Attempting Art:

an essay on intention-dependence

Michel-Antoine Xhignesse

Department of Philosophy

McGill University, Montréal

December 2016

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of
Doctor of Philosophy

Table of Contents

DISSERTATION ABSTRACT

Attempting art: an essay on intention-dependence

It is a truism among philosophers that art is *intention*-dependent—that is to say, art-making is an activity that depends in some way on the maker's intentions. Not much thought has been given to just what this entails, however. For instance, most philosophers of art assume that intention-dependence entails *concept*-dependence—i.e. possessing a concept of art is necessary for art-making, so that what prospective artists must intend is *to make art*. And yet, a mounting body of anthropological and art-historical evidence and philosophical argument suggests that not only is such a criterion unsatisfied by most of the art-historical canon, but it also rests on a false premise: concepts of 'art' are *not* shared between cultures, nor even in the same culture across time.

My dissertation aims to rectify this error by first exploring what our commitment to art's intention-dependence actually entails, and then showing that, properly understood, intention-dependence sets a number of important constraints on theories of art with respect to explanatory *desiderata* such as the success- and failure-conditions of art-attempts, the cross-cultural identification of art, and the reference of 'art' and art-kind terms.

I begin by situating art's intention-dependence in the philosophical literature on intentional action, arguing that, properly conceived, intention-dependence is a *weak* criterion which can be satisfied *either* directly or indirectly. It therefore does not necessarily entail concept-dependence. I then use this distinction to motivate a new treatment of the success- and failure-conditions for art-attempts, arguing that the extant model's emphasis on compliance with 'the manner intended' is far too restrictive to capture actual artistic practices.

I go on to show that Ruth Millikan's model of linguistic conventions supplies an independently plausible explanation of art's concept-independent origins in terms of the development of a system of indirectly intention-dependent conventions called an 'artworld'. I argue that this account of artworld development supplies us with the tools we need to distinguish *art*-kinds from other artifactual kinds.

Finally, I turn my attention to methodological issues, arguing that even though 'art' is a social kind with its roots in arbitrary and historically-contingent networks of conventions, the philosophy of art is not merely an exercise in bare conceptual analysis. In fact, there is now a great deal of evidence to show that the ways we think about 'art' are inconsistent, incomplete, imperialistic, and largely unprincipled. Yet I argue that this does not mean that the artworld data have no bearing on theories of art. Instead, I argue that our best reflective understanding of our artworld practices sets the constraints on the reference of 'art' and art-kind terms. I argue that we have no privileged epistemic access to the ontology of social kinds; our only privilege lies in our ability to determine the proper subject of our inquiries.

RÉSUMÉ DE THÈSE

Tenter l'art: un traité sur la dépendance intentionnelle

Il est pris pour acquis parmi les philosophes que les objets d'art expriment une dépendance intentionnelle—c'est-à-dire que fabriquer de l'art serait une activité qui dépendrait dans une certaine mesure des intentions d'un créateur. Mais très peu d'attention a été portée aux implications de cette présupposition. La plupart des philosophes, par exemple, prennent pour acqui que la dépendance intentionnelle implique une dépendance *conceptuelle*—c'est-à-dire, que la possession d'un concept de l'art est nécessaire pour fabriquer des objets d'art et que ce qu'un artiste doit donc avoir l'intention de faire, c'est *de fabriquer un objet d'art.* Cependant, de nouvelles recherches anthropologiques et historiques, ainsi que de nouveaux arguments philosophiques semblent indiquer que ce critère n'est non seulement pas satisfait par la majorité des œuvres dans la tradition historique de l'art, mais en outre qu'il repose sur une fausse prémisse : les conceptions de l'art ne sont *pas* identiques d'une culture à l'autre, ni même au sein d'une même culture au fil du temps.

Ma thèse tente de rectifier cette erreur. D'abord, je montre ce que notre engagement envers la dépendance intentionnelle implique réellement. Ensuite, je démontre que, proprement comprise, la dépendance intentionnelle impose de nombreuses et importantes restrictions à nos théories, surtout par rapport à leurs *desiderata*. Parmi ces derniers, je m'attarderai aux conditions marquant la réussite ou l'échec des tentatives d'art, à l'identification de l'art à travers différentes cultures et à la référence du terme 'art' ainsi qu'à celles des termes désignant les types d'art.

Je commence en situant la dépendance intentionnelle de l'art dans le contexte de la littérature philosophique sur les actions intentionnelles, soutenant que, proprement conçue, la dépendance intentionnelle constitue un critère assez faible qui peut être satisfait soit directement ou indirectement. Il n'implique donc pas nécessairement de dépendance conceptuelle. De là, j'utilise cette distinction pour suggérer un nouveau traitement des conditions de réussite et d'échec des tentatives d'art, soutenant que le modèle actuel met accorde trop d'importance à 'la manière voulue'. Les contraintes qu'impose cette 'manière voulue' sont beaucoup trop restrictives pour rendre justice à la variété des pratiques artistiques actuelles.

Je démontre ensuite que le modèle développé par Ruth Millikan pour expliquer les conventions linguistiques nous offre une explication plausible des origines de l'art en termes de développement d'un système de conventions indirectement dépendantes de l'intention de leurs créateurs. Le développement d'un 'monde d'art' peut donc être conçu *indépendamment* de la présence ou de la possession d'un concept de l'art. Je soutiens que c'est cette explication du développement d'un monde de l'art qui nous donne les outils dont nous avons besoin pour distinguer les types d'art des autres espèces artefactuelles.

Enfin, je me pencherai sur des problèmes méthodologiques, soutenant que même si l'art est une espèce sociale tirant ses origines d'un réseau de conventions arbitraires et historiquement contingentes, la philosophie de l'art ne doit pas se contenter d'une analyse conceptuelle. En fait, il existe maintenant de nombreuses preuves du fait que les façons dont nous pensons l'art sont incohérentes, incomplètes, impérialistes et généralement peu scrupuleuses. Malgré tout, je soutiens que ceci ne démontre pas que les données fournies par le monde de l'art n'ont aucune incidence sur les théories de l'art. Au contraire, ce sont nos meilleures réflexions sur nos pratiques artistiques qui orientent la référence de 'l'art' et des termes désignant les types d'art. Nous n'avons aucun accès privilégié à l'ontologie des espèces sociales, tel l'art; notre seul privilège est notre capacité à bien déterminer le sujet le bon sujet de nos investigations.

ACKNOWLEDGEMENTS

Deyr fé,
deyja frændur,
deyr sjálfur ið sama;
en orðstír
deyr aldregi
hveim er sér góðan getur.

Deyr fé,
deyja frændr,
deyr sjálfr et sama;
ek veit einn,
at aldri deyr:
dómr um dauðan hvern.

Hávamál LXXVI-VII

Chapter 1 – Introductory Remarks

Most philosophical discussions of art today take for granted the premise that art-making is necessarily intentional. Artworks, after all, are *artifacts*, and artifacts are made by agents executing some sort of plan. This fact, in turn, is used to ground the distinction between artworks and objects resulting from natural processes, which do not require any intentions at all. The result is that art-making is an *intention-dependent* activity: it necessarily requires intentional action on an agent's part; nature, by contrast, is thought to be thoroughly intention-*in*dependent. So, for example, the rustle of wind through oak leaves or the patter of raindrops on a tin roof cannot, considered by themselves, be artworks, since they are not the result of intentional action. Unless, that is, we are willing to ascribe them to Oðin or to some other god.

The sounds produced by a clarinet when someone blows across its reed, however, or of a corrugated metal sheet onto which someone drips water, *could*, in principle, count as artworks. This is not because the latter sounds have special acoustic properties which the former lack; on the contrary, the two pairs of sounds might well be acoustically indiscernible. Nor is it to say that the latter sounds definitely *are* art. Rather, the point is just that they have the *potential* to be art because they satisfy at least one necessary condition on art-making: that the artwork be the result of some sort of intentional activity.

To date, the philosophical debate over intentions in art has tended to focus on the interpretive side, on questions concerning a work's meaning or its proper interpretation. William Wimsatt and Monroe Beardsley (1946), for example, led the anti-intentionalist charge by arguing that a work's meaning is in no way determined by authorial intent. On the other hand, actual intentionalists such as E.D. Hirsch (1967) and Gary Iseminger (1992) maintain that it is, but have

struggled to explain how audiences can have epistemic access to those intentions. Splitting the difference, Jerrold Levinson's (1996) hypothetical intentionalism holds that a work's meaning is determined by an audience's best *hypotheses* about the artist's intentions, where the relevant audience is one whose members are familiar with the kinds of background information appropriate to the work in question.

Yet this kind of focus on interpretation leaves unanswered a number of interesting questions of a more ontological nature, including the very basic question of just what it is that we mean when we say that intentions are necessary for art-making. Is it, for instance, the specific intention *to make an artwork* that is necessary, or some other kind of intention, such as *to make (e.g.) music*, *to make something satisfying a particular functional role*, or just *to make an artifact*? What are the success- and failure-conditions for actions guided by these different kinds of intentions? We might also ask the related question of whether possession of a *concept* of art is necessary for art-making (concept-dependence). A great deal hinges on just this question, for if concept-dependence is true, then it may well turn out that a substantial number of allegedly artistic practices from other cultures are not, in fact, artistic. If Balinese musical practices seem to differ significantly from our own, then does that count as evidence that the Balinese concept of music is different from ours? Perhaps Balinese music is not really music after all—worse, perhaps it is not even art! A more immediate concern attaches to our own culture's artistic practices, however, since it is only comparatively recently that we have flocked to galleries, museums, and concert-halls to experience art. What if it turns out that Michelangelo and Mozart did not share their concepts of art with us? Rather than being deeply entrenched in human cultural history, could it be that 'art' is a nineteenth or twentieth century phenomenon?

Focusing on the role of intentions in the ontology of art rather than its interpretation also raises interesting questions with respect to the subject of our inquiries: should we be lavishing our attention on artworks, or on the actions that generate them? On the one hand, it is the artworks (as entities) which we hang on walls, admire, and are concerned to see and talk about. On the other hand, if the last sixty years of philosophical aesthetics is to be believed, their art-status relies on some feature of the process which generated them, not on physical properties of the works themselves.

It also raises important questions with respect to the nature of social kinds in general, and of 'art' in particular. Traditional analyses of social kinds agree that, unlike natural kinds, social kinds do not share an underlying essence: human beings are human in virtue of sharing a particular genetic profile, but Canadian citizens are not Canadian in virtue of their microstructural properties. Similarly, there does not appear to be any one set of physical properties we can identify which make artworks art. 'Art,' like 'Canadian,' seems to depend upon our adoption of certain conventions, but how and why do we adopt some conventions and not others? Does that mean that the ontology of art reflects the ways we *speak* and *think* about art, or is there some more objective measure by which we can judge theories of art? And what happens if the things we say and think about art are *wrong*, or if we later change our minds? Do some works then cease to be art, or turn out never to have been art in the first place? What determines the ontological properties of a social kind, anyway?

Variations of these questions have long been asked in other philosophical subfields—especially in metaphysics, the philosophy of action, and the philosophy of language. It is only recently, however that philosophers of art have begun to reflect on the implications for our analyses of artworks. In a 2007 paper, for example, Dominic McIver Lopes argues that art-

making often happens in the absence of an intention to make art, and that concepts are therefore not necessary for art-making. Similarly, Christy Mag Uidhir's work on the intelligibility of failed-art (2010) led him to develop a preliminary framework for art's intention-dependence in his 2013 monograph. There has even been quite a bit of recent interest in the methodology of philosophical aesthetics.[1] Yet even so, the bulk of these works have assumed art's intention-dependence without reflecting on just what that means, or what it entails for our theorizing.

This monograph represents an attempt to come to grips with just that; I take the one necessary condition on art-making that philosophers universally accept, and begin the work of unpacking its ontological significance. Let me be clear: what I am offering here is not a theory of art. It is a starting point for theories of art, one which I think promises to help us unify some of these otherwise disparate lines of argument. Art's intention-dependence may seem like a trivial property, but my aim in this monograph is to show that it actually raises significant ontological issues, and places important constraints on our theories of art.

I am, of course, starting from the assumption that art is intention-dependent. This means that my results, if correct, will only apply to theories of art that likewise subscribe to intention-dependence. One way of challenging my project is thus just to challenge the basic premise and assert that art-making is an intention-*in*dependent practice. At that point I must confess that, much as I might wish one, I do not have an independently plausible knock-down argument against art's intention-independence. But while I cannot rule it out as a *logical* possibility, it strikes me that the prospects for such an account are rather dim so far as the *actual* world—and our actual artistic practices—are concerned. Adopting it would mean, among other things, pointing to entities that are plausibly artworks, but which are the result of an entirely intention-

---

[1] e.g. Currie (2014), D. Davies (2004, 2009b, 2009c, forthcoming a), Dodd (2007, 2008, 2013), Lopes (2014, 2016), Thomasson (2005, 2006, 2007).

independent process such as natural selection, the operation of natural forces, and so on (or, conversely, arguing that artworks are natural, not social, kinds). One might then claim that a given tree in the forest, or a rock sculpted by the winds of Mars, is an artwork, independently of anyone's stipulation that we think of it as such.[2] Alternately, one could bite some bullets and claim that all things which possess a particular property (e.g. Bullough's 'psychical distance'[3] or Bell's 'significant form'[4]) are art, including many intention-independent entities. The trouble, however, is that both these models of artistic practice seem to run afoul of the (actual) history of art, which does not include any such entities.[5] Whether this history is an appropriate criterion for theory-choice is yet another issue that needs to be settled, along with the role that intentional activity plays in establishing that history.

Accordingly, I have three main goals in this monograph. The first is just to explain what philosophers mean—or *should* mean—by their commitment to art's intention-dependence. Intention-dependence is not unique to artworks; to properly grasp its implications we must understand how it operates in actions more generally. The second goal is to shift the focus of philosophical inquiry from artworks and their properties to the attempts from which they issue. This is just because art-*attempts* do all of the interesting ontological work; art*works* are just the residue left over from the process of intentional action. Finally, my third goal is to unpack the ontological significance of taking art's intention-dependence seriously, especially with respect to the issues raised above concerning the cross-cultural identification of art, recent concerns about methodology in the philosophy of art, and art's status as a social kind.

---

[2] So that it is not merely a piece of found art, but has art-status on its own merits.
[3] See Bullough (1957), especially *'Psychical Distance' as a Factor in Art and an Aesthetic Principle* and *The Modern Conception of Aesthetics*.
[4] Bell (1913).
[5] If it ever did, these soon passed out of our collective memory and practices, leaving no residue for conventions to latch on to.

Chapter 2 tackles the question of just what we mean when we say that art-making is intentional. I begin by situating our commitment to art's intention-dependence within the broader philosophical literature on intentional action in an effort to show that this commitment can be interpreted in two different senses. The stronger sense, *direct* intention-dependence (DID), requires an artist to explicitly formulate an intention to make art. The weaker sense, *indirect* intention-dependence (IID), requires only that someone directly intend to make something which, in turn, happens to satisfy the necessary and sufficient conditions for art, whatever these may be. The rest of the chapter is devoted to arguing for IID over DID, on the grounds that DID entails a commitment to art's *concept*-dependence which, in turn, motivates a descriptively inadequate model of art-making that would see us excise far too many works from art's history.

Chapter 3 focuses on what happens when agents fail to execute their goals, arguing that the answer ultimately hinges on the direction of the attempt in question (i.e. whether it was directly or indirectly intentional). Although the possibility of failure is a necessary corollary to ID, I argue, against a recent treatment of artistic failure, that the possibility of failed-*art is not*. In fact, I argue that extant treatments of failed-art theory have erred in tying the failure of an art-attempt to the work's failure to conform to the artist's intentions. These analyses of failed-art promote a view of 'the manner intended' that is far too restrictive, and which would have us do serious violence to the complex hierarchies of intentions underpinning actual artistic practices. Ultimately, I suggest that the problem stems from focusing on failed-artworks rather than failed-attempts, and go on to build an attempt-theory of failed-art that does justice to the nested sets of intentions which artists bring to bear upon their works.

The problem of concept-dependence rears its head again in Chapter 4, where I consider the recent suggestion that art-making requires a direct intention to make a work in one of the arts

(e.g. literature, music, painting, etc.). I argue that when we focus our attention on what makes a practice one of the arts in the first place, the mistake in this suggestion becomes obvious. Art-kinds cannot be analysed simply in terms of their associated physical media, since these do not suffice to give them their status as *art*-kinds. Not all applications of pigment to canvas are paintings, after all, just as not all bodily movements are dances. The missing ingredient here is a notion of convention which can do the work of marking the difference between art and non-art for a given physical medium. I find this notion in Ruth Millikan's work on the conventions of natural languages, where conventions are analysed in terms of their propagation by historical precedence. I show that Millikan's model of linguistic conventions actually explains the development of artworlds, but that kind-centred accounts of art can only access its explanatory power by way of IID.

In Chapters 5 and 6, I switch gears and reconsider my reliance on descriptive adequacy as a criterion of theory choice. Chapter 5 re-introduces concept dependence, this time in the form of the view that the ontology of art is an exercise in conceptual analysis. The idea here is just that its proper method is a description of the content of our thoughts about art. I introduce a trio of problems for this view, all of which centre on the fact that an increasingly large body of anthropological, art historical, and social psychological evidence and philosophical argument points to the conclusion that the way we think about 'art' is neither historically stable, nor geographically uniform. Ultimately, I argue for a *practice*-driven, rather than a *concept*-driven, methodology characterized by epistemic humility, and illustrate its power by applying it to the case of Balinese *mabarung*.

Finally, in Chapter 6 I tackle the question of whether our views about the nature of a social kind like 'art' can be substantially mistaken, or whether we are simply pulling rabbits out

of the same hats we stuffed them into in the first place. In light of the considerations adduced in the previous four chapters, I conclude that it is likely that our pre-reflective intuitions about art's ontology *are* confused and mistaken, just as they often are for natural kind-terms. The trick, however, is to notice that our intuitions do nothing to fix the ontological content or reference of kind-terms; it is the world which does that for natural kinds, and the social world for social kinds. Our intuitions merely offer a starting point for ontological investigations, not an end-point; they fix our domain of discourse.

Chapter 2 – Art, Intention, and Intention-Dependence

2.1 – Introduction

The division between artworks and natural entities is usually drawn in terms of intentionality. Philosophers more or less unanimously agree that artworks are artifacts or performances, while natural objects are not.[6] Artifacts and performances, of course, are the result of actions; natural objects, by contrast, are just the result of events.[7] This is not to say that nature cannot yield objects of aesthetic value or interest, nor is it to say that someone could not appropriate a natural object and present it as an artwork. If the proliferation of aesthetic theories among philosophers of the eighteenth and nineteenth centuries is anything to go by, there is a great deal of aesthetic interest to be found in nature. And if found art is an acceptable category of art, then the prospects are bright for an attempt to appropriate natural entities as artworks in their own right. Rather, the point is just that *on its own*, nature does not quite have the ability to generate artifacts or performances; some kind of intentionality is required. This requirement, in turn, helps to mark the distinction between traditional aesthetics, whose focus was beauty (natural and human-

---

[6] Consider, e.g., Savile (1969), Dickie (1984), Hilpinen (1992), S. Davies (1991), and Dean (2003). One alternative might be to think of artworks as eternal types rather than as artifacts or the products of actions. This is the default view for some of the arts, especially the multiple and performance arts, including music, literature, dance, and theatre. To my knowledge, however, only Peter Strawson (1974) has actually defended the view that all artworks have a type ontology. Gregory Currie (1989) has argued for the related claim that all art-kinds are multiple (i.e. their members are all capable of having multiple instances, even if they may in fact have only the one). D. Davies (2010b) has likewise suggested that we can profitably distinguish between works which are epistemically multiple, and those which are also provenentially multiple (all works are in principle E-multiple, but only some are also P-multiple). In any event, the point is just that even those who maintain a type ontology for particular art-kinds tend to favour accounts according to which these are 'indicated' or 'norm' types, and thus require an intentional action on the agent's part for their initiation.

