is done by the conventions governing our artistic and appreciative practices. The mere fact that some Xs are artworks does not suffice to make of X an art-kind; the missing ingredient is an appreciative practice centered on Xs and X-like works. When we ask whether a kind K is an art-kind, our interest is not in whether a particular entity E_1 falling under K is an artwork. Rather, our interest is in whether entities $\{E_1, E_2... E_n\} \in K$ are standardly taken to be artworks in virtue of belonging to K. The only way to answer this question, however, is to discover what is standard for a given community—in other words, it requires us to focus on our appreciative practices and the conventions underpinning them, rather than on the entities themselves. This is not to say that looking to our practices will supply a definite answer; our practices are contingent affairs with often-arbitrary origins and may not yet have concretized to the point where there exists a strong norm compelling us to automatically classify K as an art-kind. Though there may be significant precedents, their accumulated weight may not yet suffice to delimit a fully-operational appreciative practice.

The final objection to kind-centred analyses is what I call the *structural objection*. Kind-centred analyses are predicated on the wager that the prospects of giving theories of art-kinds are brighter than those of giving a theory of art in general. While it is certainly true that philosophers have had mixed success with developing theories of art, unfortunately the same is also true of philosophical attempts to analyse the individual arts. The result is that the prospects for kind-based theories are no better than they are for theories of art more generally. To take what is perhaps the best-developed

case in the philosophical literature, consider the notion of a 'musical work'. Is it definable in terms of necessary and sufficient conditions,[20] or not[21]? Is it abstract[22] or concrete[23]? Repeatable,[24] or singular[25]? A mental entity,[26] a historical individual,[27] an action,[28] or some kind of Platonic abstract entity[29]? The literature on the ontology of music alone is every bit as extensive and fractious as that on the ontology of art in general. This is not meant to be a decisive objection. The point, rather, is that the motivation behind the kind-centred approach does not survive scrutiny: there is no more agreement over the identity and persistence conditions of art-kinds than there is over those of art in general.

The lesson here is not that the kind-centred approach is wrongheaded, or that its responses to the coffee mug, free agent, and structural objections are wrong. On the contrary, I think there is much to like about buck-passing theories of art; in particular, I think that there is a great deal we can learn from the way in which the buck gets passed. My point, rather, has been to establish that the buck-passer's responses to these objections are unsatisfactory—not because she kicks the can down the road, but because *she does not kick it nearly far enough*.

---

[20] Jerrold Levinson, 'The Concept of Music', *Music, Art, and Metaphysics* (Ithaca: Cornell UP, 1990), 267-78 and Andrew Kania 'Definition', *The Routledge Companion to Philosophy of Music*, ed. Theodore Gracyk and Andrew Kania (New York: Routledge, 2011).

[21] Roger Scruton, *The Aesthetics of Music* (Oxford: OUP, 1997) and Andy Hamilton, *Aesthetics and Music* (New York: Continuum, 2007).

[22] Richard Wollheim, *Art and Its Objects* (Cambridge: CUP, 1980 [1968]); Nicholas Wolterstorff, *Works and Worlds of Art* (Oxford: Clarendon Press, 1980 [1968]); Levinson, 'The Concept' (1990); and Julian Dodd Julian Dodd, *Works of Music: An Essay in Ontology* (Oxford: OUP, 2007).

[23] Mag Uidhir (2013), 123-97; Nelson Goodman, *Languages of art: an approach to a theory of symbols* (Indianapolis: Bobbs-Merrill, 1968); and Ben Caplan and Carl Matheson, 'Defending Musical Perdurantism', *British Journal of Aesthetics* 46.1 (2006), 59-69.

[24] Dodd, *Works of Music* (2007).

[25] Allan Hazlett, 'Against Repeatable Artworks', *Art and Abstract Objects* ed. Christy Mag Uidhir (Oxford: OUP, 2012), 161-78.

[26] Diana Raffman, *Language, Music, and Mind* (Boston: MIT Press/Bradford, 1992) and R.G. Collingwood, *The Principles of Art* (Oxford: Clarendon Press, 1938).

[27] Guy Rohrbaugh, 'Artworks as historical individuals: lessons from photography', *European Journal of Philosophy* 11.2 (2003), 177-205.

[28] David Davies, *Art as Performance* ((Oxford: Blackwell, 2004).

[29] Peter Kivy, 'Platonism in Music', *Grazer Philosophische Studien* 19 (1983), 109-29; Dodd, *Works of Music* (2007); Jerrold Levinson, 'What a Musical Work Is', *Journal of Philosophy* 77.1 (1980), 5-28; and Stephen Davies, *Musical Works and Performances: A Philosophical Exploration* (Oxford: OUP 2001).