[7] See, e.g. Hilpinen (1992) and Rudder Baker (2004).

made), and the contemporary philosophy of art, which concerns itself primarily with the cultural practice of art-making.[8]

This distinction between art and nature mirrors a distinction that forms the cornerstone of the philosophy of action: the distinction between actions and happenings. Actions are said to involve an essential reference to some agent's intentional activity, while happenings do not. Actions are activities guided by an intention; happenings are events that occur in the absence of any intention. This distinction is somewhat coarse, since some putative actions involve a significant element of involuntary or unintentional activity (e.g. coughs, flinches). Nevertheless, the distinction is a useful one insofar as it shows us why there is such a widespread consensus over the ontological importance of intentionality: intentions are thought to play a necessary role in an artwork's coming to possess art-status. In the rest of this monograph, I shall refer to this consensus as the view that artworks are *intention-dependent*; the primary purpose of this first chapter will be to clarify exactly what that means.

I begin, in §2.2, by giving an overview of the literature on intentional action, with a special emphasis on the ontology of actions. These observations will lead me, in §2.3, to distinguish between two ideas that are frequently conflated in the ontology of art: intention-dependence and concept-dependence. In fact, I will argue that a further distinction is necessary, one which distinguishes between two species of intention-dependence. One (direct intention-dependence) entails concept-dependence; the other (indirect intention-dependence) does not. §2.4 will examine Dominic McIver Lopes's suggestion that indirect intention-dependence yields two additional models of art-making, the accidental and the incidental. Only the incidental, I will argue, gives us a plausible kind of art-making, since the accidental case does not actually feature

---

[8] The distinction was first made by Hegel in the first volume of his *Lectures on Fine Art* (2010 [1835]), but went largely unremarked by philosophers until the twentieth century.

intention-dependence. Having established incidental art's legitimacy, I will argue, in §2.5, that intention-dependence actually entails the possibility of concept-*in*dependence, and that intention-dependence is a more fundamental notion than concept-dependence. Finally, in §2.6, I will sketch out the correspondence between my treatment of intention-dependence and a recent proposal by Christy Mag Uidhir which focuses instead on the notion of 'attempts'.

2.2 – Intentional action

The distinction between actions and happenings mentioned in §2.1 runs headfirst into a serious problem. Consider the act of moving one's arm. We could describe the situation in terms of events that befall the arm (e.g. "the arm was moved"), but doing so is not especially satisfying or informative. Alternately, we could explain the arm's movement in terms of something else one does, such as contracting the relevant muscle fibres. Now, how should we explain the contraction of the muscle fibres? Presumably, we should do so in terms of action potentials sent from the brain, through the nervous system, and to the motor neurons that innervate the relevant muscle fibres. It seems that we can keep progressing in this way, with an infinite series of actions causing one another. Just as problematically, we also seem to have arrived at a series of events whose particulars are not really under intentional control in the first place. That is to say, one intends *to move one's arm*, not *to innervate this muscle fibre thusly via such and such an action potential through that motor neuron*. The point is just this: if we dig far enough beneath most actions, it seems we must come to an underlying series of involuntary events, or else to an infinite series of actions. It quickly begins to look as though we cannot act at all.

In the case of artworks, a parallel concern arises: an artwork's coming into existence seems to be logically implied by a successful act of art-making. Some other event would be

required to cause the action that causes the artwork's coming into existence, and so on until we arrive at a series of non-intentional events. It thus looks as though artworks are not the result of intentional actions at all, and that the distinction drawn between art and nature is misapplied—or, at least, that intentions do not always do the explanatory work we expect of them. This means that in order to discover what our commitment to art's intention-dependence entails we must first get clear about what it means for something to be intention-dependent.

In the context of action theory, a number of solutions to the regress identified above have been proposed. One of the earliest and most influential contemporary attempts is Donald Davidson's (1963) causal theory of action, which holds that actions are caused by desires and beliefs and which was itself inspired by earlier (and decidedly non-causal) work by Elizabeth Anscombe (2000 [1957]). For both Davidson and Anscombe, the ontology underpinning actions consists of particular objects and events. Agents are to be included among the objects, and actions among the events. In this way, in describing intentional action we are supposed to quantify over events, not properties.[9] The solution to the regress thus lies in the observation that particular actions are actually particular events which are spatio-temporally located and unrepeatable, and susceptible to different descriptions (among which is one in terms of bodily movements[10]). On this view, actions are events which are intentional under some description; under other descriptions, the action will seem unintentional.[11] This allows us to distinguish between actions that are intentional in the sense of being 'on purpose' (i.e. reflecting the right pro-attitudes on the agent's part) and those that merely feature brute intentionality. So, for example, Oedipus can be said to have married *Jocasta* intentionally ('on purpose'), since she was

---

[9] According to Davidson, we need to quantify over events to make sense of causal speech and to give a correct account of the logical form of action sentences. The data in favour of the existence of events are thus of a semantic nature.
[10] See Moya (1990: 41).
[11] It seems unlikely that any action could be intentional under *all* descriptions, however.

the king's wife and he planned to marry the king's wife. He cannot be said to have married *his mother* intentionally ('on purpose'), however, since he was unaware of his genetic relation to Jocasta. Under the first description he exhibits the right pro-attitudes; under the second, he does not. The events here are identical, despite their different descriptions. Likewise, my moving my arm is intentional ('on purpose') under the description 'contracting the relevant muscle fibres,' but not under the description 'sending an action potential through the nervous system to the motor neurons that innervate the relevant muscle fibre.' The latter event just is the former, differently described.

Despite some disagreement about whether we should quantify over properties or events,[12] Anscombe's and Davidson's ontology of actions enjoys fairly widespread support today.[13] The causal analysis of action, however, faces much more significant dissent. According to causal theories, in cases of intentional action it is one's reasons (or desires) and beliefs that cause one to act. The problem with such a view is that it is susceptible to what is known as the problem of wayward (or deviant) causes: causal chains that feature a deviation in either the agent's practical reasoning or the execution of her intentions, but which nevertheless result in the originally intended, if not 'decided upon,' course of action.

Perhaps the best known case of wayward causes comes from Roderick Chisholm, who asks us to recall that, for causal theorists, actions consist of (i) some desired result *Y*, (ii) the belief that by bringing about *X*, one can bring about *Y*, and (iii) the belief in (ii) and the desire in (i) causing the agent to bring about *X*. We can imagine a situation, however, in which the causal

---

[12] For disagreement from a causal theorist, see Goldman (1970); for similar concerns from the perspective of new volitional theories, see McCann (1974) and O'Shaughnessy (1973). Kim (1989) proposes a new model altogether.

[13] See, e.g., O'Shaughnessy (1973), Hornsby (1980), Bratman (1987), Wilson (1989), Ginet (1990), Moya (1990), Cleveland (1997), and Higginbotham et. al. (2000). A notable exception is Kim (1976), who takes actions to be ordered triples of substances, properties, and times. This leads him to posit an identity theory of events, according to which events are identical if they occur in the same time and place and instantiate the same property. This means that where Davidson would see one event with multiple descriptions, Kim would see different events.

relationship in (iii) is suspect: suppose, for example, that Captain Arthur Hastings's belief and desire so unnerve him that his body reacts in an immediate and unplanned fashion, e.g. by twitching, so that he achieves the goal envisaged in (ii) by some fortuitous accident.[14] Causal definitions of action would require us to say that the agent *X*-ed in order to *Y*, and thereby paper over the Wittgensteinian distinction between acting in accordance with a rule (or plan), and following one. What these kinds of examples show is that something more than mere belief and desire (playing a causal role in producing behaviour) are required to account for intentional action.[15]

Notice that the same problem plagues art-making: a causal model of action would hold that art-making requires (i.a.) a desire to create an artwork *A*, (ii.a.) a belief that by Φ-ing, one can generate *A*, and (iii.a) the belief in (ii.a.) and the desire in (i.a.) causing the agent to bring *A* into existence. Hastings might intend, for example, to create an action-painting by dropping a bucket of paint onto a horizontally laid-out canvas. And Hastings could be so distracted by his contemplation of how best to proceed that his grip (involuntarily) loosened (prematurely), dropping the bucket onto the canvas. Although these events accord with the causal definition of an intentional action, we should feel some discomfort at the prospect of categorizing the action as intentional since Hastings was not (yet) following his plan when he dropped the bucket. The causal chain corresponds to the artist's practical reasoning, but the action remains unintentional.

The problem can be resolved in either of two ways: by adding intention as a necessary component of intentional action, or by restricting the causal relation in some way. One means of doing the latter, which Davidson himself adopted, is to say something to the effect that the causal

---

[14] Chisholm (1966: 29-30). Davidson considers a similar case in (2001 [1980]: 79). These examples are distinct from Anscombe's in (2000 [1957]: §32 p. 57), since hers involves the commission of an error in judgement rather than an involuntary movement.
[15] Moya (1990: 117-8).

chain has to operate "in the right way."[16] Further specification is required, however, if this additional requirement is to actually solve rather than merely illustrate the difficulty.[17] Presumably, by "the right way" we should understand something like "according to the agent's plan/desires." At least, if such a strategy is to deal with counterexamples featuring internal rather than external (i.e. mental rather than physical) causal chains, then it seems most promising to unpack 'the right way' as in some way consonant with the agent's intentions. But in doing so we just return to the first strategy suggested: "the right way" is just "the manner intended," which can itself be described in a number of competing ways (some of which will cash it out as intentional, others of which will not).

The more promising option, then, is to build intentions into the definition of intentional action itself. In one of the earliest such attempts, Arthur Danto proposed that the core of any action is a "basic action," a bodily movement "which [is] *not* caused to happen by the [person] who performs [it]."[18] Basic actions are actions which are not performed by (in virtue of) doing something else: they are *sui generis*, and so can stop the regress and stand as the source of agency. If art-making is the result of the performance of some basic action or series of basic actions (say, moving one's fingers in a particular way such that the end result is a particular application of paint to canvas), then we could preserve the thesis that artworks are artifacts and, thus, the distinction between art and nature.

Unfortunately, although modified versions of Danto's proposal enjoy widespread support today, his original characterization of basic actions will not get us out of our pickle. While basic actions can solve the regress of actions identified earlier, we are still left with the problem that the movement of the fingers (the supposedly basic action) is itself the result of a series of events

---

[16] The implications of such a restriction are discussed in more detail in §3.5.
[17] Moya (1990: 118).
[18] Danto (1965: 145).

which do not seem entirely intentional: sending action potentials to motor neurons which innervate muscle fibres. In order to address this second problem, we could posit that it is the agent herself who is acting. While not necessarily false, this option does not really get us much beyond our initial conception of action as issuing (somehow) from the agent: it either lacks explanatory value, or is viciously circular. To be fair to Danto, however, this is perhaps just because we have reached rock bottom in terms of the availability of explanations, and because any questioning beyond this point cannot really be done without begging the question against the causal theorist.

Since Danto's time, much effort has been devoted to developing new volitional theories[19] which can resolve this dilemma by re-conceiving the notion of basic actions.[20] These theories propose to identify basic actions with volitions, with acts of the will, instead of with bodily movements. This is because bodily movements can occur without one's consent or endorsement, as is the case with muscle spasms. According to volitional theories, willing or trying are the basic actions, not, *pace* Danto, moving one's finger.

Volitions are a promising candidate for basic actions, since they seem to share what we take to be the intuitive features of actions: they are intentional, we are responsible for them and can control them, and they are behaviours.[21] But if volitions are conceived of as thoughts (as is the case with McCann) and, therefore, as pure actions (actions without happenings; i.e. non-physical processes), how can they be said to cause anything at all in the physical realm? Volitions would seem to fall prey to a version of the problem of interactionism: the mental or

---

[19] "New" insofar as they differ from the "old" volitional theory of British empiricism—see Moya (1990: 18).
[20] See, e.g., O'Shaughnessy (1973), McCann (1974), and Hornsby (1980). Daniel Bennett's *Action, Reason, and Purpose* (1965) develops an approach similar to Danto's.
[21] See Moya (1990: 20) and McCann (1974: 452-6).

non-physical cannot cause the physical.[22] The standard volitionist response, following O'Shaughnessy (1973) and Hornsby (1980), is to maintain that tryings are the quintessential causally basic actions: we contribute the trying, and the world contributes the rest. To raise one's arm is thus just to intentionally initiate the causal chain that results in arm-raising; actions are not events caused by volitions, they *are* volitions.[23] Without a trying, there can be no action.[24]

On the volitionist model, a work of art would be the result of a successful art-attempt. In this way, the volitionist can preserve the distinction with which we began, between works of art and works of nature. This is not to say that such a distinction is a necessary *explanans* for any theory of art, nor is it intended to beg the question in favour of art-historical descriptivism. The point, rather, is just that the explanation of the distinction between art and nature requires some reference to an agent's intentions, particularly in the form of an art-attempt. That is what it means to say that art is intention-dependent: it is the result of some art-attempt. An intention-dependent account of art must uphold this distinction on pain of contradiction. While other accounts of art are certainly possible,[25] these will not treat art as a necessarily intention-dependent activity.

The preceding characterizations are, to be sure, roughly hewn. It is not my aim, however, to defend a particular analysis of action, nor do I intend to develop a new theory of action. I prefer to remain neutral about the nature of action; the point of these remarks is just to show that we can use some of the common ground in the philosophy of action to shed light on our

---

[22] They may also fall prey to causal exclusion arguments such as Kim's (1989). Whenever a (mental) volition V is instantiated it would have to be realized by a particular physical property P, in which case P does all of the causal work. Volitions are unnecessary because we already have access to a non-volitional explanation of action.

[23] Moya 115.

[24] Note that, for Hornsby, tryings are actions that occur inside the body: her view holds that the concept of action is a causal one (Hornsby 19080: 33).

[25] E.g. some versions of the mimetic theory of art, aesthetic and other functional definitions (such as Monroe Beardsley's [1982], or those derived from Kant), aesthetic attitude theories, or some cluster theories that do not take intentional activity to be among the necessary disjuncts for art.

ontological commitments about art. If we treat art-making as an intentional activity like any other, then we begin to see the outlines of a useful distinction between activities that are directly and indirectly intentional.

2.3 – Intention-dependence vs. concept-dependence

The discussion of intentional action above leaves us in a position to explore the significance of taking art to be intention-dependent. Before we can proceed with that investigation, however, we must begin by distinguishing the claim that art is intention-dependent from the claim that it is concept-dependent. Speaking generally, we can define concept-dependence about some activity Φ as follows:

> *Concept-dependence* (CD)
> An act of Φ-ing is *concept*-dependent iff Φ-ing requires a concept of Φ.

Concept-dependence about art is just the position that art-making requires a concept of art. Anthropologists and philosophers alike are fond of claiming that cultures (including our own) can only be said to have art if they first possess a concept of art, as evidenced (perhaps) by the presence of such a word in the language, or by the existence of a class of objects that functions culturally just as art does in our (or some other) culture(s).[26] When these people say that art is concept-dependent, they do not merely mean that art-making requires the possession of some concepts in general—that much is trivially true. Rather, what is meant is that art-making requires a fine-grained concept of *art* in particular. The first sort of claim is innocuous enough; the second, I shall argue, ought to be rejected, but it must first be distinguished from the more widely accepted claim that all art is artifactual (or performative), and therefore intentional.

---

[26] Stephen Davies makes, but does not endorse, some of these suggestions in his (2000); similar views are endorsed by Denis Dutton (1995, 2000) and Richard Wollheim (1980 [1968]). David Clowney (2011) has embraced concept-dependence wholesale, based on the arguments of Paul Oskar Kristeller (1951 & 1952) and Larry Shiner (2001). Other endorsements can be found in Vogel (1997), Novitz (1998), Carroll (1993), and Dean (2003).

Speaking generally once again, we can define intention-dependence for some Φ as follows:

> *Intention-dependence* (ID)
> An act of Φ-ing is *intention*-dependent iff it is necessary that the agent act intentionally.

Applied to art-making, ID merely captures the philosophically-widespread belief that art-making requires an intention, which is what allows us to distinguish artworks from natural objects. Nothing I have yet said should seem especially controversial. But there are at least two ways of satisfying the requirement in ID, and herein lies the potential for controversy. The first, *direct* ID (DID), is satisfied when the object's creation is guided by the deliberate or explicit intention that it be a particular way, while the second, *indirect* ID (IID), is satisfied when its creation is guided by some other intention:

> *Direct intention-dependence* (DID)
> An act of Φ-ing is *directly* intention-dependent iff the agent intends to Φ.

> *Indirect intention-dependence* (IID)
> An act of Φ-ing is *indirectly* intention-dependent iff the agent intends to Ψ, where Ψ-ing entails the satisfaction of the conditions for Φ.

DID and IID define two kinds of intention-dependence. Applied to acts of art-making, DID requires an intention to make art. That is to say, DID requires that the agent operate under a particular intention with a particular content (viz. art-making) before it will allow us to call the result of her actions 'art'. By contrast, when applied to acts of art-making IID requires only an intention to make *something*, so long as that something's successful execution entails the satisfaction of the necessary and sufficient conditions for art. IID, in other words, sets a minimum threshold on what can count as satisfying the predicate 'is intentional,' namely that the agent must intend (directly) to do *something* (Ψ), although she need not be aware that Ψ-ing entails Φ-ing. Since by hypothesis all acts of Ψ-ing are Φ-ings, and Ψ-ing was intentional, then

under IID the act of Φ-ing is also intentional. Thus, the agent may be said to Φ intentionally so long as she either (1) intends to Φ, or (2) intends to do something else (Ψ) which happens to entail the satisfaction of the description Φ. The question before us—which I will endeavour to answer in the following sections—is whether *all* acts of art-making must be directly intention-dependent, or whether some can also be indirectly intention-dependent.

This discussion should help to clarify some facts about CD and ID. First, notice that ID does not entail CD: the requirements of ID can be satisfied by either DID (which is concept-dependent) or IID (which is not). Second, although IID still requires that agents possess *some* concepts, it is not, for all that, concept-*dependent*: it does not require of intentional activity that it be guided by a concept of every (or any particular) description under which its product might fall. According to IID, intentional action does not require agents to possess a level of clarity about the consequences of their actions to rival a chess grand master's. What matters with respect to IID is that we are not asking whether the agent's *Ψ-ing* was intentional (it was directly so), we are asking about her *Φ-ing*. And under the description of her action as Φ, no concept of Φ is required, since the product of the agent's action is Φ in virtue of its being Ψ. Finally, the possession of some concept Ψ is no guarantee of possession of another concept Φ, even if Ψ entails Φ or vice-versa. Suppose, for example, that being art entails being either a literary work, a work of music, a painting, or a sculpture. Even if Hastings is in possession of the concept 'art', it does not follow that he has a concept of 'sculpture' or 'music', although we might begin to doubt his competence if he lacked too many of these. Similarly, Hastings might have concepts of 'literature', 'music', 'painting', and 'sculpture' without knowing that these together delimit the concept of 'art': he would have conditions of application for the concept, but not the concept itself.