3. Conventions

Lurking behind the buck-passer's game of hot potato is a buck-stopping phenomenon capable of explaining what makes some kinds artistic: a network of cultural conventions. Before we can boil that particular kettle of fish, however, I must clarify what I mean by 'convention'.

On the traditional—which is to say, Lewisean—model, conventions are a means of solving *coordination problems*. A coordination problem is just a situation in which two or more agents are faced with a choice of alternative actions, and where an agent's choice of a course of action depends upon the actions of the other agent(s).[30] Agents have no means of communicating in advance to coordinate their actions and, until conventions develop, they have no recourse to pre-established norms or rules. So, for example, the choice of a lane in which to drive requires some kind of coordination, lest drivers collide due to their conflicting choices; the desirable outcome is to find oneself driving in the same lane as all or most other drivers going in the same direction.

An *equilibrium* obtains when no one would have been better off had they acted otherwise while everyone else acted in the same way.[31] In the absence of established laws concerning which lanes should be used, drivers must find some way to coordinate their actions. This will require them to consider not just their own preferences, but those of their fellow agents, as well as what they may reasonably expect those agents to expect from them, and so on. So, if everyone else going my way is driving on the left, I can reasonably expect other drivers driving my way to drive in the left lane, they can reasonably expect drivers other than themselves who are going their way to drive in the left lane, and they can reasonably expect that drivers other than themselves will expect other drivers driving their way to drive in the left lane, etc. I thus have good reason to drive in the left lane myself, since doing so is the best way to avoid a collision. The right lane may seem faster—perhaps it looks empty!—but so long as everyone else going my way drives (or expects to drive) in the left lane, I

---

[30] David K. Lewis, *Convention: A Philosophical Study* (Malden: Blackwell, 2002 [1969]), at 8.
[31] Ibid, 15.

would not be any better off by taking the right lane since, unless it really is empty, I would encounter another vehicle. We thus have a state of coordination equilibrium.

So, on the Lewisean model, conventions develop from agents' self-perpetuating solutions to recurring coordination problems. Relying on the notion of *salience* offers one way to resolve recurring coordination problems: one particular equilibrium stands out from the others for some particular reason, e.g., because it's uniquely good, because agents can reasonably expect one another to notice this particular equilibrium's uniqueness, by virtue of its intrinsic properties, etc. Another way to resolve these problems is by appealing to the force of *precedent*: a particular state of equilibrium is conspicuously salient to agents because it was the equilibrium reached the last time they encountered this particular coordination problem, or the last several times.[32]

Interesting as it is, Lewis's analysis of convention falls a little short when it comes to art-kinds and artistic practices, since many of these do not plausibly have their origins in coordination problems.[33] That is to say, it would be implausible to maintain that the reason human figures in Italian Mannerist paintings have elongated limbs is rooted in the fact that sixteenth-century painters faced a choice of alternative actions with multiple Nash equilibria. Rather, the standard art-historical explanation is that they were rebelling against Classical conventions of form and established new conventions of their own, featuring cryptic body language, exaggerated poses, and elongated proportions.[34] Nor did Classical conventions of depiction have anything to do with coordination problems. Similarly, Wayne Davis has observed that the conventions governing fashion do not seem reducible to coordination problems, either.[35] The convention that dresses, gowns, and skirts are for women and trousers are for men, for example, does not appear to solve any kind of historical

---

[32] Ibid, 35-6.

[33] cf. Lopes, *Beyond Art* (2014), 148, who says that conventions are explanatorily unsatisfactory since artistic practices may not solve coordination problems.

[34] Marilyn Stokstad, *Art History* (New York: Harry N. Abrams, Inc., 1995), at 712-19. The language of 'conventions' here comes from Stokstad herself.

[35] Wayne Davis, *Meaning, Expression, and Thought* (New York: CUP, 2003), at 225.

coordination problem, even if it solves other kinds of problem. So not all of the behaviours we class as 'conventional' are reducible to coordination problems; this means that the Lewisean analysis of convention does not get at the root of the phenomenon.