We can see, then, that intention-dependence is the more fundamental notion here. Concept-dependence is not an independent feature of theories of art: it can only be secured derivatively, by an appeal to DID. In fact, concept- and intention-dependence are often intimately intertwined in just this way. Consider George Dickie's version of the institutional theory, according to which an artwork is an artifact made by someone who intends it for presentation as a candidate for appreciation to an artworld public whose members recognize it as such.[27] The fact that the work is an artifact intended for presentation as art secures its intention-dependence, while the content of the artist's intention (viz. 'that it be received as art') as well as the public's recognition of the work as falling under the concept 'art' secure its concept-dependence by making essential reference either to the concept 'art,' or to a concept of one of the arts (what Dickie calls 'artworld subsystems').[28]

Institutional theories are intention-dependent because they require the artist to do what she does intentionally, and they are concept-dependent because they require that the artist understand what she is doing as contributing to a culture of *art*-making (or *music*-making, and so on for the other arts), and because they require that the work's audience understand it as just such a contribution. Without some suitably fine-grained concept of art, no such contributions (or recognitions of such contributions) are possible. It is this same concept-dependence, however, that gets institutional theories into a spot of trouble by grounding a common objection to them: how can we identify artworks in cultures whose practices—whose 'artworld subsystems'—differ

---

[27] See Dickie (2000).
[28] For a more detailed treatment of ID in the context of reducing 'art' to a disjunction of the arts, see Ch. 4.

significantly from our own?[29] Institutional definitions are thus circular by nature, since they require a concept of art in order to get off the ground.[30]

To be sure, this is hardly a devastating objection; institutionalists have marshalled a variety of responses, but exploring these lies beyond the scope of my present project. My point here has just been to show that concept-dependent theories of art must get their concept-dependence via DID. This, in turn, shows us that intention-dependence is the more fundamental notion. Just how that intention-dependence gets cashed out, however, is up for debate. In the rest of this chapter I will argue that DID places constraints on art-making that are far too onerous, and would require extensive revisions to the art-historical canon. Instead, I will argue that we should take the ontology of action as our guide and use IID to set the minimum threshold for art-making. Doing so will allow us to accommodate both the history of art and the possibility of DID.

2.4 – Accidents and Incidents

I observed earlier that philosophers of action are generally agreed that a behaviour *B* is an action iff *B* is intentional under some description. Although actions are not intentional under *all* descriptions, it is important to note that what it is to be an action is just to be intentional under *some* description. Likewise, while it follows from their artifactuality (or performative nature) that all artworks are the products of actions, it does not follow that they are all made with the specific intention to make art. By way of analogy, consider the process of hiking a trail. Suppose Dame Celia Westholme's intent is to hike the Orion trail: she arrives at the trailhead, notes the trail

---

[29] I will tackle the problem of the cross-cultural identification of art and art-kinds in more detail in Ch. 5.
[30] Though not, Dickie thinks, viciously so (2000: 101-2). They are only as circular as any other social or cultural entity's definition would be. For a more detailed exploration of this problem with institutional theories, see D. Davies (2004), Ch. 10.1, esp. 248-9.

markers, and says to herself: "I'll take this trail." Now suppose that, unbeknownst to her, an inept junior park ranger has (consistently) swapped the signs for the Orion and Cygnus trails. Consequently, once her hike is completed Dame Westholme has hiked the Cygnus and not the Orion, despite her intention to the contrary. Although she unwittingly hiked the Cygnus, it would be misleading to say that Dame Westholme did so entirely *unintentionally*, since it would invite us to conflate two descriptions of her intention. She did not hike the Cygnus 'on purpose,' but she did do so intentionally: she succeeded relative to the description of her intention as one to hike *this* trail, even if she failed relative to the description of her intention as one to hike *the Orion*. The point is this: she intentionally hiked the trail she hiked, but she did not intentionally hike the Cygnus. Rather than say she hiked the Cygnus *un*intentionally, we should say she hiked it *non*-intentionally. In much the same way, we might say that Oedipus intentionally married Jocasta, although he did not intentionally marry his mother (i.e. he was unaware that Jocasta = Oedipus's mother). Under some descriptions, these acts of hiking and marrying have the property of being intentional, while under others they do not.

What Dame Westholme and Oedipus do are actions because there is some description of those events according to which what they did was intentional (viz., hiking *this* trail, marrying *that* woman, Jocasta). What we are interested in, however, is not the question of whether the event was an action. Rather, we are interested in the very description under which the event is intentional, and it is this interest that compels us to distinguish between 'intentional' used to mean 'on purpose,' and used to mean 'with some intention or other.' One way of marking this distinction is in terms of degrees of intentionality. According to the weaker sense (labelled IID in §2.3), the product of the agent's activity merely accords with her plan. According to the stronger sense (labelled DID in §2.3), the product was generated by following her plan: it involves a

normative commitment to doing things a certain way.[31] The issue here is what plan the agent follows, rather than what follows logically from her plan: we should say that Dame Westholme indirectly intends to hike the Cygnus, just as Oedipus indirectly intends to wed his mother. I will continue to explore what is actually going on in these cases in §2.5. For now, however, let us turn back to the question of art-making.

It seems that someone might, in principle, succeed in making an artwork despite never intending to do any such thing. In fact, following Lopes (2007)'s lead, we can distinguish at least two ways in which one can non-intentionally make art. First, we have 'accidental' art:

> *Accidental Art (ACC)*
> S accidentally makes an F just in case S intends to make a G, an F is not a G, S fails to make a G, and, in failing to make a G, S makes an F.[32]

A non-art example will serve to make this clear.[33] Suppose that Hastings (S) is a neophyte canoeist who intends to make a cruising stroke (G). A cruising stroke is not a J-stroke (F). Finally, Hastings fails to make a cruising stroke and, in so doing (because, in fact) he completes a J-stroke instead (an easy mistake to make). We can then say that he has accidentally J-stroked, since he neither completed a cruising stroke nor intended the J-stroke.

Second, we have 'incidental' art:

> *Incidental Art (INC)*
> S makes an F incidentally just in case S intends to make a G, S does not intend to make an F, S makes a G, and in making a G, S also makes an F.[34]

---

[31] For a development of the idea that intentional action is (or can be) normative, see Moya (1990).

[32] Lopes (2007: 8)

[33] I offer a non-art case here to show how ACC is meant to work in principle. In fact, substitution of art-cases (e.g. with a marble pillar (non-art) for G and a sculpture (art) for F) into ACC's schema yield results which are hard to countenance as art, for reasons I will explore below. For the time being, it suffices to observe that there is something strange in the idea that one might just 'end up' with an artwork unless one is already in possession of a theory of the necessary and sufficient conditions for art. In the absence of such a consensus it seems rather more plausible to say that what one ends up with is an object which one can then turn around and *present as* art (thus leaving the question of its art-status undecided). Doing so, however, would undermine F's accidentality; it is for this reason that I argue below that ACC *cannot* actually yield artworks.

[34] Lopes (2007: 9)

To illustrate, suppose that being a triangular prism[35] is, in fact, sufficient for art-status, and suppose that Dame Westholme comes from a culture without any art or knowledge of art. Finally, suppose that she gets lost while (mistakenly) hiking the *Cygnus*. A seasoned outdoorswoman, she uses her shoelaces and emergency blanket to erect a shelter in the shape of a triangular prism. Unbeknownst to her, Dame Westholme has created a work of art, even though she never intended anything of the sort: hers is an incidental artwork. It is not accidental, since she directly intended to give her shelter that shape, and succeeded. It is incidental because the task she directly intended to accomplish also happens to satisfy the sufficient conditions for art, despite never being intended as art. It is art indirectly (hence: incidentally). Of course, cases of incidental art need not feature so much coincidence; indeed, most will be far more mundane, featuring the development of works that fit more or less neatly into existing categories of art.

According to Lopes, both accidental and incidental art-making are possible means of making art non-intentionally—that is to say, the process features no intention to make *an artwork* and is therefore not dependent upon possession of the concept 'art.'[36] Given the ontology of actions outlined above and the account of intention-dependence developed in §2.3, however, we are in a position to see that Lopes is only half right: only one of the two cases is intention-dependent, and that is the incidental case.

Let us begin by examining accidental art-making. When an *F* is made *accidentally*, it is not a *G* and is only *F* non-intentionally. Although an intention to make a *G* is present, it fails, and in failing results in an *F* entirely non-intentionally. This kind of case is parallel to that of Dame Westholme above who, recall, intended to hike a particular trail but failed to do so and, in so failing, ended up hiking another instead. Just as we can find a description of these events such

---

[35] I have in mind here the geometric shape, not the optical instrument.
[36] Lopes (2007: 8-9).

that Dame Westholme's actions are intentional (by specifying the content of her intention so that it describes *this* particular trail ostensively, rather than by a name), so too can we find a description according to which our agent's accidental art-making turns out to be intentional: she intentionally followed her plan, not realizing it was flawed and would result in an artwork instead (although she did not intentionally follow a flawed plan).

So accidental creation nevertheless features an action, not a mere happening. The important question is whether the final product, the putative artwork, is an intentional object. It is not, since it derives its art-status neither from being intentionally endowed with it, nor from the agent's successfully following a plan (and thus endowing it with its properties—the very properties by virtue of which it resembles an artwork—intentionally). If the product of accidental creation were an artwork, that result would tell against the distinction between art and nature. This is because under the *G*-description, the action is a failure. This is important because it is only the *G*-description which features any intentionality at all, and under that description the action fails; considered solely under the *F*-description, it would be a mere happening. So, if we take intention-dependence to be a necessary feature of artworks, then what Lopes calls "accidental art" cannot, in fact, be art, since its products are not intention-dependent. The failure of the intention to make a *G* renders the product of that action non-intentional. In much the same way we do not typically hold people (morally) responsible for the accidental consequences of their actions.[37]

By contrast, when an *F* is made *incidentally*, it is made with the (successfully realized) intention that it be a *G* and not with the intention that it be an *F*. Consider, for example, medieval Church frescoes: there is no doubt that their primary intended purpose was to serve as devotional

---

[37] See, e.g., Thomas Nagel's *Mortal Questions* (1979, esp. ch. 3, "Moral Luck") and Bernard Williams's *Moral Luck* (1981), and the discussion they have prompted in ethics. Sherri Irvin comes to much the same conclusion in her (2012: 250), where she observes that accidentality undermines art-status.

objects, and their makers might not have shared our concept of art,[38] and yet their work has been unequivocally incorporated into our art historical canon. These would be incidental artworks, along with classical statuary, cathedrals, epic poetry, and (perhaps) pre-historic cave paintings.

Now, incidental creation is intention-dependent because one both intends to make a G and succeeds in doing so; the fact that G is also F does nothing to undermine the intentional character of the action or its product (although it was not intentional that it satisfy some other description of the product). And even though incidental art is intention-dependent, it neither entails nor requires concept-dependence: one need not have any concept at all of the thing that is incidentally created in order to create it. Nor need one acquire a concept of that thing, or recognize that it falls under such a concept, once it has been created. All that is required is a concept of the properties one intends to instantiate by one's act of creation. Whether one also draws the inference that things with *these* properties are of *that* functional kind is irrelevant. The incidental art case thus maps neatly onto that of indirect intention-dependence (IID) outlined in §2.3 above.

Accidental art-making, however, maps onto neither IID nor DID; it is not even weakly intentional, since the resulting F is inconsistent with the agent's plan to G (since that plan failed). Although incidental art-making does exhibit the stronger, normative, sense of intentionality (since one's plan goes off without a hitch), it is important to note that the fact that the product is F in addition to being G was *not* an intended part of that plan. So although the activity of G-ing is intentional in the stronger, direct sense, its product's art-status does not derive from that intention, but rather from the fact that following the plan yielded an object (G) which also satisfies the conditions for being an F. So, G-ing is thus indirectly F-ing. G's art-status (its satisfying the additional description 'art') is indirectly (incidentally) intentional, despite the

---

[38] For variations on this theme, see Baxandall (1988 [1972]), Goehr (1992), and Shiner (2001).

object itself having been produced directly intentionally. Finally, deliberate art-making is, of course, intentional in the stronger sense of following a particular, normative, plan of action.

The difference between accidental and incidental art-making, then, lies in that the accidental case features a *failed* intention, while the incidental case features a successful intention (to make an object not directly intended to be art, but which qualifies as such nonetheless). This difference is important, since the accidental case features (1) no art-making intention, and (2) a failed, non-art, intention. Thus, while accidental creation turns out not to be intention-dependent (because one neither intends to make an *F*, nor does one succeed in what one *does* intend to do), incidental creation *is*, at least because it accords with the weaker sense of intentional action stemming from IID.

2.5 – Separating Concept- and Intention-Dependence

Given an Anscombian/Davidsonian ontology of actions, these observations should come as no surprise. Indeed, I remarked in §2.2 that, on such a model, the event constituting Oedipus's act of marrying his mother can be described as both intentional (marrying Jocasta) and non-intentional (marrying his mother). Recall that both Anscombe (2000 [1957]) and Davidson (2001 [1980]) observed that a single action can be given a number of differing descriptions, some rendering it intentional, others not. For Anscombe, it is an agent's *knowledge* of what she does that generates the descriptions under which we call that action the execution of her intentions;[39] the result is that, for her, even 'incidental' actions should be counted as non-intentional (since they are not tied to the agent's knowledge of what she does).

---

[39] Anscombe (2000 [1957]), §48, p.87; see also §25.

Similarly, for Davidson and his followers it is the agent's *beliefs and desires* that inform this description and thus *cause* the action in question,[40] so that any action whose causal effects are unforeseen (or those which should be but are not) and which do not feature in a primary reason come out as non-intentional.[41] The later Davidson,[42] Carl Ginet,[43] and Michael Bratman,[44] however, all endorse the somewhat different thesis that even though the expected side effects of an intended action might not have been intended by the agent, they are nevertheless effects she brings about intentionally. This observation holds equally well for side effects which are not expected but *should* have been or which follow given the logical consequences of an agent's knowledge.[45] So, as Ginet points out, a broom given hard use is intentionally worn down even though the agent may not have explicitly desired to wear it down.[46]

We are now in a position to see that all this talk of 'non-intentional' actions is misleading, since it invites us to conflate the accidental and incidental cases. Instead, we should say that Oedipus's act of marrying his mother is *incidental* (to borrow Lopes's terminology) or 'indirectly intentional.' It exhibited some degree of intentionality, but not the strong or direct variety: although he possessed the direct intention to marry a particular woman, "Oedipus's mother" was not (supposed to be) among her descriptions. Actions of this type contrast with actions whose results are accidental, since the accidental result is not tied to any intention (whether direct or indirect) to bring it about. The intentions in the accidental case fail to bring about the intended result—we might thus more appropriately characterize the accidental case as

---

[40] Davidson (1963: 685-7).
[41] See, e.g., Davidson's discussion of the prowler in *Actions, Reasons, and Causes* (1963, esp. 686-8), where he concludes that alerting the prowler is unintentional.
[42] Davidson (2001 [1980]: 50).
[43] Ginet (1990: 76).
[44] Bratman (1987: 123-4).
[45] For some skepticism, see Chisholm (1966) and Moya (1990, esp. Ch. 12).
[46] Ginet (1990: 76); see also Bratman (1987: 123-4).

one of deviant or failed intentionality, a possibility which I will explore in more detail in Ch. 3. It is the degree to which the agent's plan directs her activity that matters here.

So far, I have argued that art-making is intention-dependent, and that this commitment can be understood to require either direct or indirect intentionality. I have argued that indirect intentionality supplies the best way of understanding this commitment, since it allows for the possibility of incidental art-making. Theories of art that allow for incidental art have no problem explaining the art-status of works made by other cultures, or made under the aegis of some concept that does not match our concept of fine art.[47] What is more, they make this room without sacrificing the possibility of directly intention-dependent artworks. If I am right, then this means that art-making actually requires the possibility of concept-*in*dependence. If 'art' has a common metaphysical essence, as theories of art have commonly assumed, then art-status *must* be concept-independent, since an entity could come to satisfy the conditions required by that essence incidentally. The significance of my argument so far is that even if 'art' has no common metaphysical essence (as seems likely), we nevertheless cannot rule out the *possibility* of concept-independence.

So any theory of art that precludes the possibility of incidental art has taken a wrong turn at the level of spelling out its commitment to intention-dependence.[48] If we assume that intention-dependence is a *necessary* component of art-status, then it follows that concept-dependence is not: intention-dependence is the more fundamental notion here, and it cannot be had any other way. Concept-dependence, on the other hand, can be derived from intention-dependence through directly intention-dependent actions. Consider again the institutional theory of art, according to which to be art is to count as art in a context *C*, usually by being presented

---

[47] Unlike, e.g., at least some versions of the institutional theory.
[48] The reason for this has to do with the facts of our artistic practices, which regularly accept and accommodate incidental works. I will discuss this issue in more detail throughout Chs. 4, 5 and 6.

for appreciation to a public whose members recognize it as art.[49] Such a presentation requires a directly intention-dependent action, and such a recognition requires the possession of the relevant concept, so concept-dependence is a necessary condition for art-making according to the institutional theory. Institutional theories are also committed to intention-dependence insofar as they assume that artworks are artifacts or performances.[50] But if, as I have just argued, a proper account of art's intention-dependence must license its incidental creation (since it is the weak sense of intentionality that supplies the threshold which must be met for some event to count as an act of Φ-ing), then the institutional theorist is in the unenviable position of being forced to choose between her commitments. If she rejects the thesis that IID is sufficient for art-making then she must embrace concept-dependence wholesale, so that her commitment to ID just boils down to a commitment to the fundamentality of CD. At that point, the institutional theorist will either (1) have to deny that entities produced without our concept of art are artworks, or (2) accept the extensional inadequacy of her theory.

This is not to say that intention-dependence comes at the cost of rejecting brute concept-dependence for *any* artifact-type. As Lopes observes, it is not plausible that some artifacts—cyclotrons and automobiles, for example—could be made incidentally, since they are so complex that there are no other artifacts sufficiently close to them which one could be making instead.[51] One could, however, incidentally make artifacts belonging to a subtype of these types, since there exist other subtypes in the vicinity: *isochronous* cyclotrons or *synchro*cyclotrons, for instance, or *sports* cars.

Lopes's examples show us that no matter how one makes art—deliberately (under a concept of art; following a plan) or incidentally (in accordance with a plan)—one's activity is

[49] Dickie (2000, esp. 99).
[50] Dickie (1984).
[51] Lopes (2007: 10).

necessarily intentional in at least the weak sense, even though the possession of specific artistic intentions is not itself necessary for art-making (as with, e.g., prehistoric cave-paintings). The fact that such specific intentions are unnecessary rules out concept-dependence as a necessary condition on art-making, since to be concept-dependent is just to possess exactly such a specific intention. Lopes's distinction is therefore of special interest to philosophers of art, because it gives us the tools to explain how artworks can be made in temporally or geographically distant cultures which do not share our concepts and values: their practices coincide with ours either accidentally, deliberately, incidentally, or not at all. If it is either of the middle two, then we have art; if the first, then we do not. If their practices do not coincide with ours at all, or introduce significant disanalogies,[52] then we face the difficult decision of whether to include them as a new kind of artistic practice, or to exclude them altogether. Under the description 'art-making,' their output is at least weakly intentional and therefore satisfies the requirement of intention-dependence. It can count as art, provided it also complies with whatever *other* standards exist for what counts as art, if there are any—standards which, it must be said, need not be internal to their own culture.

Should we accept Lopes's account? I have argued that a commitment to intention-dependence should force us to abandon accidental art-making as a plausible option. But do we have any positive ground for thinking that an amended Lopesian account presents us with a more desirable model of art-making than strictly concept-dependent theories do? At least one advantage of this approach is that it helps us to preserve the art-historical canon, and to explain some of its inclusions. Even if, for example, it should turn out that a neolithic fertility statuette was not made with any concept of art in mind, that statuette can still turn out to be art (although

---

[52] See, e.g., S. Davies (2007)'s observations concerning Balinese cultural practices. I will have more to say on this topic in Ch. 5, esp. §5.6.

it is not necessarily so), thereby justifying its inclusion in our art-historical canon. Similarly, even if it was made for and served primarily (or even entirely) religious purposes, and even if its contemporary audience had no concept of art like ours, Piero's *Resurrection* fresco can still count as art, provided it meets whatever other criteria a theory of art requires of it. Although I am not currently in the business of supplying a fully-fledged theory of art, the important point here is that no theory of art consistent with the framework I have provided needs to prejudge the issue against such objects counting as artworks. It may still turn out, on a given theory, that these objects are not artworks,[53] but at least they will not have been consigned to the junk heap based simply on the psychological states of their makers.