To be sure, many philosophers and theorists have attempted to extend Lewis's analysis of convention to social world conventions. Howard S. Becker, for instance, has argued that we need to appeal to Lewisean conventions to explain the artworld's operations and organization.[36] Becker is especially concerned to explain collective action and cooperative activity in the artworld, such as that between artist, audience, and artwork, between musicians tuning their instruments, as it occurs in the act of interpretation, or in making decisions about beats, blocking, and focus in theatre. Crucially, he argues that conventions "allow people who have little or no formal acquaintance with or training in the art to participate as audience members,"[37] thanks in part to their socialization, or to their knowledge of the history of similar genres, media, works, characteristic styles and periods, artworld gossip, etc.[38] While I think Becker is largely correct, his arguments are premised on the routine conflation of the technical notion of a coordination problem with cooperation more broadly, so that many of his examples of 'conventions' showcase no coordination problems whatsoever.

Happily, another analysis of convention is available, although it has long languished in obscurity: Ruth Garrett Millikan's 'natural' conventions. On this model, conventions are patterns of behaviour that reproduce largely due to the weight of precedent, meaning that they need not be responses to coordination problems, though they can be.[39] By 'reproduction' of a pattern of behaviour P_1, Millikan just means the generation of a new pattern P_2 with roughly the same form as P_1. The relationship between P_1 and P_2 is governed by counterfactual dependence: if the model (P_1)

[36] Howard S. Becker, Art Worlds: 25th Anniversary Edition, Updated and Expanded (Berkeley: University of California Press, 2008 [1984]) 40-67.
[37] Becker, *Art Worlds*, 46.
[38] Becker, *Art Worlds*, 48
[39] Ruth Garrett Millikan, *Language, Thought, and Other Biological Categories: New Foundations for Realism* (Cambridge: MIT Press, 2001 [1984]), at 10.

had differed in certain respects, so too would the copy (P$_2$) have done. A model does not determine all of a copy's parts, but it should determine those which are functionally salient. Only those tokens reproduced due to precedent will count as conventional, and the answer to the question of which convention each token instantiates will depend on the precedent from which it was derived.[40] The Lewisean analysis of convention presupposes that the participants are rational, and requires that they have a fairly sophisticated network of higher- and lower-order beliefs. Millikan's model, by contrast, requires no beliefs or rationality whatsoever. Even when it comes to the subclass of natural conventions that includes coordination problems, no rationality is required: "coordination conventions proliferate," Millikan argues, "because, rationality aside, people learn from experience exactly as other animals do".[41]

Millikanian conventions arise when we select actions over and over again due to the weight of precedent. Mere regularity of solutions to coordination problems is not sufficient to generate a convention, however. This is just because regularities of outcome which obtain from courses of action developed independently of previous outcomes are merely accidentally regular.[42] Coordination *conventions* arise when precedent increases the salience of a solution in such a way that agents choose the pattern of behaviour over and over because it is one which was chosen in the past. Specific patterns of behaviour might stand out as especially salient for a number of reasons, including their ties to the human perceptual apparatus, complementarity of function and ease of processing fluency, piggy-backing on evolutionarily-selected capacities and traits, reinforcement of or by existing practices and institutions, or for purely arbitrary factors (in which case the resulting categorization is effectively random).

---

[40] Ibid 3, 17.
[41] Ibid, 10.
[42] Ibid 3-4.

Conventions are thus fairly arbitrary creatures:[43] different historical accidents would have yielded different conventions, meaning in turn that different populations are likely to have different conventions (so long as we adjust for cultural transmission). So, for instance, virtually all human cultures which have been studied have objects and practices that employ the same kinds of vehicular media as those which we, in our cultural context, classify as 'art.' Closer inspection, however, reveals that putative art-kinds do not match up neatly across cultures. The Balinese, for instance, appear to have music, since they use tools to produce sounds. But their musical tradition places much more emphasis on its devotional and participatory elements than ours does. Although we can easily recognize the similarities between the vehicular medium of Balinese 'music' and our own, the fact remains that the network of conventions governing the Balinese practice is relatively alien to Western audiences who are unacquainted with Balinese sonic and religious practices. The similarity of vehicular media does not necessarily translate to a similarity of the appreciative practices governing the creation and use of a given vehicular medium.[44]

What is more, even if two cultures share traditions with a similar medium profile and appreciative practice, that is no guarantee that it is an *artistic* practice in both cultures.[45] The influence of conventions thus motivates the conclusion, *pace* Kendall Walton, that art-kinds are not in principle pellucid to perception alone.[46] We can easily imagine the development, through the accretion of art-historical precedent, of an appreciative practice whose artworks' appreciative kind is not discernible based solely on the sense-impressions they generate. Just think of Walton's infamous

---

[43] Such that a convention's function, if it has one, might just as well have been satisfied by some other pattern of behaviour. See Ruth Garrett Millikan, 'Deflating Socially Constructed Objects: What Thoughts Do to the World', *Perspectives on Social Ontology and Social Cognition*, edited by Mattia Gallotti and John Michael, (New York: Springer, 2014, 27-39).

[44] None of this is to say that the Balinese do not have music, or that their 'music' is not art. The point, rather, is just that these categorizations must rest upon a closer inspection of and reflection upon the relevant underlying conventions and practices. See Stephen Davies, 'Balinese Aesthetics', *Journal of Aesthetics and Art Criticism* 69.1 (2007), 21-9.

[45] See David Davies, 'Dominic McIver Lopes: Beyond Art', *Notre Dame Philosophical Reviews* (2015).

[46] c.f. Kendall Walton, 'Categories of Art', *The Philosophical Review* 79 (1970), 334-67, at 337-42. Walton grants that some categories of art (such as literature) might not be obvious to perception alone, but he thinks that most will be.

guernicas, works whose painted surfaces, when viewed from the right angles, look like exact copies

of Picasso's *Guernica*, but whose bas-relief dimensions are crucial to their proper appreciation.[47] Even

though the proper way to appreciate guernicas is through one's senses, doing so is not enough to

individuate flatter ones *as guernicas* rather than as paintings. Lopes himself admits as much when he

argues, *contra* Danto, that we are *not* all Testaduras now, because we have no trouble recognizing

individual works as art. We stumble, he thinks, when it comes to categorizing the hard cases as

instances of a *particular* art-kind (e.g. *4'33"* as *music*, *Fountain* as *sculpture*, etc.).[48]

It is worth noting, at this juncture, that this account of *convention* is, in many respects, similar

to the account of social *norms* given by Cristina Bicchieri.[49] Bicchieri argues that social norms are

predicated on imitation, itself grounded in one's behavioural expectations of others, as well as one's

perception of others' behavioural expectations. There is thus a key difference between Millikan's

minimalist account of convention and Bicchieri's: on the Millikanian picture, imitation is just one of

several ways of generating the regularity of behaviour that results in conventions, and imitation need

not be rational or propositional in nature. It is also worth noting that although Bicchieri's

conventions are all the result of coordination problems, she does not think that all social norms

result from conventions. Ultimately, the question of which is the bedrock concept, convention or

normativity, is a chicken-and-egg-herring which need not occupy us here; my point is simply that

kind-centred theories of art should pass the explanatory buck on to an account of social conventions

(or norms, if you prefer).

Now, none of this is to say that conventions are not sometimes manifested (and thereby

reinforced) deliberately, that is to say, that agents never consciously decide to instantiate existing

conventions. This happens when, for example, Hastings and Poirot play a game of chess, when

---

[47] Walton, 'Categories', (1970).
[48] Lopes, *Beyond Art*, 71-4.
[49] See *The Grammar of Society: The Nature and Dynamics of Social Norms* (New York: CUP, 2006), esp. 8-41.

Parliament meets, or when we wear what we wear because it's fashionable. In these cases, participants follow explicit sets of (procedural) rules which describe the conventions in question. In following these rules, participants reproduce the relevant conventions.[50] But conventions can also be manifested in the absence of any deliberate intention to that effect. Sometimes, they are instantiated just because it is easier or more natural to instantiate an extant pattern of behaviour than it is to develop or instantiate an equally viable alternative. This is the case when, for example, Westerners use a fork at the table rather than chopsticks, a decision that is usually made without much forethought. The choice between fork and chopsticks certainly solves a problem, but not a *coordination* problem.

Perhaps most often, however, conventions are manifested by their easy familiarity, because they are 'what has always been done' or because they are simply what everyone else expects: the precedent is so well established that no other options present themselves to the agent as salient (a fact reflected in their status as *ex post facto* rationalizations). Consider the widespread cultural practice of a woman taking her husband's surname upon marrying him. There is no good, ultimately non-arbitrary reason for the practice to persist at a time when women are considered autonomous agents, and yet it remains the traditional norm in much of the Anglophone world, where it is adopted largely unreflectively. Elsewhere, historical accidents have counted against this convention. Since the French Revolution, for example, French law has stipulated that individuals must keep the names on their birth certificates, resulting in a new, contrary convention. The older convention was similarly undermined in many French colonies, including Québec, where women wanting to adhere to it must go through the same formal procedure as those changing their names for other reasons—and Québec's registrar of civil status does not recognize 'marriage' as a legitimate reason for name change! In Spain and Korea, by contrast, the custom has been for women to keep their original

---

[50] Millikan, *Language: A Biological Model* (Oxford: OUP, 2005), at 14-15.

surnames. We can see, then, that some conventions are negatively reinforced, as with some of the laws surrounding name changes. Others are positively reinforced by a society's laws, as with the conventions which govern marriage: so, for example, until 2005 in Canada, 2013-14 in most of the UK, and 2015 in the USA, 'marriage' could only obtain between two adults of different sexes. But conventions are not immutable; they are frequently broken (especially when several parties are involved), and they can change over time as different precedents accrue more and more social heft, or as their reinforcing institutions change or are dismantled.[51] The particular character of a convention depends upon the use to which we put it, on the aims of the relevant cultural practice.