At a minimum, a good theory of art should be compatible with one or more plausible explanations of the origins of art. Ideally, it should adopt or suggest such an account. The account of intention-dependence that I championed above achieves this result thanks to the possibility of incidental art. While we cannot ascertain whether some early artifacts (such as cave-paintings) are directly (art-) intention-dependent, it is obvious that they are at least indirectly intention-dependent. At the very least, they seem to count as incidental artworks because they were clearly intentionally created (this is obvious from the fact that they are representational) and they seem to comply with our standards of art. In this capacity, they helped to ground the formation of a cultural practice centred on their production.[54] My account treats art-making as an intentional activity like any other, and posits that it shares the same basic ontology. Human beings do not need to realize that they are creating art in order to make

---

[53] For more on this point, see Ch. 3, which addresses the possiblity of failed-art, and Ch. 5, where I argue that our judgements about art-status and about an art-kind's ontological properties are prone to serious error.
[54] For more on just how this might work, see Ch. 4, esp. §4.4.

artworks; they need only engage in the intentional production of objects (or performances), and let the concretization of the practice do the rest.[55]

2.6 – Attempt-dependence

Finally, let us turn to a cognate treatment of intention-dependence, which has been proposed in terms of 'attempt-dependence'. Christy Mag Uidhir (2010, 2013) invites us to notice that intention-dependence underlies a more general constraint on art-making: attempt-dependence. According to Mag Uidhir, art-making is not just an intention-dependent practice, it is *attempt-dependent*—that is to say, all art is necessarily the result of some art-attempt. To understand what he means by this, we must first observe that an attempt is just an action guided by the intention that a particular goal obtain [in the manner prescribed][56]—essentially, attempts are just what new volitional theories of action called "tryings". Attempts thus consist of a goal and an intentional action (acting with the intention that *F*), and an attempt will count as successful so long as the goal obtains [in the manner intended].[57] An attempt to become a doctor will therefore consist of a series of actions (viz. whatever is required to realize the goal in the appropriate manner, such as—in the case of a medical doctor—completing a series of exams and a residency) undertaken so as to realize the goal of becoming a doctor. Like doctor-attempts, what counts as an art-attempt will be determined by what it is for something to be an artwork.

---

[55] It is worth observing that the ontology of the objects in question is not dependent on our social practices; these practices merely serve to identify those objects it is whose ontology we are concerned to investigate. For a more detailed treatment of this point, see Ch. 6.

[56] Mag Uidhir (2013: 17). This requirement is Mag Uidhir's. I have square-bracketed it in this section so as not to pre-judge the issue of whether it is an appropriate restriction. A more detailed discussion of this point follows, and is taken up again in Ch. 3 (esp. §3.5), where I argue that it leads Mag Uidhir into choppy seas.

[57] Mag Uidhir (2013: 17-8).

Mag Uidhir (2013) invites us to distinguish between two kinds of attempts, which he labels attempts *de re* and attempts *de dicto*.[58] The former, which he calls 'attempts to Φ', are what I call *direct* attempts: attempts taking Φ as an explicit goal. The latter kind he calls 'Φ-attempts', and these name all attempts (including attempts to Φ) the satisfaction of whose success conditions entails the satisfaction of the conditions for Φ. But Mag Uidhir's Φ-attempts do not quite map on to what I have characterized as *indirect* attempts, although these are included in the class of Φ-attempts. This is because the class of Φ-attempts is much larger than that of indirect attempts. On Mag Uidhir's model, an art-attempt is any attempt whatsoever, the satisfaction of whose success-conditions entails the satisfaction of the conditions for being art. This means that direct art-attempts (attempts to make art) are included in the class of art-attempts, as are indirect attempts (attempts whose success entails the satisfaction of the sufficient conditions for art, whatever those might be). By contrast, direct attempts do not fall under the class of indirect attempts, since the latter class already assumes that it is something else (Ψ, not Φ) which is directly attempted. Indirect attempts are just another subclass of art-attempts, one meant to capture what is left over when we subtract direct attempts from the class of art-attempts. It is evident, then, that even on Mag Uidhir's model the weaker (indirect) sense of intentionality supplies the minimum threshold for what counts as an art-attempt.

An example should help to make this clear. Consider doctor-attempts: the class consisting of all attempts the satisfaction of whose success-conditions entails the satisfaction of the conditions for being a (medical) doctor. Let us forget about residencies and suppose, for the sake of simplicity, that passing the right set of exams suffices to become a doctor. One way to become a doctor, then, is just to go through medical school with the aim of sitting and passing the right

---

[58] Mag Uidhir (2013: 21). Although Mag Uidhir mentions the *de re*/*de dicto* distinction in a subheading, it does not appear elsewhere in his monograph. I am thus left to reconstruct that attempts *de re* are his *attempts to Φ*, while attempts *de dicto* are his *Φ-attempts*.

examinations. In this way, one directly attempts to become a doctor. But suppose a hapless Hastings, planning to become a *lawyer*, unwittingly shows up to the medical examination room and forms the intention *to pass the exam before him*, unaware that it is a medical, not a bar, exam. Even though Hastings has no intention to pass a medical exam, he does intend to pass *the exam before him*, and satisfying the success conditions for passing the exam before him entails the satisfaction of the conditions for becoming a doctor. Hastings therefore engages in a doctor-attempt without attempting to become a doctor: his indirect attempt suffices to count as being in the relevant attempt-class for doctoring. Indirect intentionality can thus suffice to ground a Φ-attempt.

2.7 – Conclusion

I began this chapter by arguing that the distinction between artworks and natural objects hinges upon the role of intentions, and thus runs parallel to a distinction, in the philosophy of action, between actions and happenings. In Anscombe's and Davidson's ontology of actions, events are only ever counted as an action if they are intentional under *some* description. Following Anscombe's and Davidson's example, I argued that an object can only count as an artwork if we can give some description of it according to which it is the product of intentional action. From there, it became imperative that I distinguish between the notions of intention-dependence and concept-dependence; while the latter entails the possibility of the former, I argued, the reverse does not hold. This distinction, in turn, grounded the conclusion that, properly understood, intention-dependence does not require the possession of some specific intention to or concept of *F*; rather, it merely requires that one intend to do *something*, succeed in so doing, and, in so succeeding, also satisfy the requirements for *F*-ing.

In order to explain the fact that art-making, like Oedipus's wedding, can be intentional under some descriptions and non-intentional under others, I introduced Lopes's account of two kinds of non-intentional art-making, accidental and incidental art-making. While accidental art-making features the failure of an intention to G and no intention to F (to make art), incidental art-making was said to feature a successful intention to G, no intention to F, and the fact that Gs satisfy the description F.[59] There, I argued that a commitment to intention-dependence should preclude us from counting Lopes's accidental case as one of art-making, since it features a *failed* intention to G (and no intention to F or produce an F-equivalent product). What Lopes's account introduces, however, is the possibility of incidental art-making, which results in an artwork that is at least indirectly intention-dependent. Armed with such an account, the philosopher of art can maintain that art is intention-dependent despite the fact that much of art history seems to feature no explicit art-making intentions.

Finally, I showed that Mag Uidhir's language of attempt-dependence actually mirrors that of intention-dependence which I have been using so far. Where it differs, however, is in maintaining that the action must be guided by the intention that the goal obtain *in the manner prescribed*, a constraint that was dismissed when I initially discussed causal theories of action. The re-introduction of this constraint highlights a tension with respect to actual artistic practice and plan-following that will be a central focus of the next chapter. There, I will explore how this tension leads to a problem with the possibility of failure in general, and of failed-art in particular.

---

[59] Note that although this means that all Gs are Fs, the converse does not necessarily hold.

Chapter 3 – Failures of Intention and Failed-Art

3.1 – Introduction

"Failed-art" is a term coined by Christy Mag Uidhir (2010) to describe non-art entities that are non-art as the result of a failed art-attempt. Even though they are not artworks, failed-artworks are not trivially non-art in the manner of other artifacts or natural objects; the category is an informative one insofar as these entities' non-art status results from the particular way in which the art-attempt that generated them failed. Although we often wield art-status evaluatively in ordinary discourse, so that 'art' is synonymous with '*good* art' and non-art with '*bad* art,'[60] we ought to take more care with our language. Closer inspection reveals that we can usefully distinguish between *failed*-art and *good* or *bad* art. Good art is a subset of artworks which we judge favourably, and bad art is a subset of artworks which we judge unfavourably. Failed-art, however, is a subset of *non*-art. To say of bad art—as we often do—that it is a "failure" is just to double down on one's negative evaluation of the work and its execution of the goal one attributes to the artist. The category of failed-art, by contrast, describes those entities which exhibit an ontological failure to count as art. Although there is much that could be said about the mechanisms of evaluation and community agreement,[61] this chapter is occupied with a question much narrower in scope: what does failed-art look like, ontologically-speaking?

Mag Uidhir helpfully suggests a framework for thinking about failures, which he argues must satisfy three conditions:

> (1) Attempt Condition: An object *w* is the product of an *F*-attempt,

---

[60] Occasionally, scholars who ought to know better do so as well. Kuspit's (2004), for example, is an exercise in just such a conflation.

[61] See, e.g., Robert Brandom (1994, esp. Chs. 3 and 8), and D. Davies (2012), who teases out this thread in Brandom.

(2) Non-Art Condition: w is not[62] F in the manner intended, and

(3) Failed-Art Condition: w is not F in the manner intended as the result of the F-attempt in (1).[63]

With this general framework in place, it is a simple enough matter to generate a theory of failed-*art* by substituting 'art' for 'F' in the schema above. Mag Uidhir goes on to distinguish between two types of failed-art, *simple* and *complex* failures. Simple failures occur when the product w of an F-attempt is *non-art* (non-F) *simpliciter*;[64] so, for example, when the result of my attempt to chisel a bust is a heap of rubble rather than any kind of statue. Complex failures occur when the object w *is* "art"[65] (F; or, at least, accords with its creator's plan), but not "in the manner intended,"[66]—that is to say, not as a result of the proper following of its creator's plan (and, thus, is *not* art after all).

This chapter examines failed-art theory in more detail. I begin, in §3.2, by identifying what I take to be two fundamental senses of 'failure,' the *conformative* and *performative* senses. Armed with this distinction I will show, in §3.3, that an attempt's *direction* (i.e. whether it is directly or indirectly intentional) constrains the types of failure which are salient to it. I will argue that this means, among other things, that although the possibility of failure is a necessary corollary of art's intention-dependence, *pace* Mag Uidhir the possibility of failed-*art* is not. §3.4 attacks the non-art condition on the grounds that it falsely purports to be exhaustive, and that it places too much emphasis on 'the manner intended,' so that it threatens to label most actual artworks as failed-art. These considerations lead me, in §3.5, to carefully consider the complex

---

[62] In both his (2010) and (2013) Mag Uidhir phrases (2) and (3) such that "*w does not possess F*" in the manner intended. I have amended this to 'is not' throughout (including in condition (3)) to smooth the substitution of 'art' for 'F' later. I take this to be an entirely friendly amendment with no unfelicitous repercussions for Mag Uidhir's account.

[63] Mag Uidhir (2010: 394).

[64] Mag Uidhir (2013: 19-20).

[65] This formulation is a little misleading. It results from substituting 'art' for 'F' in Mag Uidhir's schema. The entity in question satisfies the success-conditions for an attempt to make an F, but it is actually *non-F* because it is not F in the manner intended, i.e. as a result of the maker's successful execution of her plan of action. We should thus say that the entity in question *seems* to be F, but in fact is not.

[66] Mag Uidhir (2013: 20).

hierarchy of intentions underpinning most art-making in an effort to articulate just what is meant by 'the manner intended.' Ultimately, I suggest that part of the problem stems from focusing our attention on failure as a property of artworks rather than of art-attempts. In §3.6 I turn my attention to amending the attempt-theory of failed-art so that it focuses on the attempts in question, rather than the products of those attempts. The result is a theory of failed-art that does a better job of capturing the realities of art-making.

3.2 – Two senses of 'failure'

Before getting into the details of failed-art theory, a few preliminary remarks are in order. Following a long tradition in action theory (the essential moves of which were outlined in Ch. 2), I assume that an attempt is just an action guided by the intention that a particular goal obtain in the manner prescribed.[67] Attempts thus consist of a goal and an intentional action (acting with the intention that *F*); attempts are successful when the goal obtains in the manner prescribed, and unsuccessful when the goal either does not obtain, or does not obtain in the manner prescribed. An *art*-attempt, then, is just any attempt whatsoever, the satisfaction of whose success-conditions entails the satisfaction of the conditions for being art. Some of these attempts will be deliberate, such as when Artemisia Gentileschi set out to paint Judith with the head of Holofernes. Others will be incidental, such as when prehistoric peoples painted cave walls, despite presumably having no concept of fine or "high" art.[68] Notice that what counts as an art-attempt is determined by what it is for something to be art. Thus, for an institutionalist, location in and acceptance by an artworld public would be required; for an intentional-historical account it would be the intention that it fit into some historical way of regarding art, and so on for other

---

[67] Just what falls under "the manner prescribed" is problematic, but must await a fuller treatment in §3.4.
[68] For a more detailed discussion of the different ways in which the success-conditions for art can be filled out, see Lopes (2007) and §2.4 of this monograph.

theories of art. The possibility of failure is dependent upon the success conditions of the enterprise in question, whatever these conditions might be; it is the epistemological role of a theory of art to supply just these success conditions.

There is, of course, a trivial sense in which art-attempts can fail: when, for example, an artist's creative activity is permanently disrupted. Should Amyas Crale formulate the intention to paint a portrait of his mistress but die before he can put brush to pigment then it is uncontroversial that he has failed to create an artwork. Should he die mid-process, we might have either an unfinished work that still counts as an artwork (like Coleridge's *Kubla Khan*), or no artwork at all (arguably, Duchamp's Woolworth Building is just such a case). Similarly, if someone intends to assert "P or Q" but dies after saying 'P,' her assertion is either incomplete, or it is not an assertion in the first place, despite its appearance to the contrary.[69]

While some unfinished works are nonetheless artworks in their own right, the question of interrupted activity invokes the problem of vagueness. These are artworks-in-progress which were never completed; how far along does the artist have to get before we are willing to say that an art-attempt has taken place, and how much more is required before we are willing to class it as a *successful* attempt resulting in an artwork? One possibility is that a mere intention suffices but this seems too inclusive, and too disconnected from any actual actions on the agent's part.[70] If, for instance, Hastings intends to paint a picture but never gets around to putting brush to canvas, then we do not have an art-attempt in the first place; the intention that a goal obtain in a particular manner is present, but the action is missing. The product (such as it is) will certainly be *non*-art, but it cannot be *failed*-art. A better option is to count as art-attempts only those

---

artistic plans which the agent is directly in the process of executing.[71] So, for example, going out to buy some paints and canvas would not constitute an art-attempt, but putting pigment to canvas would.

With that in mind, it should come as no surprise that one might fail to execute one's intention in different ways. One way, which I shall call the *performative* sense of failure, results from not fully or properly performing the requisite action. A second, *conformative* sense, results from not actually bringing about the desired consequences.[72] Suppose, for instance, that Hercule Poirot asks his audience whether they have any questions about the murder he has just solved, and that Hastings intends to ask a question. Suppose further that, in order to do so, Hastings must first raise his hand to signal that he would like to ask a question. If for some reason Hastings does not successfully raise his hand, then his attempt fails *performatively* because he does not actually perform the action required to fulfill his intention (viz. to ask a question). If he does raise his hand (performative success) but Poirot misreads, ignores, or simply does not see his signal, then Hastings has failed *conformatively* because he has not brought about the desired result. Performative failures result from uncompleted or wrongheaded actions, while conformative failures result from completed actions that do not bring about the desired consequences.

The extent to which some particular course of action is a success or failure thus depends on the answers to two questions: (1) did the agent do what she set out to do in order to bring about the intended results, and (2) did the agent's actions actually bring about the intended result? We often succeed in the former while failing in the latter, as in Hastings's case. But what

---

[71] For a defense of this kind of attempt liability (in the context of criminal law), see Sarah Paul's excellent (2014).
[72] David Davies pointed out this distinction to me in conversation, and does so again in his (forthcoming b). Davies's formulation is in terms of success/failure-in-agency and in-action; in order to avoid potential confusion (since both types of failure attach to actions) I have chosen to call the successes and failures in question *conformative* (in-action) and *performative* (in-agency).

if someone were to achieve the desired outcome without successfully performing the action intended? Suppose that Miss Felicity Lemon, somewhat uncharacteristically, intends to scare Poirot with a fake can of peanut brittle. Miss Lemon brings it to her office with the intention of giving it to Poirot, but fails to do so for one reason or another. Now suppose that, hunting for a snack, Poirot comes across the can in Miss Lemon's office much later and opens it up only to make the terrifying discovery that it contains a spring-loaded cloth snake. While the results *seem* to conform to Miss Lemon's intentions, the series of events leading up to them does not reflect a successful course of action that she undertook, since her earlier attempt to scare Poirot was aborted. Because the result does not issue from an action on her part, Miss Lemon is not fully responsible for it: Poirot's scare is merely accidental.[73] What this case serves to illustrate is that performative success is a precondition for conformative success. Exactly how a particular attempt fails, however, depends upon the success conditions of the attempt in question.

3.3 – Direction and the attempt condition

By distinguishing between failures of conformity and failures of performance, we are now equipped to see that the direct and indirect senses of intentional action (i.e. attempts) yield different types of failures. Courses of action undertaken directly-intentionally can fail either conformatively or performatively, while indirectly-intentional courses of action can only fail performatively.

Suppose, for example, that Hastings attempts the following:

Counting
Count (out loud) from 1 to 5, inclusively.

---

[73] As Nagel (1979) and Williams (1981) might put it, it is a matter of moral luck. See also §2.4.

Hastings's attempt here is directly intentional: he explicitly takes *Counting* as his goal. In order to succeed, Hastings must (1) begin counting, and (2) offer the sequence (1, 2, 3, 4, 5). If Hastings does not first begin counting (e.g. he intends to do so, but never gets around to it), then he has not acted in the first place and, therefore, cannot be said to have failed (or succeeded). Once Hastings actually begins to perform the requisite action we have a proper *Counting*-attempt, which means (given the distinctions introduced in §3.2) that his attempt can fail in either of two sorts of ways:

> (A) Because his attempt is aborted (i.e. it is abandoned or interrupted partway through) or goes awry.
> (B) Because the sequence offered does not conform to the goal governing the performance in (1).

If Hastings falls into trap (A) because he loses interest in *Counting*, then we have a trivial case of a failed *Counting*-attempt: Hastings attempted to satisfy *Counting*, did not produce a sequence satisfying *Counting*, and produced no result because he never actually completed the actions required to satisfy *Counting*. The failure here is purely performative. Likewise, if Hastings falls into trap (B) and sincerely offers the sequences (1, 2, 4, 3, 5), (1, 2, 4, 5), or (1, 2, 5), he will have failed in his *Counting*-attempt. The failure in this second case is conformative, meaning that Hastings will have successfully satisfied the attempt condition and generated a sequence of numbers, but that sequence does not satisfy his goal (it is $\neg F$). What is more, the product of his actions is $\neg F$ precisely because of the deviant way in which he went about trying to execute his goal (*Counting*). We thus have an instance of failed-*Counting*.

So much for the directly-intentional cases; now, what about those which are indirectly intentional? Suppose we modify the example so that rather than *Counting*, Hastings instead attempts

> *Stating*
> State all and only the natural numbers between 1 and 5, inclusively and in ascending order.

Once again, Hastings's attempt is directly intentional, since he is explicitly taking *Stating* as his goal. Its failure and success conditions will thus mirror those outlined for *Counting* above. What is special about *Stating*, however, is that it is satisfied by the same sequence of numbers as *Counting*, and vice-versa. The two differ intensionally, but the successful execution of the one entails the satisfaction of the success-conditions for the other, so we can describe direct attempts at the one as indirect attempts at the other.