4. How conventions beget artworlds

Conventions manifest themselves in all sorts of ways in artistic contexts. They can influence our interpretation and evaluation of artworks: in literature and film, for example, genre has a profound effect on what we accept as true in a story, as well as on how we judge the work in question. The transparency of symbols is likewise founded upon convention, so that representations of dogs in wedding portraits typically stand in for fidelity, while upside-down national flags can symbolize distress, a state of war, or act as a symbol of protest. Conventions can even affect how and what we *see* in a picture. To take John Dilworth's example, engraving frequently uses the technique of cross-hatching to indicate tone and shading in monochromatic prints, but nobody familiar with the art-kind mistakes the cross-hatching for a layer of mesh covering the scene depicted.[52] But conventions can also manifest as more complex collections of attitudes, beliefs, dispositions, and practices; let us call such a collection a *world*. Allow me to explain.

---

[51] See Millikan, *Biological,* 16 and also Carroll, 'Forget', 7-8.
[52] John Dilworth, 'Medium, subject matter and representation', *Southern Journal of Philosophy* 41 (2003), 45-62, at 50.

In defending his institutional theory of art, George Dickie posits that the 'artworld' is an untidy collection of culturally-constructed systems.[53] The collection is an untidy one because no person or group of people needs to have consciously decided which things are and aren't art: art just comes about as a result of people's behaviour over time.[54] We can see the traces of this kind of socialization if we consider the ways in which our children are initiated into artistic practices today:

> Art teachers and parents teach children how to be artists and how to display their work. Children are taught how to draw and color and how to put their drawing on the refrigerator door for others to see. What children are being taught are basic *cultural roles* of which every competent member of our society has at least a rudimentary understanding. These cultural roles are, I believe, invented very early on in primitive societies and persist through time into all structured societies.[55]

We initiate our children by giving them paints and pencil crayons, instruments, lessons, and sheet music, by having them rehearse plays and read novels and write stories in school. We reinforce these amateurish efforts by celebrating the fruit of their labours, by attending their performances and recitals. With practice, some of them progress to larger stages. The rest of us may not go on to create much art, but we are nevertheless introduced to the conventions surrounding the generation and appreciation of art. The artworld, Dickie observes, is a repository for many different kinds of cultural roles: artist and public, to be sure, but also agents, critics, curators, dealers, entrepreneurs, movie-goers, philosophers of art, promoters, theorists, and many more.[56]

Dickie's story is a plausible one, but it is a story which began *at least* 40 000 years ago, and we have virtually no information about its twists and turns until the Classical age, ~2500 years ago. It would be helpful, then, if we could move beyond toy examples to see similar 'worlds' developing elsewhere in the real world, at a time for which we have access to fairly extensive records. And it is

---

[53] See George Dickie, 'The Institutional Theory of Art,' *Theories of Art Today*, ed. Noël Carroll, (Madison: The University of Wisconsin Press, 2000), 93-108.
[54] Ibid, 100.
[55] Ibid, 101.
[56] Ibid, 102.

in this connection that I wish to point to Tiziana Andina's account of the development and decline of a 'tulipworld' in seventeenth-century Europe.[57]

The story begins around 1610, when flowers became fashionable accessories throughout Paris, especially at the court of Louis XIII. Although roses were initially deemed the height of fashion they were soon supplanted by tulips, which were judged to be more beautiful. The fashion spread throughout Europe and found a special home in the United Provinces, where it centred on acquiring and wearing a beautiful but extremely rare tulip, the *Semper Augustus*, of which only a handful of bulbs existed. Because the *Semper Augustus* was nearly impossible to obtain, shrewd botanists began to breed their own varieties of tulips. As the market developed, sellers employed respected artists to compile illustrated catalogues in an effort to showcase the varieties on offer. It did not take long for people to stop caring about the bulbs and flowers themselves, or about the status they might confer upon their wearers. They began, instead, to care primarily for the bulbs' and flowers' monetary value, which increased at a frenzied pace until the market collapsed in 1637.