The direct attempt constituting the indirect attempt can, as we saw above, fail in either of two ways (performatively or conformatively). But what about the co-extensive indirect attempt in each case? To the extent that Hastings counts by *Stating* his success or failure is incidental, since he does not attempt *Counting* directly. If Hastings fails to satisfy *Stating* then the product of his action (call it *P*) is bound to be a non-*Counting*, too, given that *Counting* and *Stating* have the same success conditions. To determine whether *P* is also a *failed-Counting*, however, we must determine whether the fact that it is a non-*Counting* is dependent upon Hastings's actual attempt.

This is clearly so when Hastings's attempt is aborted or goes awry: the reason *P* is a non-*Stating* (and, thus, a non-*Counting*) is because Hastings's attempt was somehow inadequate. We thus have a performative failure to both State and Count. But now suppose that Hastings 'fails' because he offers the sequence (1, 2, 4, 3, 5). Clearly, this is a conformatively failed-*Stating*. It is not, however, a conformatively failed-*Counting* because *Counting* was never the goal of Hastings's attempt. A successful *Counting* would have *coincided* with Hastings's goal, but it would not itself have been his goal. He might not even be aware that *Counting* is possible, after all! There is simply nothing connecting this description of the consequences of Hastings's actions to his intentions. The same will be true of any number of other equivalent descriptions of

Hastings's action—e.g. *Listing* ("list the numbers..."), *Mentioning* ("mention the numbers..."), *Setting* ("make an ordered set of the numbers..."), etc. So, while attempts which are indirect in nature might fail performatively, they do not fail conformatively.

On the weak reading of art's intention-dependence that I advocated in Ch. 2, art-attempts can be either direct or indirect. Their products, correspondingly, are either deliberate or incidental. We have just seen that although all attempts can fail, differently directed[74] attempts can fail in different ways. Direct attempts can fail both conformatively and performatively, while indirect attempts can only fail performatively. We can now discern two consequences for artworks. First, *pace* Mag Uidhir only the failure of a *direct* art-attempt can result in failed-art; the failure of indirect art-attempts just results in *non*-art. Second, these observations suggest that some non-art objects (especially failed-art) may sometimes be mistaken for art-objects,[75] and vice versa.[76]

None of this is to say that the formulation of the Attempt Condition above is problematic. On the contrary, the Attempt Condition is perfectly true, so long as we bear in mind that attempts can have different directions, and that differently directed attempts can fail in different ways. The important point to notice is that the weak reading of art's intention-dependence that I defended in Ch. 2 does not, by itself, entail that *any* artworks are strongly intention-dependent, or that the class of art-attempts includes any attempts to make art. If the only art-attempts in a world are indirect, then the only way in which they can be said to fail is performatively. If, however, we also have direct attempts, then these can fail either performatively or

---

[74] I use 'direction' of attempts to refer to whether an attempt is direct or indirect. The intended contrast is with *kinds* of attempts (i.e. whether the attempt in question is an *F*-attempt, a *G*-attempt, etc.), and with *types* of success and failure (viz. conformative or performative).

[75] Mag Uidhir's "complex" failures (2010: 394).

[76] Recall, e.g., the anonymous janitor at London's Eyestorm Gallery who, in 2001, threw out Damien Hirst's (untitled) pile of beer bottles, ashtrays, and coffee cups. This is a paradigm case of mistaking a putative art-object for non-art.

conformatively. At a minimum and *pace* Mag Uidhir, then, an intention-dependent theory of art must license the possibility of *performative* failure. The possibility of conformative failure is only raised when the class of art-attempts also includes direct attempts.

3.4 – Some problems with the non-art condition

We are now in a position to understand what goes wrong with the non-art condition, which states that in order to count as failed-art, a thing must be either non-*F* or *F* in an unplanned manner.[77] The problems here are twofold: (1) it purports to be exhaustive, and (2) the conditions placed on failures threaten to infect most art-making with that label. Let us begin with the first of these.

Recall that, according to Mag Uidhir, 'simple' failures occur when the product *w* of an *F*-attempt is non-*F*.[78] Simple failures are thus what I call "conformative" failures: the agent performs the requisite actions but the product of her attempt is non-art for whatever reason—that is to say, the product does not satisfy certain conditions external to the agent's activity which are necessary for its (intended) art-status. If the institutional theory of art is correct, for example, then this might be because the end result is deficient in some way that permanently precludes it from artworld uptake. Hastings might thus be so deceived about his artistic prowess that the end result of his efforts is wholly opaque and unintelligible to posterity.

'Complex' failures, on the other hand, are supposed to occur when the object *w* appears to satisfy *F*, but does not do so "in the manner intended."[79] This last fact is then taken to

---

[77] Given Mag Uidhir's commitment to the possibility of both complex and simple failed-art, the disjunction here must implicitly underlie his articulation of the non-art condition, even if the final version of his definition of failed-art does not make any explicit reference to the second disjunct. The presence of the first disjunct explains simple failures, while the second explains complex ones.
[78] Mag Uidhir (2013: 19-20).
[79] Mag Uidhir (2013: 20).

undermine *w*'s *F*-status in the first place. On the model I have suggested, complex failures should be classed as performative failures: the artist fails to successfully pull off the intended course of action, and this fact scuppers her artistic plans. The result is just non-art, though it might look like art. Because performative success is a precondition for conformative success (see §3.2), any pretence the product might have towards art-status is undermined by the failure of the attempt that preceded its production. Imagine, once again, that Hastings intends to ask a question and begins to raise his hand, but (perhaps because he feels intimidated) lowers it before Poirot can see his signal. Well aware of Hastings's chronic befuddlement, however, Poirot nevertheless calls upon Hastings to ask his question. Hastings has failed to perform the requisite action, and the fact that he gets to ask his question anyway is purely a matter of luck: it is not sufficiently closely tied to his intentional activity. The point is precisely that Hastings does not act in accordance with his goal, and so it is a mistake to think that he was conformatively successful.

Notice, however, that the failure of performance already secures the product's status as a failed-*x*; the issue of its uptake as art does nothing to specify the mechanism of failure. Thus, whereas "simple" failures pick out the entire class of conformative failures, "complex" failures merely name the subclass of performative failures whose products *seem* to be artworks but in fact are not. Another relevant subclass of performative failures would be those that do not enjoy uptake as art—that is to say, those which are not art (because the requisite actions were not performed, or not performed properly) and which do not fool anyone into thinking that they might be. Let us call these 'uncomplicated' failures to distinguish them from their complex counterparts.

Closer inspection of a complex failure will, as per the arguments of §2.4, reveal it to be the accidental product of an art-attempt, thus defeating its claim to art-status. While it is easy enough to imagine such cases in principle, it is a different matter entirely to identify such cases in practice. The task would be much simplified if one subscribed to a concept-dependent theory of art, since one might then point to the artifacts of other cultures as issuing from attempts lacking the guidance of the requisite concepts.[80] But if, as I argued in Ch. 2, we ought to abandon concept-dependent theories in favour of those whose intention-dependence is indirect, then the matter of identifying such failures becomes rather more fraught. Uncomplicated failures, on the other hand, are intuitively much more populous since they do not require uptake as art. They might include, for instance, Duchamp's Woolworth Building, the plan for which he formulated but never actually executed, or someone's abortive attempts to carve a figurine, paint a scene, or sculpt some marble. Their identification is likewise much more easily accomplished. The first problem with Mag Uidhir's characterization, then, is just that it presents the illusion of exhaustiveness. His "simple" failures name the whole class of conformative failures, but his "complex" failures only pick out one kind of performative failure, leaving the other kind (uncomplicated failures) unanalysed.

The second—and, I think, more serious—problem with Mag Uidhir's characterization of failed-art concerns its reliance on "the manner intended." Any account that places such a hefty premium on so vague a notion owes us some explanation of just what it entails, and no such explanation is forthcoming from Mag Uidhir. One plausible explanation is that "the manner intended" amounts to a requirement that the action be directly intended, since Mag Uidhir builds no appeal to "the manner intended" into his definition of art-attempts but makes explicit

---

[80] I shall have much more to say on the cross-cultural identification of art and art-kinds in Ch. 5.

reference to it in his account of failure.[81] This would be effectively the same move which was supposed to relieve the causal theory of action from the problem of wayward causation (see §2.2), and it is about as informative. Suppose that Hastings intends to create an action-painting by dropping a bucket of paint onto a canvas. Unfortunately, he is so distracted in his admiration of a passing sports car that he involuntarily loosens his grip and drops the bucket onto the canvas. The result looks like an artwork—in fact, it looks much as Hastings hoped it would!— but is actually *non*-art because the chain of events leading to its creation does not correspond to Hastings's intentions. In Mag Uidhir's terms, it is a piece of complex failed-art; in mine, a performative failure. The causal chain corresponds to the artist's practical reasoning, but the end product nevertheless seems unintentional.

To explain this discrepancy we could introduce (as Davidson does—see §2.2) a restriction to the effect that the intended goal must obtain "in the manner intended." In much the same way that Danto and later the New Volitonists tried to build intentions into the definition of intentional action by identifying a class of basic actions, it might thus be argued that Mag Uidhir identifies art-attempts as the fundamental unit of art-making. "The manner intended" should therefore be understood in terms of producing a result in a particular way, under the guiding hand of direct intentionality. With respect to the case above, Mag Uidhir might then say that although Hastings intended to produce a particular arrangement of pigment by deliberately pouring the bucket's contents onto the canvas, that is precisely what he failed to do and so the result is complex failed-art. Although this is a plausible line of response, the identification of "the manner intended" with direct intentionality risks setting conditions on art-making that are so strict as to be descriptively inadequate. Allow me to explain.

---

[81] cf. Mag Uidhir on art-attempts (2010: 387) vs. on failures (2010: 394).

3.5 – The manner intended

Thus far, my talk of the intentions underpinning art-attempts has been overly simplistic. I have been assuming (for the sake of clarity) that we were only dealing with one or two specific, goal-directed intentions. Artistic practice, however, is much messier than toy examples allow, and often involves processes of trial and error, revision and editing, and a fair amount of luck. As Michael Baxandall has put it,

> A static notion of intention, supposing just a preliminary stance to which the final product either more or less conforms, would deny a great deal of what makes pictures worth bothering about, whether for us or for their makers. It would deny the encounter with the medium and reduce the work to a sort of conceptual or ideal art imperfectly realized.[82]

Similarly, Denis Dutton has observed that "in the sense of using skill to produce a preconceived result, creative artists strictly speaking *never know what they're doing*."[83] The point here is not that artistic activity is random, but rather that the making of artworks—indeed, the performance of just about any action—requires the agent to make any number of discrete decisions, revisions, cancellations of previous actions or intentions, decisions not to do certain things, and so on. It also typically involves many different (sometimes competing) levels of intention: to make an artwork, paint a portrait, paint a photorealistic portrait, paint a photorealistic portrait of *X*, to paint a portrait that will change the artworld or which will infect its audience with melancholy, to use these brushes and those paints, etc.[84] As Baxandall observes, each of these moments has an impact on the final product; the artist's intentions undergo some amount of modification in the process of creating the work in question. The final work stands as a declaration, *ex post*

---

[82] Baxandall (1985: 63).
[83] Dutton (2009: 227-8).
[84] In their discussions of how to cash out a composer's intentions for the purposes of conforming to them for the sake of authentic performances of musical works, both Randall Dipert (1980) and Aron Edidin (1991: 413-5) have observed that we can actually discern several different levels of nested intentions. The point I am making here is that these observations should be generalized to an altogether new context, namely that of art-making as a whole.

facto, of the artist's intentions, but it also stands as one episode in "a serial performance of problem-stating and problem-solving."[85]

The goals of artistic practice, then, often do not obtain in quite "the manner intended"— at least, not in the sense of an overall prescription. The manner, rather, is negotiable at every step in the process of making. Just think of poor Apelles, whom Sextus Empiricus reports flew into a rage of despair when he found that he couldn't paint the foam on his horses just so.[86] In his fit of pique, he threw a sponge at his painting and thereby achieved the very effect he had failed to reproduce using all of his skill. The point of Sextus Empiricus's anecdote is just that Apelles *needed* a fortuitous accident to finish his painting as he wanted it. It strains credulity to think that the culmination of Apelles's efforts is a work of complex failed-art simply due to the minor element of chance involved in its creation—especially if he pulled off the rest of his painting without a hitch. When chance affects an element central to a work's identity, we might reasonably expect to call such a work failed-art because the strength of coincidence is such that without it, the action itself would be counted a performative failure. This is not the case with Apelles's painting, however. The problem with Mag Uidhir's model is just that it would have us overstate the impact of luck in Apelles's case: the amount of luck is too minor in the context of Apelles's action to infer that Apelles's work as a whole is either unintentional or a failure.

One could try to salvage Mag Uidhir's requirement by pointing out that although one's plan might have to change in response to particular challenges faced in the production of the work, or as one develops a more concrete idea of what it is one wants to do, one still ends up following *a* plan, and producing a work according to it. Plans need not be static things that entirely precede their author's action. In that case, however, there is nothing to stop us from

---

[85] Baxandall (1985: 66).
[86] Sextus Empiricus (1996: Bk. I, §12, p. 7).

describing *all* action as issuing from 'plans' formulated just moments before and in the course of acting. While such a solution has some merits, it still misses the point of Sextus Empiricus's anecdote: accidental results are part and parcel of the process of artistic creation. The stricter the requirements of "the manner intended," the less capable it is of handling such accidents without counting the whole work as accidental; conversely, the looser its requirements, the less effective it is in excluding deviant performances.

Part of the problem here stems from treating failure as a property of particular *works* rather than of particular *attempts*. There is something misleading in talking about failed-art, as though the entities in question exemplified the property of failure in the same way that an apple exemplifies redness or an orange roundness. Imagine we did the same for our talk of artworks, so that instead of 'art' we said 'successful-art' (i.e. the product of a successful art-attempt). The construction looks like it predicates success of 'art,' just as 'failed-art' looks like it predicates failure of 'art.' It thus runs the risk of being mistaken for an evaluative rather than an ontological claim. Appropriately flagged, this is a mostly harmless side effect: it is much easier, after all, to use 'failed-art' as a component part of a sentence than it would be to use 'product of a failed art-attempt.' Unfortunately, the locution lends itself to obscuring the fact that it is the *attempts* which do all the work here; the individual works just take the ontological credit. Success and failure are properties of *attempts* (or of the events that contain them), not of the products of those attempts (except analogically speaking). The products owe their existence to the attempts which generate them, and these attempts have their success- and failure-conditions set, as we have seen, by the nature (whether direct or indirect) of the goal which holds them together. Properly speaking, the failure in 'failed-art' attaches to the art-attempt, not to the art.

When we think of failure in terms of failed-art rather than failed art-attempts, it is easy to forget that we are actually giving a very general characterization of real-world practices. Indeed, thus far my talk of art-attempts has been at the general level: I avoided endorsing any particular combination of necessary and sufficient conditions for art other than intention-dependence, while recognizing that the actual application of this schematic treatment of intentionality to 'art' would require some such commitment to a theory of art. But nobody needs a theory of art to see that actual artistic practices are much more complex than the preceding characterization makes them out to be. This is why Mag Uidhir runs into trouble with "the manner intended."

The point of Baxandall's intervention is to show that our actual artistic practices are governed by a complex hierarchy of intentions, some more general in scope (e.g. to make an artwork), and others less so (e.g. to paint a portrait of *X*). More general intentions seem capable of surviving more changes than less general intentions, since their generality leaves them open to more ways of being satisfied. Changes at the level of general intentions entail changes to the more specific intentions falling under them—so, for example, if Hastings decides to sculpt Haraldr Harðráði rather than paint him, he must also change his intentions with respect to how he was going to represent him. Changes at the level of more specific intentions, however, need not entail changes at more general levels: if Hastings decides to paint an abstract rather than a photo-realistic portrait of Harðráði, then his intention to paint a portrait of Harðráði remains unchanged, as does his intention to generate a painting.

Baxandall's observations at the beginning of this section hint at a better explanation of what happens in Apelles's case: the finishing touch is an entirely accidental one but, once the effect was achieved, Apelles licensed it by incorporating it into his work *and letting it stand*. In so doing, Apelles declared his intentions—but the declaration was made *ex post facto*, not *ex*

54

ante. The question of whether some particular result conforms to one's intentions is not one that is necessarily settled beforehand, but rather one that is usually avowed after the fact. Planning is not a static activity that entirely precedes an agent's action; it is a dynamic process sensitive to the challenges of the occasion. Plans—especially artistic plans—develop as the agent pursues them, tests their limits, and overcomes the accidents and obstacles that crop up. What matters is not whether the final product is exactly as envisaged by the agent at the outset, but rather that its final form be linked to her goal by a chain of intentional actions.

3.6 – The attempt theory of failed-art

Where does all this leave failed-art theory? Recall Mag Uidhir's presentation:

> (1) Attempt Condition: An object *w* is the product of an *F*-attempt,
> (2) Non-Art Condition: *w* is not *F* in the manner intended, and
> (3) Failed-Art Condition: *w* is not *F* in the manner intended as the result of the *F*-attempt in (1).[87]

The first thing to observe is that conditions (2) and (3) are seriously flawed. We just saw that the non-art condition is too demanding in the first place. Human beings in general, and artists in particular, are not guided in their work by just a single intention, but rather by any number of overlapping intentions, some broad in scope and others far more task-specific. The more tasks that need to be accomplished in order to satisfy one's goal, the more intentions are required of a person. The result is that the non-art condition must be weakened to reflect the dynamic nature of art-making.

The non-art condition also faces another problem, however. The point of (2) is just to establish *w*'s non-art status (since *w* cannot be failed-art if it is actually art). Strictly speaking, however, (2)'s stipulation of "the manner intended" only allows for what Mag Uidhir calls complex failures (my performative failures); the schema neglects to capture all of the cases of

---

[87] Mag Uidhir (2010: 394). I have preserved here the felicitous modifications I introduced in §3.1.

failure which Mag Uidhir identifies, let alone any others which we might be able to discern. In order to make room for his simple (my conformative) failures, not to mention my 'uncomplicated' failures, we need to be able to countenance the possibility that w is not F *simpliciter*. That means either leaving the locution 'is not F' unanalysed, or breaking it down into a disjunction between conformative and performative failure (e.g. 'w results from an F-attempt that fails either conformatively or performatively').

Condition (3) faces a similar problem. The point of the failed-art condition is to relate the entity's non-art status to the way in which the attempt that generated it goes awry (viz. either conformatively or performatively). Condition (2)'s invocation of "the manner intended," however, performs the failed-art condition's work for it already: it ensures that w's non-art status is the result of the agent's incorrect performance of her F-attempt. So (3) is redundant.

Finally, notice that Mag Uidhir's attempt-theory of failed-art takes an *object*, w, as its proper subject: the salient questions, from (1)-(3), all pertain to w's properties.[88] Yet, as I argued at the beginning of this section, it is the F-attempt which does all of the work for which w takes the credit. F-attempts are a necessary component of Mag Uidhir's attempt-theory, but they are not its proper focus. We can avoid these pitfalls simply by re-casting the theory so that it is concerned with attempts, rather than with the entities that issue from them. Consider the following disjunctive schema for an attempt-theory of failed-art:

> *Failed-Art Theory A*
> (1a) An agent A initiates a direct F-attempt. (Attempt condition)
> (2a) A's F-attempt fails conformatively or performatively. (Failure condition)
> (3a) Because of the way in which A's F-attempt failed in (2a), the product p of A's actions is non-F. (Failed-art condition)
>
> or

---

[88] This focus on objects and their properties is curious, given Mag Uidhir's claims to offer us an *attempt* theory of art. These recur throughout his (2013) monograph, especially in the concluding remarks of each chapter, but his meaning is most clearly explained on pages 6-7. To be sure, attempts are a necessary component of his theory; it is just that his focus appears to be on objects and their properties rather than on the actual attempts that generate them.

Failed-Art Theory B
(1b) An agent *A* initiates an indirect *F*-attempt. (Attempt condition)
(2b) *A*'s *F*-attempt fails performatively. (Failure condition)
(3b) Because of the way in which *A*'s *F*-attempt failed in (2b), the product *p* of *A*'s
actions is non-*F*. (Failed-art condition)

For the sake of simplicity of exposition, we can aggregate theories A and B into a unified failed-art theory AB:

Failed-Art Theory AB
(1ab) An agent *A* initiates a direct or indirect *F*-attempt. (Attempt condition)
(2ab) *A*'s *F*-attempt fails. (Failure condition)
(3ab) Because of the way in which *A*'s *F*-attempt failed in (2ab), the product *p* of *A*'s
actions is non-*F*. (Failed-art condition)

The disjunction in (1ab) makes explicit the fact that not all attempts are the same, a fact which was lost in Mag Uidhir's failed-art theory. The reason we should make explicit reference to the direction of an attempt in (1ab) is that it is this direction which will determine the type of success or failure which attends it in (2ab). Because the object *w* has dropped out in favour of a focus on the *F*-attempt, we must also trade the non-art condition for a failure condition like (2ab). Mag Uidhir's non-art condition was too demanding, since it assumed that the direction of failure had to be direct ("the manner intended"), but only allowed for performative failure. With (2ab), that problem is remedied: the requirement is not that the product of the *F*-attempt be non-*F*, but rather just that the *F*-attempt fails. And as we know from the foregoing discussion, if *A*'s *F*-attempt was direct then that failure can be either conformative or performative; if indirect, then it must be performative. Finally, by removing 'the manner intended' from the failure condition, (3ab) is left to perform informative work: it tells us that in order to get failed-art, the product of *A*'s *F*-attempt must be non-*F* due to the way in which *A*'s *F*-attempt failed in (2ab).

In this way, we can explain Apelles's horses and give an exhaustive account of the different types of failed-art while preserving our commitment to the dynamism of artistic practice. Some failed-artworks are failures because something went wrong with the execution of

the art-attempt (performative failures), and some because of a mismatch between the actual and intended final products (conformative failures). So far as Apelles is concerned, we can discern at least three separate levels of intention in his story, all operating over different spans of time. At the most general level of description, Apelles's horses are the result of a directly intentional process (a process with an explicit intention to paint a picture containing horses as a component part) that takes place over the duration of his painting activity. Call this his picture-attempt (*P*-attempt). Because Apelles's *P*-attempt is direct, his efforts can fail either conformatively or performatively (still speaking at the most general level). At this level, however, we have no failure: Apelles ultimately successfully painted his picture and it looked as he anticipated it would.

Under the umbrella of his art-attempt, however, we can also discern a secondary direct intention to paint Poseidon's horses in a particular way; call this his *H*-attempt. Apelles's *H*-attempt only covers a portion of the time covered by his *P*-attempt (viz. the span over which he is concerned with painting his foamy horses). Apelles uses all his skill to render the horses (his *H*-attempt is performatively successful), but the end result is not quite as he wants it: the *H*-attempt is, at least initially, a conformative failure. I shall return to this attempt in a moment.

Finally, in the course of painting his picture Apelles undertakes a tertiary attempt when he flings his sponge at the offending picture; call this his *F*-attempt. Flinging the sponge is directly intentional, and is in fact successful: Apelles throws his sponge at the painting and hits it, as intended. The success conditions for Apelles's *F*-attempt (that the sponge hit the picture, or perhaps the horse-part of the picture) do not entail the satisfaction of the conditions for his *H*-attempt, so Apelles's *F*-attempt does not incidentally satisfy his *H*-attempt. The fact that the *F*-attempt resulted in satisfying the goal of his *H*-attempt is entirely accidental.

Why doesn't the conformative failure attached to Apelles's *H*-attempt travel up the line to infect his *P*-attempt? Quite simply, because the *P*-attempt does not end there! Apelles has not given up his *P*-attempt, he merely has a fit of rage that involves undertaking a new attempt. It remains open to Apelles to give up, to erase the accident and try again (perhaps with some new techniques), or simply to accept the result of the accident and let it stand as a retroactive declaration of his intentions, *à la* Baxandall.

To the extent that Apelles painted foamy horses by flinging his sponge at his painting, his success is accidental because Apelles does not attempt to paint his foamy horses by throwing his sponge at the work (i.e. the *F*-attempt is neither a direct nor an indirect *H*- or *P*-attempt). If his *F*-attempt had failed (e.g. because the sponge missed its target), Apelles would be left with his initial failed *H*-attempt plus a failed *F*-attempt, but no more or less. By itself, the failure of the *F*-attempt would have no additional effect upon the product of his earlier attempts—unless, perhaps, it caused him to abandon his *P*-attempt entirely (thus resulting in a performative failure).

Ultimately, the effect achieved follows from Apelles's plan because he changes his plan to accommodate it. Nor is this a unique result: Apelles had to change his plans a number of times in the course of executing his *H*-attempt, since he was never quite successful in achieving the intended effect. Mag Uidhir's model would have us class the product of Apelles's artistic activity as a piece of complex failed-art (or, at least, a complex failed *H*-attempt). It should be clear by now, however, that this is the wrong result because it asks us to evaluate the success or failure of an activity before its proper termination. It does not do justice to the many different attempts over time that we group together as Apelles's *P*-attempt.

3.7 – Conclusion

My aim in this chapter has been to explore the impact of failed-art theory, which seemed like a necessary corollary of art's intention-dependence. I showed that Mag Uidhir is right to argue that only the results of art-attempts have the potential to count as failed-artworks. And, so long as we are agreed that art-making is an intentional activity (as opposed to, say, a natural one), then he is also right that any intention-dependent theory of art must allow for the possibility of failure. By focusing his attention on the properties of the *objects* which result from art-attempts, however, Mag Uidhir unfortunately neglects the influence of an attempt's *direction*. The result, I have argued, is a flawed attempt condition, a non-art condition that both falsely advertises its exhaustiveness and over-invests in "the manner intended," and a redundant failed-art condition.

This oversight creates a number of additional problems for Mag Uidhir's failed-art theory. For one thing, Mag Uidhir thinks that the failure of an indirect art-attempt results in failed-art.[89] Armed with the distinction between different types of failure, however, we saw that indirect attempts can only fail performatively, thus resulting in *non*-art rather than *failed*-art. The products of such failures lack the intentional connection to the right goal they would need in order to count as failed-art. For another, Mag Uidhir overestimates the power of intention-dependence when he argues that any 'substantively'[90] intention-dependent theory of art must license the possibility of *simple* failure (i.e. what I am calling conformative failures).[91] Let us briefly turn our attention to this last case, to showcase the influence an attempt's direction can have.

---

[89] Mag Uidhir (2013: 28 fn14).
[90] The notion of 'substantive' informativeness does a great deal of work in Mag Uidhir's theorising. Its articulation, however, is surprisingly vague and context-sensitive: a necessary condition is substantively informative if it tells us something interesting about the thing under discussion, and trivial if it does not (2013: 4-6).
[91] Mag Uidhir (2013: 28).

Mag Uidhir illustrates his claim with an analysis of Jerrold Levinson's intentional-historical theory of art,[92] which he takes to claim that

> *w* is an artwork if and only if *w* is the product of an attempt to be regarded in way *r*
> where way *r* is a way in which pre-existing artworks are or were correctly regarded.[93]

According to Mag Uidhir, the success or failure of art-attempts does not matter for Levinson's theory; all that is required for art-status is that the entity in question be the product of an art-attempt.[94] So long as one non-passingly forms the intention to make art, then one cannot fail to make art.[95] The result is that, according to the intentional-historical theory, either all art-attempts are successful or their failure is irrelevant. Since both options are ridiculous, Levinson's intentional-historical theory cannot be intention-dependent.

We do not need to scrutinize the details of Levinson's theory to see where this analysis goes wrong. In fact, we need only recall from §2.6 that attempts consist of a goal and an intentional action (acting with the intention that *F*), a fact which Mag Uidhir himself endorses.[96] To form the intention to make art is to supply a goal, but not yet to provide an intentional *action*. Similarly, forming the intention to lie is not quite *to lie*, although it is certainly a necessary component of lying. As soon as one actually *offers* a work *w* for regard in any way that pre-existing artworks are or were correctly regarded (*r*), however, one's action might fail in either of two ways. It might fail because one does not do enough to motivate *w*'s falling under *r* (performative failure), or it might fail because one's actual offering does not conform to one's intentional goal (conformative failure).[97] Although Mag Uidhir is right to say that any account

[92] Levinson (1979).
[93] Mag Uidhir (2013: 28).
[94] Mag Uidhir (2013: 29).
[95] Mag Uidhir (2013: 30, fn. 19)
[96] Mag Uidhir (2013: 17-8).
[97] Perhaps, e.g., one is ambiguous about what it is that one is offering up for regard, or unclear in one's specification of the intended regard. (If one is offering a work for regard in a way that does *not* correspond to a way in which pre-

that identifies mere intentions to Φ with Φ-attempts will violate the requirements of intention-dependence, his own analysis of Levinson's theory demonstrates no such violation. Intention-dependence necessitates nothing more than the possibility of performative failure, which even the intentional-historical theory of art can concede.

The moral of the story is just this: the weak reading of art's intention-dependence which I defended in Ch. 2 requires, at a minimum, that art-attempts be indirect. This, in turn, means that art-attempts can at least fail performatively—if they are direct, however, then they can fail either conformatively or performatively. So any intention-dependent theory of art must at least allow for the possibility of *performative* failure. In practice, this is a rather weak requirement which most extant contemporary theories of art can easily satisfy. It is another question entirely whether intention-dependent theories must also allow for the possibility of failed-*art*. If my characterization of the effect an attempt's direction has on its failure-conditions is correct, then the answer is in the negative: only failed *direct* art-attempts can result in failed-art; failed *indirect* art-attempts will just result in non-art. I argued in Ch. 2 that the proper articulation of art's intention-dependence does not require of art-attempts that they be direct; this means that, *pace* Mag Uidhir, failed-art is *not* a necessary consequence of art's intention-dependence.

---

existing artworks are or were correctly regarded, of course, then one does not actually have an *r*-attempt in the first place.)

Chapter 4 – Medium, Kinds, and Conventions

4.1 – Introduction

David Clowney (2011) has recently mounted a defence of art's concept-dependence on the epistemic grounds that the alternative strains credulity. Even if we concede that our concept of 'fine art' is the product of a unique cultural and historical context, he argues, artists who are geographically and temporally distant from us must have had *a* concept of 'art'. Otherwise, we would be committed to the nonsensical claim that artists operating without a concept of art can make artworks without knowing what they are doing. The claim is nonsensical, Clowney thinks, for the following reason:

> At this point we might remind ourselves that it is impossible to make art without making music or painting or sculpture or poetry or one of the other things on that open-ended list. (To recognize the truth of *this*, we can depend on the ordinary use of the word 'art.') And it is impossible to make any of those things without doing so intentionally—that is, without having at least some idea of what one is doing. So premodern people did not make music, painting, and the like without knowing what they were doing.[98]

In other words, it is impossible to be unaware that one is creating art because one is deliberately intending to instantiate an entity that falls under an art-kind. Let us call this position *Clowney's Gambit*; if successful, then it sets the stage for thoroughly concept-dependent theories of 'art.'

There is a kernel of truth to the idea that, if art is intention-dependent, then one cannot make art entirely unintentionally. That kernel makes for a thin gruel, however. There are two ways of reading Clowney's Gambit. The first is just a straightforward espousal of direct intention-dependence (DID): there can be no art-making without the possession of the particular concept 'art' (or 'statue,' 'lithograph,' 'music,' etc.). If that is the claim, then it fails for the reasons already outlined in Ch. 2. As I argued in §2.3, the right way to think about this

---

[98] Clowney (2011: 312).

requirement is along the lines suggested by indirect intention-dependence (IID), since it is the only option on offer which captures the relevant range of art-historical phenomena. A more interesting possibility is inspired by IID: prospective artists must deliberately attempt to instantiate some particular set of properties, but they need not do so with the additional knowledge that this set of properties satisfies some other description(s). Talk about trying to make music or painting, in this context, would just be shorthand for trying to instantiate particular sets of properties which satisfy certain descriptions ('music,' 'painting,' etc.), without directly trying to satisfy those descriptions. As Peter Kivy has observed, there is no great mystery in the fact that one can use two pontoons to stabilize a water craft without ever possessing the concept of a catamaran.[99]

I will not dawdle further on the particular content of Clowney's remarks. Instead, I want to pick up on an idea prompted by Kivy's discussion of Clowney's Gambit: maybe the problem here is that we cannot produce necessary and sufficient conditions for 'art' in general, whereas it seems relatively straightforward to do so for individual art-kinds.[100] I noted in Ch. 2 that the success-conditions for particular kinds of attempts are parasitic upon one's theory of the kind of attempts in question, so that the aesthetic attitude theory of art will yield different success conditions from the institutional theory, and so on. It could be suggested, however, that I have gone about this the wrong way up until now. Perhaps the more fruitful strategy would be to focus my attention on art's constituent media, since the boundaries of their associated concepts seem much more readily accessible. As Dominic McIver Lopes observes, we seldom disagree about

---

[99] Kivy (2012: 72).
[100] Kivy (2012: 72).

whether the hard cases are *art*; what we balk at is their categorization under a particular art-kind.[101]

In order to counter Clowney's Gambit I will begin, in §4.2, by showing how some philosophers of art have suggested we might go about developing kind-centred theories of art. It will quickly become clear, however, that we cannot simply analyse art-kinds in terms of their constituent matter. In §4.3, I will follow Timothy Binkley's suggestion that the missing ingredient is the notion of *convention*. I will then sketch out the two main contenders for an analysis of convention, David Lewis's account in terms of coordination problems, and Ruth Garrett Millikan's 'natural' account. I ultimately cast my lot with Millikan's analysis of conventions in terms of historical precedence because of the relative paucity of plausible coordination problems at the heart of artistic conventions.

In §4.4 I apply the Millikanian analysis of convention to the history of art in an effort to flesh out the notion of 'artworlds' as collections of attitudes, beliefs, dispositions, and practices. §4.5 tackles the three main objections to kind-centred theories of art: the coffee mug, free agent, and structural objections. I argue that kind-centred accounts need access to conventions in order to address the coffee mug and free agent objections; unfortunately, they still founder on the structural objection, even if it is not entirely devastating to the kind-centred project. Finally, in §4.6 I show that kind-centred theories of art are as committed to intention-dependence as any other theory of art. What is more, I argue that such accounts need access to indirect intention-dependence if they are to appeal to the power of convention as a resource against the coffee mug and free agent objections.

---

[101] Lopes (2014: 71-4).

4.2 – Medium, art-kinds, arts, and appreciative kinds

Clowney is not alone in thinking that theories of art might be improved by a focus on art-kinds rather than on art in general. The idea was first proposed systematically by Dominic McIver Lopes, who in a 2008 paper[102] argued for the following 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.[103]

'The arts' here just designates the various art-kinds: media such as painting, sculpture, dance, literature, music, etc. Lopes's thought is that philosophers can develop definitions of art by going through the back door, so to speak: defining 'art' looks like a futile endeavour, so we should instead be aiming to give theories of the various arts. On this model, something is a work of art if 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 should suffice to tell us why a particular sculpture is art; the notion of 'art' in general has nothing more to contribute.[104]

Lopes's bottom-up approach should be distinguished from the more common, top-down way of understanding art-kinds, which reverses the order of explanation so that art-kind-membership is determined by one's theory of art. Christy Mag Uidhir (2013) has recently attempted a new defense of just such a structure, proposing that an activity such as painting can only count as an art-kind if the way in which something satisfies the conditions for being a painting also satisfies the conditions for being an artwork, whatever those might be. As he puts it,

> While no doubt media and practices play a role with respect to any robust analysis of art forms, the operative notion of what it is to be an art form must chiefly be in terms of works; i.e., the products of intentional actions and their corresponding sortal-descriptions.[105]

---

[102] Lopes develops these ideas further in his 2013 monograph.
[103] Lopes (2008: 109).
[104] Lopes (2008: 122).
[105] Mag Uidhir (2013: 95).

In other words, the suggestion is that we should look to the properties of, e.g., paintings to determine the requirements for membership in the art-kind 'painting'. What makes a painting 'art' is some property of the painted work by virtue of which it satisfies the necessary and sufficient conditions for 'art'. With respect to new art-kinds, the framework would work as follows. Suppose we want to determine whether 'video game' is an art-kind. Whether any currently existing video games are artworks is largely irrelevant; what matters is whether any video games can satisfy the necessary and sufficient conditions on 'art' by virtue of the way in which they satisfy the conditions for being a video game.[106] The trouble, of course, is that properly filling out the schema Mag Uidhir articulates requires us to first settle on a theory of art—a feat which philosophers have yet to manage. The prospects for Lopes's project thus seem brighter, since it only requires us to begin with descriptions of particular sortals rather than of 'art' in general. But just what *are* art-sortals, anyway?

The most obvious candidate is what we typically call art's '*media*,' i.e. kinds such as dance, literature, music, painting, photography, sculpture, etc. The trouble is that philosophers and non-philosophers alike often mean very different things by their use of 'medium.'[107] One common use of 'medium' is to group artworks together on the basis of their underlying material arrangement. So, for example, 'painting' is taken to designate the more-or-less two-dimensional application of pigment to a surface, usually canvas, wood, paper, or wet lime plaster, 'sculpture'

---

[106] Mag Uidhir (2013: 97).

[107] A terminological note is in order here. The present section discusses the development of the distinction between what Joseph Margolis (1980) and D. Davies (2010a) call a work's 'artistic medium' and its 'physical' or 'vehicular' medium. Mag Uidhir (2013), however, uses 'art form' and 'art sortal' interchangeably to designate a work's artistic medium; Lopes (2008 & 2014), on the other hand, identifies a work's artistic medium with its 'appreciative kind' and its vehicular medium with its 'medium profile.' Because these terms are all intended to capture the same notions, and in an effort to simplify comparisons between these accounts, in subsequent sections I will use the terms 'art-kind' (for artistic medium) and 'vehicular medium' (for physical medium), except in direct quotations. This move on my part follows Robert Stecker's suggestion (2010 [2005]: 284-5) that we distinguish between a medium (everything belonging to a physical medium) and an art-form (a subset of a medium, all of whose tokens are artworks).

refers to more-or-less three-dimensional manipulations of matter, usually in ceramic, metal, stone, or wood, and so on for the other media. When conceived in this way, many philosophers have treated a work's medium as largely incidental to the work's identity or value: a work's medium, they think, has nothing much of its own to contribute to the appreciation of that work. It is merely a vehicle for conveying the content an artist intends to communicate.[108] For expression theorists such as Croce (1922) and Collingwood (1938), for instance, the value of art lies in the act of artistic expression, not in the materials or techniques the artist employs to communicate that expressive content.