Over the span of twenty-seven years we can observe the birth and death of a cultural practice, the development and decline of a tulipworld. Andina observes that, like the artworld, the tulipworld has a plethora of roles for individuals to occupy, many of which it shares with the artworld.[58] Experts (i.e. botanists/curators) and specialized authors (i.e. gardeners and cultivators/artists) are responsible for the identification and generation of the cultural commodity, merchants and travellers play a central role in its distribution, florists and dealers act as market intermediaries and set value, and as a result both tulips and artworks come to be seen as valuable for their worth as trading currencies or investments, not just their aesthetic interest. Both worlds are thus populated by characters with similar functional roles, and develop markets characterized by

---

[57] See Tiziana Andina, *The philosophy of Art: The Question of Definition - From Hegel to Post-Dantian Theories*, Trans. Natalia Iacobelli, (New York: Bloomsbury, 2013), at 52-6.
[58] Ibid, 58-9.

speculative forces. If we were to observe a similar practice in another culture, but with gold and silver arm-rings as a focal point rather than tulips, we would be able to tell whether arm-rings had the same kind of status in that culture as tulips do in our own by looking at their functional role in that culture's social institutions.[59]

The most important thing to observe, however, is that artworlds and tulipworlds develop through the establishment of precedential practices and a process of gradual mutual reinforcement. Eventually the network of attitudes, beliefs, dispositions, and practices becomes so complex that it appears to have a life of its own, and agents may no longer even be aware of the force that precedent exerts over them. All of these cultural practices and attitudes are part of the tulipworld, and of the artworld in kind: the hobbyist who grows tulips in her garden is as much a part of that world as the dedicated tulip breeder, and the folk artist is as much a part of it as the art critic for *The Nation*. Institutions may only officially recognize or sanction certain kinds of cultural practices, but they are just a contingent outgrowth of worlds (*pace* Searle). The notion of a world is intended to be much broader in scope, so that it includes all cultural activities relevant to the development and establishment of the relevant concept.

With respect to the development of the tulipworld, it is worth noting that it is largely arbitrary, *at first*, which kinds of flowers come to have so much social significance. It was also arbitrary that flowers acquired this social significance in the first place; it could just as easily have been penannular brooches or arm-rings, as has actually been the case in the past. Tulips established themselves as the kind of choice because they were the flowers which socially influential courtiers and nobles chose to wear, and those courtiers and nobles acquired their own trend-setting influence through a long history of precedents rooted in French social conventions. Once tulips became the established kind of choice the nascent tulipworld acquired many more layers of complexity in a very

---

[59] David Davies also develops this idea for 'legal currency' in the context of a 'commerceworld' in his *Art as Performance*, 248-9.

short span of time, especially in terms of economic conventions and roles. the focus on tulips thus became *less* arbitrary over time: the reason everyone cared about them was because they were fashionable, valuable, etc. Had it not been for the tulip market crash of 1637 and the concomitant dismantling of the tulipworld, this cycle might have perpetuated itself indefinitely. The point here is just that parallel considerations apply to the case of the artworld, whose gradual development and longevity have granted it a greater measure of stability.

5. What makes a kind an art-kind

Where does all this leave us? So far, I have argued that what makes a kind an *art*-kind is a matter of convention. In fact, this will be true for *any* social fact or kind. But what makes some conventions *art*-relevant, and others not?

At this point, readers may be tempted to adopt an institutional account of the social world, perhaps alongside an institutional theory of 'art'. This is the path John Searle has recommended: social-world conventions are governed by institutional facts, which exert deontic powers over us— powers marked by terms such as 'duties', 'permissions', and 'rights', among others.[60] These institutions are built up out of collectively-intended cooperative activity, which results in the collective imposition of agentive functions on the world. The result is that an institutional fact depends for its existence on continued human cooperation and recognition of that fact's status as a social fact.[61] On the Searly picture, "the purpose of institutional reality is to create new powers,"[62]

---

[60]John R. Searle, 'What is an institution?', *Journal of Institutional Economics* 1.1 (2005), 10.
[61] John R. Searle, *The Construction of Social Reality* (New York: The Free Press, 1995), at 26, 38; and 'Are There Social Objects?', *Perspectives on Social Ontology and Social Cognition*, edited by Mattia Gallotti and John Michael, (New York: Springer, 2014), 17-26, at 18-21.
[62] Searle, 'Are there', 19.

and those powers are always propositional in structure because they must figure in rational deliberation.[63]