An equally common use of 'medium' tries to capture the thing which does the work of mediating between a work's artistic vehicle and its artistic content. So, for example, Monroe Beardsley has pointed out that 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). Our proper interest is thus not with the bodily motions considered in themselves (as the work's vehicular medium), but rather considered as 'movings' and 'posings.'[109]

There is a danger, in this connection, of according too much significance to a work's 'medium.' It is this tendency, for instance, which resulted in the elaboration of the 'doctrine of medium specificity' in the early twentieth century. Those who subscribe to the doctrine maintain that each medium is uniquely suited to the communication of different kinds of aesthetic or conceptual content. As Morris Weitz characterizes it, they argue that "Each art... has a specific

---

[108] Hence the use of the terms 'artistic vehicle' and 'vehicular medium' to refer to a work's physical medium. The intended contrast is to a work's 'artistic' medium, as we shall soon see.
[109] Beardsley (1982: 34-5).

function, which gives it its uniqueness; and this function is determined by the nature of the medium."[110]

Although the doctrine's origins trace back at least as far as Gotthold Ephraim Lessing's *Laocoön: an essay upon the limits of painting and poetry* (1905 [1766]), it came into its own in the artworld with Clement Greenberg. For Greenberg, the essence of modernism lay "in the use of characteristic methods of a discipline to criticize the discipline itself, not in order to subvert it, but to entrench it more firmly in its area of competence."[111] The doctrine held some sway even before Greenberg's landmark essay, however. Consider Rudolf Arnheim's infamous (and *Laocoön*-influenced) diatribe against the "talkies." Film and sound, he thought, work separately to express the same thing, thereby distracting the viewer and muddying the message. For Arnheim, any combination of media "must serve to express something that could not be said by one of the media alone."[112] As a rule, he thought that one medium must dominate the other and take the lead;[113] in fact, he thought it was precisely the limitations of a medium which made it aesthetically interesting, since they forced a medium to develop its own distinctive style.[114]

These philosophical defenses of the doctrine of medium-specificity stretch credulity for present-day audiences who are intimately familiar with the "talkies" and with mixed-media works. They do serve to highlight an important point, however: we cannot go on identifying 'medium' simply with a work's physical matter. As Lopes observes, painting is one of the arts,

---

[110] Weitz (1950: 120), quoted in D. Davies (2010a: 185). Note, however, that nothing in Lopes's account subscribes him to the doctrine of medium-specificity. The view that every artwork has a medium does not entail that every work belongs to one and only of the arts. See Lopes (2014: 17).
[111] Greenberg (1960: 1)
[112] Arnheim (1938: 215).
[113] Arnheim (1938: 223).
[114] Arnheim (1938: 229).

but this should not be understood to mean that every product of the activity of painting is an artwork.[115]

We can make a start by looking to Joseph Margolis's useful distinction between a work's *physical* and its *artistic* medium.[116] The physical medium is just whatever physical matter composes the artwork in question. The artistic medium, on the other hand, consists of "compositional ingredients that are themselves informed by the purposiveness of the entire work."[117] That is to say, a work's artistic medium consists of a series of intentionally-produced elements which are embodied in the work's physical medium.[118] In painting, for example, the physical medium would be pigment on some surface, while the artistic medium would be "a purposive system of brushstrokes—that is, *placed* or 'intended' colors and forms."[119] Sometimes, as with painting or sculpture, the artistic medium may be physically embodied; this is not necessarily the case, however. In dance or music, after all, the elements (dance steps and notes) are not quite as concrete as pigments and surfaces:[120] they are embodied in the performer's actions, in the perturbations of a listener's visual or auditory space.

A work's artistic medium does not quite supervene upon its physical medium, however: the relationship between the two is rather more complex. This is because the same matter can ground entirely different artistic media.[121] Consider Kendall Walton's infamous guernicas (1970): 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.

---

[115] Lopes (2008: 124). Lopes later calls this the 'coffee mug objection' (2013: 17). I will discuss this problem in more detail later, esp. in §4.3, §4.5, and §4.6.

[116] Margolis (1980). Note that although Beardsley (1958: 82) is skeptical that our talk of 'medium' can be sufficiently regimented to perform useful critical work, he appears to endorse a version of Margolis's distinction in his (1982) essay on dance.

[117] Margolis (1980: 42).

[118] See Michael Baxandall's remarks (1985: 66), which I discuss in Ch. 3 (esp. §3.5).

[119] Margolis (1980: 41-2).

[120] D. Davies (2010a: 182-3 & 2011: 11).

[121] The argument that follows comes from D. Davies (2004: 58).

70

Guernica's colours, figures, and the way they are depicted matter a great deal to its proper appreciation as a painting, while its topology does not. Evaluated as a guernica, however, the topology is everything while the colours and figures are largely irrelevant. Each artistic medium is tied to a different set of properties relevant to the proper appreciation of works in that medium, and the assignment of appreciation-relevant properties to an artistic medium is purely a matter of convention.[122]

If guernicas seem too far-fetched, then one need only look to music for an example closer to home. It is simple enough to imagine a group of Pythagoreans or Platonists for whom the interesting and artistically-relevant aspects of music have to do with its mathematical rather than its audible properties. Indeed, in the course of arguing that instrumental music is neither expressive nor representational, the 19[th]-century musicologist Eduard Hanslick felt compelled to add that music can rightly be appreciated for its mathematical form, though its beauty remains independent of that form.[123] Somewhat more contemporaneously, consider the culture surrounding the new musical subgenre known as "algorave," which focuses on live coding, and in which participants direct their attention to a screen displaying the code as it is generated by the performer.[124] Fractals can likewise be appreciated for either their aesthetic or their mathematical properties, as evidenced by the culture of appreciation surrounding the works of Jason Padgett, Vicky Brago-Mitchell, and Carlos Ginzburg.

---

[122] This is a far cry from the doctrine of medium-specificity rejected above. The idea is not that each artistic medium is specially suited to certain kinds of artistic content. The idea, rather, is that each artistic medium occupies a particular (contingent, conventional) role or roles in our cultural practices, and that these roles in turn reflect upon and inflect the way in which we engage with their associated works.

[123] Hanslick (1891 [1885]: 91). Although he is dismissive of the architectonic views of 'amateurs' and 'sentimental authors' like Hegel and Oerstedt, Hanslick nevertheless appears to think that mathematics has a key role to play in the *physical* study of music, and that music can be appreciated for its mathematical features so long as one does not make the mistake of thinking that mathematical thinking explicitly underlies composition, or accounts for music's *beauty* (1891 [1885]: 92).

[124] The subgenre is generally acknowledged to have been created (or discovered, indicated, or instantiated—as you please) by Alex McLean and Nick Collins, who gave their first public performance from the subgenre in 2011. See Cheshire (2013). I am indebted to Eric Murphy for bringing this subgenre to my attention.

What these reflections serve to demonstrate is that the notion of 'medium' is not simply tied to a work's physical composition. Something more is needed to explain just what a work's art-kind is supposed to be. In the next two sections, I will develop Timothy Binkley's suggestion that this 'something more' is inextricably bound up with a network of cultural *conventions* which act in concert to determine the ways in which certain ideas or aesthetic qualities can be conveyed.[125] Before tackling the way conventions operate in artistic practices, however, I must clarify what is meant by 'convention.'

4.3 – Conventions

I noted in §4.2 that any theory of an art-kind ought to be able to explain why some of the entities belonging to that kind are artworks while others which seem to share a vehicular medium with 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*.[126] In an effort to respond to the coffee mug objection, Lopes argues that art-kinds are fundamentally *appreciative* kinds,[127] which is just to say that they are the products of appreciative *practices*.[128] The exact nature of each of the arts thus depends in part on the value a particular community attaches to each of the members of the kind in question.[129] 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').[130] Artistic media—what Lopes calls 'appreciative kinds'—are not individuated by a single vehicular medium, but by what he calls a

---

[125] Binkley (1977: 269).
[126] Lopes (2014: 17).
[127] Lopes (2014: 121).
[128] Lopes (2014: 148).
[129] Lopes (2014: 130).
[130] Lopes (2014: 139).

'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.[131] 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.[132] So, it is the social practices from which bizen-yaki results which 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.[133] In the rest of this section, I will develop the notion of convention which I think is the most promising way of cashing out these social practices.[134]

On the 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).[135] For many coordination problems, 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

---

[131] See Lopes (2014: 140, 144, 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 (2014: 195); for a similar argument against singularity of medium, see Carroll (2003: 5-6). The notion of a medium profile conveniently sidesteps the problem of deciding how fine-grained we need to be about medium (see Carroll 2003: 7).

[132] Lopes (2014: 182).

[133] See Lopes (2014: 148).

[134] Inspired by Rawls (1955), Lopes gives a minimal characterization of a social practice in terms of conformity to certain rules so long as others who engage in the practice do so too. He goes on to say that "The rules need not be conventions, as they may not solve a coordination problem" (Lopes 2013: 148). Lopes clearly has Lewis's analysis of convention in mind; as I shall shortly demonstrate, I favour Ruth Millikan's analysis for reasons of descriptive adequacy. In particular, Millikan's model makes no appeal to rules of murky origin, and requires nothing so strong as conformity. Precedent can exert its weight indirectly from the cultural background, or directly from the agent's intentional stance. It is also a fairly lax criterion allowing for significant divergence and evolution; conformity to rules is not nearly so flexible. Finally, it owes no fealty to the current actions of other members of a culture.

[135] Lewis (2002 [1969]: 8).

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.[136] 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 right, I can reasonably expect other drivers driving my way to drive in the right lane, they can reasonably expect drivers other than themselves who are going their way to drive in the right lane, and they can reasonably expect that drivers other than themselves will expect other drivers driving their way to drive in the right lane, etc. I thus have good reason to drive in the right lane myself, since doing so is the best way to avoid a collision. The left lane may seem faster (perhaps it looks empty), but so long as everyone else going my way drives (or expects to drive) in the right lane, I would not be any better off by taking the left 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. One way in which recurring coordination problems can be resolved is through *salience*: one particular equilibrium stands out from the others for some special reason, e.g. because it's uniquely good, or because agents can reasonably expect one another to notice this particular equilibrium's uniqueness, etc.[137] Another way in which these kinds of problems may be resolved is through the force of *precedent*: a particular state of

---

[136] Lewis (2002 [1969]: 15). Note that Lewis stipulates that coordination problems must allow for at least two states of coordination equilibrium; otherwise, the coordination problem at hand is a trivial one (2002 [1969]: 22).
[137] Lewis (2002 [1969]: 35).

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.[138]

Interesting as it is, Lewis's analysis of convention falls a little short when it comes to the kinds of conventions governing artistic practices. This is because it is often difficult to see what kinds of coordination problems the conventions in question might have aimed to resolve. After all, what kinds of coordination problems could plausibly result in the conventions governing membership to the classes of painting or music? Similarly, as Wayne Davis (2003) has observed, most of the conventions governing the world of fashion do not seem to solve any coordination problems at all, since we do not care very much about how other people dress.[139] The convention that dresses and skirts are for women, for example, does not appear to solve any kind of historical coordination problem. So it seems that not all conventions arise as solutions to coordination problems.

Ruth Garrett Millikan's analysis of 'natural' conventions aims to describe exactly these kinds of conventions, as well as those involved in coordination problems. On Millikan's view, conventions need not be responses to coordination problems, although they may sometimes be; instead, they are patterns of behaviour that reproduce due to the weight of precedent.[140] Lewis's 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 of the participants:

---

[138] Lewis (2002 [1969]: 36).
[139] Davis (2003: 225).
[140] Millikan (2001 [1984]: 10). 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) 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. See Millikan (2001 [1984]: 3, 17).

"coordination conventions proliferate," Millikan argues, "because, rationality aside, people learn from experience exactly as other animals do."[141]

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. This is just because regularities of outcome which obtain from courses of action developed independently of previous outcomes are merely accidentally regular.[142] 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.

Conventions are thus fairly arbitrary creatures: different historical accidents would have yielded different conventions.[143] This in turn means that different populations are likely to have different conventions, so long as we adjust for cultural transmission, which works to unify conventions across cultures. 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. As we will see in §5.6, the Balinese appear to have music, but their musical tradition places a more substantial emphasis on its devotional and participatory elements. 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 musical conventions. The similarity of vehicular media, then, does not necessarily translate to a similarity of the appreciative practices governing the creation and use of that

---

[141] Millikan (2001 [1984]: 10).
[142] Millikan (2001 [1984]: 3-4).
[143] Millikan (2001 [1984]: 7).

vehicular medium.[144] What is more, even if two cultures share a practice with a similar medium profile and appreciative practice, that is no guarantee that the practice is an artistic practice in both cultures.[145] The influence of conventions thus motivates the conclusion, *pace* Walton, that art-kinds are not in principle pellucid to perception alone.[146] 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 the guernicas from §4.2: even though the proper way to appreciate guernicas is through one's senses, doing so is not enough to individuate them *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.).[147]

In response to Clowney's Gambit, it should be acknowledged that conventions are sometimes manifested (and thereby reinforced) deliberately, that is to say, agents sometimes consciously decide to instantiate existing conventions.[148] This happens when, for example, Hastings and Poirot play a game of chess, or when Parliament meets. In these cases, participants follow explicit sets of (procedural) rules which describe the conventions in question; in

---

[144] 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. For a more detailed discussion of this topic, see Ch. 5, esp. §5.6.

[145] D. Davies (2015) raises this objection in his review of Lopes's (2014). I shall have more to say on this issue in Chs. 5 and 6. For now, suffice it to say that our groupings of certain activities under the concept 'artworks' is ultimately arbitrary. Faced with the practices of other cultures, we must make a decision either to reject their art-status, or to accept it. If we accept it, we must further decide whether they fall into one of our existing art-kinds, or whether they motivate the development of a new art-kind altogether.

[146] Walton (1970, §II). To be fair, Walton does grant that some categories of art (such as literature) might not be obvious to perception alone, although he does think that most are.

[147] Lopes (2014: 71-4).

[148] This is just directly intention-dependent (DID) activity.

following these rules, participants reproduce the relevant conventions.[149] 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 imitate a given 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.

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.[150] Consider the widespread cultural practice of a woman taking her husband's surname upon marrying him. There is no good, 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 West, where it is adopted largely without second thought. 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 (note, however, that 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 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

---

[149] See Millikan (2001 [1984]: 14-5).
[150] Millikan (2001 [1984]: 8).

in Canada and 2015 in the USA, 'marriage' could only obtain between two adults of different sexes. 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.[151] As Noël Carroll so aptly puts it, "the medium does not fix the parameters of style, but stylistic ambitions dictate the production or reinvention of the medium."[152] The particular character of a convention depends upon the use to which we put it, on the aims of the relevant cultural practice.

4.4 – The development of 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 instance, 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 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.[153] Similarly, Ernst Gombrich pithily observes that "To say of a drawing that it is a correct view of Tivoli does not mean, of course, that Tivoli is bounded by wiry lines."[154] But conventions can also manifest as complex collections of attitudes, beliefs, dispositions, and

---

[151] Millikan (2001 [1984]: 16). See also Carroll (2003: 7-8) on the frequent alteration of art-kinds.
[152] Carroll (2003: 8).
[153] Dilworth (2003: 50).
[154] Gombrich (1984 [1960]: 73), quoted in Dilworth (2003: 61, fn 13).

practices; let us call such a collection a *world*. It is to this notion that I must now turn my attention, if I am to offer a satisfactory answer to Clowney's Gambit.

In defending his institutional theory of art, George Dickie (2000) posits that the 'artworld' is an untidy collection of culturally-constructed systems. 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. 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.[155]

We initiate our children by giving them paints and coloured pencils, instruments, lessons, and sheet music, by having them rehearse plays, read novels, and write stories in school. We reinforce these amateurish efforts by celebrating the fruits of their labour, by attending their performances and recitals. With practice, some of them progress to bigger 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. Dickie's artworld is a repository for many different kinds of cultural roles: artist and public, to be sure, but also critics, theorists, philosophers of art, curators, entrepreneurs, dealers, agents, promoters, and many more.[156]

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

---

[155] Dickie (2000: 101).
[156] Dickie (2000: 102).

years ago. It would be helpful, then, if we could see similar 'worlds' developing elsewhere, at a time for which we have access to fairly extensive records.

It is in this connection that Tiziana Andina (2013) offers her account of the development and decline of a 'tulipworld' in seventeenth-century Europe.[157] 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 centered on acquiring and wearing a beautiful but extremely rare tulip, the *Semper Augustus*, of which only a handful of bulbs were in existence. 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, and 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 1634.

In the span of twenty-five years, then, 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.[158] 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

---

[157] Andina (2013: 52-6).
[158] Andina (2013: 58-9).

as valuable for their worth as trading currencies or investments, not just for their aesthetic interest. Both worlds are thus populated by characters with similar functional roles, and develop markets characterized by speculative forces.[159] 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 to their functional role in that culture's social institutions.[160]

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 and development. Eventually the network of attitudes, beliefs, dispositions, and practices is 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* magazine. Institutions may only officially recognize or sanction certain kinds of cultural practices, but they are just a contingent outgrowth of worlds. 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.[161] What is more, they are neither static nor monolithic agglomerations of cultural practices. As Aron Edidin observes, our artistic, critical, evaluative, and interpretative practices—and sometimes even works themselves—change with the accumulation of more and newer works of art.[162]

---

[159] Andina identifies 'worlds' with markets and market forces. To my mind, this overemphasizes the importance of a small number of cultural roles. For one thing, the development of markets seems like a contingent rather than a necessary consequence of the establishment of worlds. For another, it seems just as likely that *markets* sometimes drive the development of worlds, rather than the other way around.

[160] D. Davies develops this idea for 'legal currency' in the context of a 'commerceworld' in his (2004: 248-9).

[161] In this connection, see also §5.5.

[162] Edidin (1991: 410).

With respect to the development of the tulipworld, it is worth noting that it is entirely arbitrary, at first, which kinds of flowers come to have so much social significance. It is also entirely 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. Tulips establish themselves as the kind of choice because they are the flowers that 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 (especially in terms of economic conventions and roles) in a very short span of time. 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.

4.5 – The trouble with kind-based accounts

Accounts proposing to analyse 'art' in terms of the arts face three main problems.[163] The first, which Lopes calls the 'coffee mug' objection (and which we saw in §4.2), stems from the fact that in ordinary speech the designation of a particular entity as falling under a 'medium'[164] carries with it the tacit assumption that the entity in question is an artwork. 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

---

[163] S. Davies offers versions of each of these objections in his (2008).
[164] Ordinary speech, of course, does not distinguish between a work's physical and its artistic medium. If it did, we would not be so ready to conflate kind-membership with art-status.

principled reasons to refrain from affixing the label of 'art' to all the works falling under a vehicular medium.

As we saw in §4.3, Lopes suggests that we look to our social practices for the answer to the coffee mug objection. On this score, Lopes is exactly right: only an examination of our appreciative practices will reveal why some ceramic objects are artworks, and others are not. Close examination of these practices (as in §4.4) reveals that they consist of several layers of nested conventions. The origins of these conventions are often entirely arbitrary (especially when they do not solve coordination problems), and are held together by the weight of historical precedent. 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 (was) not our custom to engage with coffee mugs and horned pansies in that fashion. Or, at least, not unless someone else goes (went) to the trouble of making them seem especially aesthetically or socially salient by, for example, declaring of an ordinary mug that it is a piece of found art, or that, henceforth, the horned pansy is the height of fashion. Coffee mugs certainly can be artworks, just as pansies can be fashionable, but our default assumptions about and attitudes towards them are that they are not, and we demand that claims to the contrary be substantiated somehow.

The second major problem facing an analysis of the arts is posed by what Lopes calls 'free agents,'[165] entities that are clearly artworks, but which do not seem to belong to any particular pre-existing art-kind or appreciative practice. 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 shared a bed (literally understood) in the designated period. Works such as *Everyone* defy the strategy of explaining 'art' by reference to theories of the arts because they

---

[165] Lopes (2014: 17).

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 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 conventions.

The other option, according to Lopes, is to assign the work to a 'new' art-kind such as conceptual art, whose medium profile is especially characterized by language or sets of ideas (especially about art itself).[166] 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 impressionism, post-impressionism, and aestheticism. 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. 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*

---

[166] Lopes (2014: 196-8).

Whistler's Mother. The work's artist-given title indicates that its proper subject is not the woman depicted (and her relation to the artist), but rather the formal properties which she embodies. These considerations suggest that the question we should be asking of Lopes is whether conceptual art's precursors were without an art-kind (and were therefore not artworks) until the 1960s and '70s, when conceptual art was conceived.[167]

Certainly not. Lopes quite rightly observes that the best interpretation of such free agents is as pioneering works:[168] 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 previous ones. 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[169]—Duchamp himself referred to it as a sculpture in a letter to his sister in 1917.[170] 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 are 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,

---

[167] There is, of course, no doubt that the painted works of Whistler, Kandinsky, and Picasso are *paintings*. My point here is simply to motivate the possibility that some works belonging to a kind might plausibly pre-exist the concretization of that kind. Whistler, Kandinsky, and Picasso thus set in motion the development of a series of conventions which make Duchamp's readymades possible. And, in fact, these readymades are plausibly interpreted as works of conceptual art existing half a century before conceptual art concretizes as an art-kind.