I do not have the space here to delve into the Searly picture. One of its chief merits, however, is that it shifts our focus away from individual social objects to the institutional facts underpinning their creation. In this respect, it echoes my own admonition to look to the processes underpinning the creation and presentation of artworks for their ontological properties, rather than on the artworks themselves. It also offers us a relatively robust top-level guarantee that our art-kind categorizations are substantially correct: correctness of categorization depends upon the application and co-application conditions for the kind-term which are collectively developed by competent participants in the relevant institutional facts and structures.[64] It certainly seems right to me that institutions can exert quite a lot of normative force. Whether they are the only—or primary—vehicle of conventional agglomeration, however, is another matter. My suggestion here is more basic, and does not require much intentional activity at all: the deep ontology of 'art' and art-kinds, along with other social kinds, is rooted in arbitrary and historically-contingent networks of convention. Sometimes these conventions will congeal into a world which allows for the further development of institutional facts. But conventions can exert normative pressure on us in the absence of deliberate intention or institutional regulation—just think of the convention that dresses and gowns are for women, or that you answer the phone by saying 'Hello?'.

It seems to me that, prior to the development of a world, all we have are bare conventions clustered together for largely arbitrary and contingent reasons. As worlds develop around them however, these conventions begin to acquire more and more regulative force. Eventually, once we

---

[63] Ibid, 19-20.
[64] Amie Thomasson articulates this kind of suggestion in 'The Ontology of Art and Knowledge in Aesthetics', *Journal of Aesthetics and Art Criticism,* 63.3 (2005), 221-9, at 223; and 'Debates about the Ontology of Art: What Are We Doing Here?', *Philosophy Compass,* 1.3. I offer reasons for skepticism in Michel-Antoine Xhignesse, 'Fake Views—or Why Concepts are Bad Guides to Art's Ontology', *BJA* 58 (2018), 193–207.

have robust, long-standing institutions, *then* we can say that these govern what counts as an appropriate instance of the practice in question. But institutions are themselves under constant pressure to change, to accommodate new instances and reject old paradigms.

All this suggests to me that we can pose the question in two distinct ways: diachronically, or synchronically. Diachronically: why do some conventions produce art-kinds and institutions, while other conventions do not? All social facts require conventions for their existence, so what, if anything, is special about art-relevant conventions? Synchronically: why aren't the conventions which govern the production and reception of an ordinary coffee mug art-relevant, whereas those governing the production and reception of bizen-ware *are*? Why are *these* conventions art-relevant, while *those* are not?

In isolation, the answer to the diachronic question is, quite simply, that there is *nothing* intrinsically special about art-relevant conventions. There is no deep ontology for us to discover here, nothing that plays the same explanatory role as homeostatic property clusters do for natural kinds.[65] That our conventions are the way they are is ultimately arbitrary and contingent: a different set of historical accidents could just as easily have seen us class drawing, painting, and print-making together, perhaps even treating them all as singular, rather than multiple, art-forms. It is not arbitrary, of course, that our artistic practices typically exploit features of the human sensory apparatus rather than, e.g., sense-modalities only available to star-nosed moles. But the particular forms that this exploitation takes, and the priority we assign to different ways of exploiting the same sense-modalities, *are*, ultimately, arbitrary and historically contingent. The most satisfying answer we can have to the diachronic question is not especially satisfying: because that is what we do, and what we have done for a long time, now. We simply do not have the kind of epistemic access to the

---

[65] I argue that this ontology of conventions is sufficiently stable to ground the reference of social-kind terms in Michel-Antoine Xhignesse, 'Social Kinds, Reference, and Meta-Ontological Revisionism', *Journal of Social Ontology* 4 (2018), 137–156.

prehistoric foundations of our artistic practices we would need to give a satisfying answer to the diachronic question. All we have are some archaeological and biological data, plus increasingly good documentation the closer we get to the current era; and while that is enough for us to make educated guesses, it is not enough for satisfying answers.[66]

What this account *does* offer us is a unified and minimalist strategy for answering *both* the diachronic and the synchronic questions, provided the relevant data are epistemically accessible. The basic strategy is simply to pass the buck down to the conventions and institutions governing the relevant chunk of the artworld.

If, for example, we want to know why cinema is an art-kind, then, as the buck-passer suggests, we should start by looking at the appreciative and other practices informing our engagement with, and production of, film. In other words, we can give a top-level functionalist explanation of cinema's art-status by appealing to the role it plays in present-day culture. For a more robust answer, however, we need a *proceduralist* explanation: we need to trace back cinema's origins to uncover the conventions underpinning its ancestor practices—a fairly simple task for a relatively new and well-documented art-kind like cinema. If we want to know why cinema counts as one of the arts, we need to understand the functionalist justifications of cinema's inclusion among the arts; these appealed to its similarity to photography, for which, in turn, functionalist justifications appealed to its similarity to painting, drawing, and print-making, and which explained its mechanical attributes as extensions of earlier use of the camera obscura, itself predicated on the perforated *gnomon*, etc. Similarly, if we want to know what makes Sacha Gervasi's *Anvil: The Story of Anvil* (2008) a documentary film, then we need to both consult our contemporary documentary practices and trace back the genre's origins in the conventions governing cinéma vérité and direct cinema.

---

[66] See, in this connection, Stephen Davies's groundbreaking *The Artful Species* (Oxford: OUP, 2012), which goes as far as we can towards pinpointing the prehistorical origins of our artistic practices and sensibilities.

To return to the coffee mug objection: the bizen-ware fits into the context of a tradition that takes pieces of its kind to be the proper targets of certain kinds of appreciation which fall under the purview of 'art.' The ordinary mug, on the other hand, does not fall into any such practices. It certainly *could*—its non-art status is not a metaphysical necessity—but it *does not*. And if we want to know exactly why that is, then we need to investigate the origins of our appreciative practice and of the ceramics-world, which together supply us with a genetic basis on which to ground the distinction.

In other words, the answers to these questions become increasingly *less* arbitrary over time, as our initially arbitrary practices congeal into conventions, which consolidate into worlds, which can then generate institutions capable of exercising regulative force. Even if what makes a kind an art-kind is, at first, largely or entirely arbitrary, that arbitrariness is increasingly conventionalized as the artworld develops, and as our network of institutions and practices becomes more and more robust.

VI – Conclusion

The main challenge facing kind-based accounts is to provide some sort of reason for thinking that a particular kind *K* is an art-kind, since there are so many more kinds than art-kinds. I have argued that Lopes is right to suggest we look for the answer somewhere in the vicinity of our appreciative practices, but his framework lacks the resources offered by a further reduction down to systems of interlaced conventions. Every work of art is made of something—that is to say, every work of art uses a vehicular medium to convey its artistic content. Our artistic and appreciative practices group some of these works together arbitrarily, not just on the basis of their vehicular medium or the sense-modalities they exploit but on the basis of fit into extant and historical artistic and appreciative practices. As new works emerge, they are classified and judged according to the ways in which their predecessors were, and over time this accretion of precedent congeals into a convention, which in

turn plays an important role in the composition of an artworld. The answer to the question of why an ordinary coffee mug is not a work of art is the same as the answer to the question of why *Viola Cornuta* lacked *Semper Augustus*'s social significance: quite simply, because it is and was not our custom to engage with coffee mugs and horned pansies in that fashion.

It may seem somewhat unsatisfactory to ground the nature and properties of art-kinds in a process that is often arbitrary and contingent, patched together by the weight of historical precedent. But therein lies the beauty of a model based on the salience of precedent: as precedents accumulate over time and congeal into institutions and worlds, the operations of the conventions they regulate become increasingly *less* arbitrary and contingent. The fact that their distant historical origins were arbitrary and contingent does not entail that their current operations are, too; rather, it simply means that we must take care to ask the right questions, and work to trace back the development of the ontology(ies) that gave rise to them.[67]

This account also has the distinct advantage that it does not require much from agents by way of a direct intention to make art, or to make a work belonging to one of the arts. It simply requires them to make something as they do because that is what others did—being art comes later, once the practice has concretized further. My account shifts responsibility for kind-membership away from individual agents and their propositional attitudes, and away from the work's physical properties, placing it squarely on the conventions governing the work's contexts of appreciation and creation. The basic idea is hardly novel—Timothy Binkley first suggested, in 1977, that cultural conventions act in concert to determine the ways in which certain aesthetic qualities or ideas can be conveyed.[68] And, as we have seen, various elements of the suggestion have been taken up more recently by the likes of Becker, Bicchieri, Clowney, Lopes, Mag Uidhir, and Searle. What was missing was an account of convention that is not shackled to coordination problems.

---

[67] For more on this point, see Xhignesse, 'Fake Views'.
[68] See Timothy Binkley, 'Piece: Contra Aesthetics', *Journal of Aesthetics and Art Criticism* 35.3 (1977), 265-277.

So, what makes a kind an art-kind? Broadly put: a conventional atmosphere.[69]