[168] Lopes (2014: 202).

[169] 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" (1996: 153).

[170] Naumann and Obalk (2000: 47). 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."

they must be art in virtue of belonging, however poorly, to some 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 contentious in the first place.[171] In other words, free agents show us that the kind-centred account flirts with circularity.

The danger in kind-based accounts 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. Mag Uidhir succumbs to this temptation when he observes that the "operative notion of what it is to be an art form must chiefly be in terms of works; i.e. the products of intentional actions and their corresponding sortal-descriptions."[172] This is exactly the move which leads us back towards the doctrine of medium-specificity. As I argued earlier, 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.

S. Davies puts the point especially clearly when he argues, with respect to architecture, that

> To assert that architecture is an art form just is to say that architects are artists who, as such, are bound to produce artworks, if not always of the best kind, in performing their role. Whereas, if architecture is not an art form, some architects might acquire the standing of artists and some buildings might be made to be artworks by their architects, but the architect is not an artist and does not produce art by performing her *standard* role.[173]

The mere fact that some *X*s are artworks does not suffice to make of *X* an art-kind; the missing ingredient is an appreciative practice centered on *X*s 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

---

[171] Recall Lopes's diagnosis from (2008: 120-1) and (2013: 73-4).
[172] Mag Uidhir (2013: 95).
[173] S. Davies (2007b: 136). Mag Uidhir appears to take S. Davies's observations on board (2013: 98), but somehow misses the point that what makes a sortal an *art*-sortal is rooted in a community's appreciative practices. Consequently, his focus remains exclusively on the works themselves, as on (2013:95). Note also that while Mag Uidhir cites S. Davies (2007b) on (2013: 98), the citation does not conform to Davies's actual remarks in his (2007), although it does capture the gist of S. Davies's remarks. A formatting error is probably to blame.

artwork. Rather, our interest is in whether the 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: because they are rooted in conventions, our practices are arbitrary affairs and may not yet have concretized to the point where there exists a strong norm compelling us to automatically classify a kind like architecture as an art-kind. Though there may be significant precedents, their accumulated weight may not yet suffice to delimit a fully-operational appreciative practice.

Kind-based accounts founder in the face of a third objection, however, which we can call the *structural objection*. Kind-based accounts 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 seem no better than they do for other kinds of theories of art. 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,[174] or not[175]? Is it abstract[176] or concrete[177]? Repeatable,[178] or singular[179]? A mental entity,[180] a historical individual,[181] an action,[182] or some kind of Platonic abstract entity[183]? The

---

[174] Levinson (1990), Kania (2011).
[175] Scruton (1997), Hamilton (2007).
[176] Wollheim (1968), Wolterstorff (1980 [1968]), Levinson (1990), Dodd (2007).
[177] Mag Uidhir (2013, esp. Chs. 4 & 5), Goodman (1968), Caplan and Matheson (2006).
[178] Dodd (2007).
[179] Hazlett (2012).
[180] Diana Raffman (1992), Collingwood (1938)
[181] Rohrbaugh (2003).
[182] D. Davies (2004).

literature on the ontology of music alone is every bit as extensive and fractious as that on the ontology of art more generally. To be sure, the structural objection is not decisive against kind-centred analyses: their prospects are dim, but not nonexistent. The real power of the structural objection is simply that it strips away much of the rationale behind the kind-centred project, even if it does not destroy it outright.

4.6 – Medium, Convention, and Intention-Dependence

Now for the crucial question: what does this detailed tour of kind-centred analyses of 'art' have to do with intention-dependence? The first thing to observe is that attempts to analyse 'art' in terms of art-kinds will mobilize exactly the same account of intention-dependence (ID) I offered in Ch. 2. A theory of painting, for example, will still take painting to be (at minimum) indirectly intention-dependent (IID). That is to say, in order for a work to count as a member of the class of paintings (*qua* art-kind, not *qua* vehicular medium), it will have to be the result of a painting-attempt. This is because, as we saw in §4.3, theories of art-kinds must have the resources to explain why only some of the entities in a vehicular medium count as art. Any theory of an art-kind which rejects ID will face extra difficulties when it comes time to explain why some non-ID entities are artworks, while others (some of which might even appear identical) are not.

Just as with 'art,' IID gives us a better characterization of the development of art-media than DID because it does not require concept-dependence. Recall that kind-centred analyses of 'art' are predicated on the principle (R), which holds that

> (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.[184]

---

[183] Kivy (1983), Dodd (2007), Levinson (1980), S. Davies (2001).
[184] Lopes (2008: 109).

If we assume DID for (R), then we will be unable to account for the first works in an activity P_n. This is because there would be no extant activity P_n prior to the work x_1, and so x_1 could never be intended as a work in activity P_n. What is more, if the accounts of conventions and worlds I gave in §4.3 and §4.4 are substantially correct, then DID-(R) would entail a substantial sacrifice of descriptive adequacy since, as we saw, conventions can exert their weight indirectly.[185] In order to intend to make an artwork, an artist operating under DID-(R) would first have to intend to make a work belonging to an art-kind, by virtue of which her creation would then count as 'art'.

It is historically unlikely that most art-kinds developed as the result of an individual's (or several individuals') direct intention to instantiate an object belonging to that kind. The more likely story, rather, is that people were interested in particular features of certain entities and that their interest was contagious, so that others were interested in replicating the same kinds of features in entities of their own. The mimicry may well have been directly intentional, but the instantiation of entities belonging to a particular kind need not have been. Consider this 'Just So' story by way of example: Australo-Hastings finds a supply of nuts. He knows a tasty treat lies inside the hard shells, but does not know how to get at it. One day, he sees Australo-Poirot bashing a nut's shell with a stone. A famished Australo-Hastings thus picks up a stone of his own and uses it to smack a nut until the shell gives way. Australo-Poirot *may* have the concept of a tool (or whatever), but he certainly need not have it (yet); Australo-Hastings probably does not, however. Foremost in Australo-Hastings's mind is the need to get to the nut: he directly intends to mimic Australo-Poirot's behaviour, but only incidentally does he develop a tool of his own.

---

[185] That is to say, since one need not explicitly conceive of oneself as following a precedent in order for that precedent to exert its influence. If, for example, that precedent is so well established that no other options present themselves to the agent as salient, the agent will act in accordance with precedent as a matter of course, as in the case of matrimonial surname change in §4.3.

As other Australopithecines follow suit, the kind 'tool' (or whatever) concretizes, and is eventually separated out from the activity of bashing nuts (allowing, e.g., for the attachment of a handle to increase efficacy).

All of these considerations weigh in favour of IID for (R), so that an artist need not have any concept of the art-kind to which her work ultimately belongs.[186] Under the influence of IID-(R), she can go about her work by intending to create 'art' in general, without ever having to worry about the work's proper kind. The question of its kind need not be settled beforehand; it can (and should) be determined after the fact of its instantiation, once its various properties have been established. One last example will serve to make the difference clear. Suppose we accept the usual story about *Fountain* (1917), according to which it is a work of conceptual art. DID-(R) would require Duchamp to think of what he does as instantiating conceptual art, not sculpture, despite the fact that conceptual art did not yet exist as an art-kind. IID-(R), by contrast, allows for Duchamp to *think* he is instantiating sculpture even if it later turns out that what he actually instantiated is a work of conceptual art.

At the general level, intention-dependence has another, more important role to play in the analysis of 'art' in terms of art-kinds: it supplies the framework of an answer to coffee mug-style objections. Imagine a simple coffee mug (non-art) and an identical-looking piece of bizen-ware (art). By appealing to intention-dependence, we can explain why the bizen-ware is art (assuming it is) but the ordinary mug is not: the bizen-ware is the result of a direct or an indirect art-attempt, whereas the ordinary mug is not. This is not to say that it can never be art—it was, after all, produced by an intention-dependent process, so it at least satisfies that necessary condition—just

---

[186] Lopes (2014: 78) agrees, although he seems to think this means that under kind-centred analyses, 'art' is not essentially intention-dependent at all, since not all artworks are necessarily products of an intention to make entities falling under that particular kind. It should be clear, at this point, that this only follows if we fail to distinguish between DID and IID, since IID models intention-dependence for just such cases.

that, on its own, it is non-art. This is obviously not a full answer to the coffee mug objection, but it is perfectly tractable as a partial answer. The rest can be filled out, as Lopes suggests, by looking to our appreciative practices and to our artworld(s), which supply us with a genetic basis on which to ground the distinction. In brief, 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, in point of fact, it does not. We can take things a step further by analysing our appreciative practices in terms of networks of conventions rooted in the weight of precedent.[187] Although doing so ultimately does nothing to diminish the arbitrariness of the (collective) decision to count something as an artwork or not, it does at least supply an explanation for that arbitrariness which does not preclude the possibility of the mug's counting as art under counterfactual conditions.

But, it might be thought, couldn't we eventually develop an appreciative practice centred on natural objects and processes, which would then underpin an art-kind characterized by intention-*in*dependence? Worlds are complex beasts, after all, and the conventions underpinning them are arbitrary and capricious. It would not take very much for someone to offer a sunset up for aesthetic consideration. Indeed, this is just what Europeans did under the influence of Romanticism. Romanticism's emphasis on emotion, the sublime, and the glorification of nature gave us a view of the arts as centring on the production of aesthetic experiences. In the visual arts, this began in the 1760s with a revolt against the burgeoning industrial revolution. The result was the preponderance of landscape painting, which came to feature ever more turbulent vistas

---

[187] In a similar vein, Sherri Irvin (2012) has argued that we distinguish between at least some art-kinds (viz. installation art and performance) on the basis of the sets of conventions underpinning their production and appreciation. The result, she suggests, is that it might prove more (ontologically) profitable to treat artworks as sets of norms or instructions propagated by means of historical precedent.

and reached its zenith in the mid-19[th] century.[188] Romantic literature likewise reinforced the cults of sensibility and nature by taking these as principal motifs,[189] which in turn motivated the aesthetic theories of the German Idealists.[190]

But even though the Romantics were preoccupied with aesthetically appreciating intention-independent phenomena, they did not take these to be artworks in themselves. Rather, they glorified them as *subjects* of art; the *artworks* were still the (intention-dependent) products of their actions. This example should serve to show that the mere existence of an appreciative practice does not suffice to distinguish its objects as constituting an *art*-kind. Not all appreciative practices, after all, pick out *art*-kinds; the question is, what makes a particular appreciative practice *artistic* as opposed to something else?

The answer, I have argued, is to be found in the complex network of conventions from which that practice issues. The set of practices that count as artistic is a contingent one with historically variable membership, but the fact remains that, in the actual world, ID is a necessary feature of these practices. This is not to say that future generations might not start counting natural entities and processes as falling under the extension of 'art'. The point, rather, is that they would have to overcome a heavy chain of historical precedents to do so. So heavy, in fact, that we should say that they were talking about something else entirely: it would be schmart, not art.[191] It is a difficult thing to delimit the boundaries of a social kind since, unlike with natural kinds, we cannot appeal to an underlying essence. Doing so calls for a decision on our part, a process for which I will suggest a model in §5.6. My contention here is just that ID supplies us

---

[188] Think, e.g., of the works of Francisco Goya (1746-1828), Caspar David Friedrich (1774-1840), Ivan Aivazovsky (1817-1900), and J.M.W. Turner (1775-1851).
[189] Think, e.g., of the works of William Wordsworth (1770-1850), Samuel Taylor Coleridge (1772-1834), Ralph Waldo Emerson (1803-82), and Henry David Thoreau (1817-62).
[190] Chief among them Kant (1987 [1790]), Hegel (2001 [1835]), and Schopenhauer (2010 [1818] & 1966 [1818]), all of whom locate the value of art in the aesthetic experience it engenders (albeit in very different ways).
[191] I will return to these issues in Chs. 5 and 6.

with more stable grounds for doing so than most others we might dream up. And so long as kind-membership is a matter of convention, then IID is the sense of ID that is at stake. Anyone is perfectly welcome to an intention-independent or concept-dependent theory of art or the arts, should they want one. The challenge in that case will be to explain their revisionary commitments in such a way as to justify the loss of descriptive adequacy when it comes to describing our *actual* appreciative and artistic practices.[192]

4.7 – Conclusion

I opened this chapter with Clowney's Gambit, which attempts to construct a case for art's concept-dependence by appealing to an analysis of 'art' in terms of art-kinds. The task of refuting Clowney's Gambit required me to sketch out the frameworks which have been offered for kind-based theories of art, as well as the two main philosophical analyses of convention. The primary 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.[193]

I have argued that Lopes is right to look for the answer somewhere in the vicinity of our appreciative practices, but that his framework is missing 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 somewhat arbitrarily, not just on the basis of their vehicular media but also on the basis of fit into extant and historical artistic and appreciative practices. As

---

[192] I will take up the dispute between descriptive and revisionary ontologies of art in more detail in Chs. 5 and 6. Ultimately, however, I think that the price of adopting an intention-independent ontology is just too high, especially since the intention-dependent ontology I am defending has the conceptual resources needed to explain the conditions under which intention-independence could have come to characterize our appreciative and artistic practices. The proponent of intention-independence, by contrast, would need a full-blown error-theory to account for our current embrace of ID.

[193] See, e.g., Young (2016). This is effectively just another version of the coffee mug objection.

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 kind-based conventions, which in turn play an important role in the composition of an 'artworld'.

What I am offering is the framework for an analysis of art-kinds in terms of systems of conventions. It is an admittedly vague framework, since I have only gestured to the particular conventions underpinning individual art-kinds. But it differs from existing frameworks in an important way: it places responsibility for kind-membership squarely on the actions and systems of convention governing a work's contexts of creation and appreciation, rather than on properties of the work itself.

The upshot of all this is that it is not artworks or even art-kinds themselves which do the hard work of securing the work's art-status. That work is done by our artistic and appreciative practices, which are complex systems of conventions governing, among other things, our appreciation of, beliefs about, dispositions toward, and interpretations of artworks. The art*works* themselves are largely irrelevant, except insofar as their bare existence is required. Similarly, art-*kinds* are no more than a heuristic device wielded bluntly, an unreliable shortcut to art-status. The real work is done by the conventions which supply the necessary pretext for ascriptions of art-status and kind-membership, and which offer a handy roadmap to proper appreciation. Ontologies of art have tended to overemphasize the importance of the work itself, which is just a trace of the action that generated it and derives its art-status as an echo of the weight of actions that came before. The kind-centred accounts I sketched in §4.2 run the constant risk of overemphasizing the role of the physical object at the expense of the long line of actions from which it results, that is, at the expense of the conventions governing our artworld-practices.

Of course, the fact that the analysis of art requires this sort of emphasis on history and sociology leaves it vulnerable to critiques targeting the arbitrary and exclusionary nature of that history and sociology.[194] The next chapter sees me turn my attention to such critiques with an eye to determining whether we can be justified in appealing to descriptive adequacy as a constraint on theory choice in the ontology of art.

---

[194] Lopes makes this very point on page 119 of his (2013).

Chapter 5 – Two Approaches to Descriptivism

5.1 – Introduction

The previous chapters have been primarily concerned with charting the consequences of taking art's intention-dependence seriously. I argued in Ch. 2 that this means distinguishing intention-dependence from concept-dependence, and rejecting the latter as too restrictive of artistic practices. In Ch. 3 I showed that a descriptively adequate treatment of art-attempts requires us to distinguish between the conformative and performative senses of success and failure, and to focus our critical attention on the success/failure of art-attempts rather than on the properties of the resulting entities. Finally, in Ch. 4 I tied my account of art's indirect intention-dependence to an independently attractive treatment of the genesis of art in terms of Millikanian conventions.

In each of the previous three chapters, I have argued for the profitability of turning our attention away from the properties of artworks themselves, and towards the actions that generate them and the network of social practices in which they occur. In this chapter and the next, I want to shift the focus of my attention to the methodological question of how we can ground claims about social kinds like 'art'.

In the philosophy of art, this debate takes the form of a dispute between descriptivists and revisionists over the subfield's proper methodology. The dominant approach, descriptivism, is concerned to describe the actual structure of our thoughts about art. Accordingly, it takes those thoughts as the standard by which theories of art are to be adjudicated: the more of our thinking a theory captures, the better. In a turn to what I call concept-driven ontology, some descriptivists have even gone so far as to argue that competent users of art-kind terms cannot be mistaken

about the ontology of 'art' since the concept's content is dependent upon how and what we think about it in the first place.

The alternative, revisionism, has not proven especially popular among philosophers of art, although it is now the default position in metaphysics more broadly. Revisionists caution that our conceptual schemes are mere accidents of history, and thus advocate deference to general metaphysical and scientific principles. According to revisionists, we ought to defer conceptual analysis until we understand the kinds of entity at issue. At that point, we can analyse the relevant concepts under the guidance of our best metaphysical and scientific theories, deferring to these when our results clash with our intuitions. In contrast to descriptivists, revisionists argue that the right answers to ontological questions are not determined by what we say or think about them. Consequently, our common-sense views about art may well be substantially mistaken, and should not figure prominently in our investigations of that concept.

This chapter aims primarily to develop a case against the concept-driven approach to ontology that does not rely on the sorts of considerations I introduced against concept-dependence in Ch. 2. If I am successful in doing so, then the end result should be an added reason to reject the notion that art is concept-dependent, and to favour instead its indirect intention-dependence.

I will begin, in §5.2, by considering the case for a concept-driven descriptivism, focusing on Amie Thomasson's landmark defense of that approach. The next three sections, §5.3, §5.4, and §5.5, will build a case against descriptivism as a concept-driven ontology. In §5.3, I will introduce the problem of 'art's' conceptual instability, which gives us reason to doubt that 'art's' content is stable across the uses to which it was put in the history of our own culture, let alone across the uses of other cultures. If concepts are a guide to ontology, then 'art's' history offers us

an embarrassment of riches. Against the response that what we are interested in is the concept that competent users of the terms in our present culture deploy, §5.4 will present the problem of conceptual imperialism. 'Conceptual imperialism' names the tendency on the part of even competent users of 'art' and art-kind terms to assimilate unfamiliar practices under the banner of familiar concepts without regard for the actual features of those practices—or, conversely, of exaggerating the differences between similar practices. The problem, then, is that even competent users of a term seem susceptible to substantial errors when analysing the practices of other cultures, resulting in mistakes about the concept's extension and functional role. §5.5 will introduce the problem of conceptual inclination, which concerns the influence of artworld precedents on our nascent ("folk") ontologies of art. I argue that this gives us further reason to doubt that our intuitions about art and art-kinds actually reflect the concepts' ontological content rather than conventionally salient features of our thinking. Finally, in §5.6 I will suggest an alternative descriptive strategy, the practice-driven approach to ontology, which takes our artworld practices—not concepts—as its starting point. What is more, I will illustrate the practice-driven approach's plausibility by applying it to the problem of the cross-cultural identification of art, focusing in particular on the case of Balinese *mabarung*.

5.2 – Descriptivism as concept-driven ontology

The distinction between descriptive and revisionary ontologies owes its origins to Peter Strawson's *Individuals* (2003 [1959]), which aimed to explain the human ability to refer to individual entities in the world around us. During the course of his investigation, Strawson found it necessary to distinguish between the descriptive and revisionary approaches to metaphysics. Descriptive ontologies, he thought, are primarily concerned to describe the *actual* structure of

our thoughts about the world, while revisionary methodologies aim to offer a better structure.[195] The main method of descriptive metaphysics is conceptual analysis: according to descriptivists, our metaphysical theories are answerable to our judgements about our individual practices, which supply the data points our theories aim to explain. Revisionary methodologies, on the other hand, begin with our category-judgements and use these as the basis for judgements about individuals.[196]

A theory of mystery novels, for example, might begin by canvassing several novels in the genre and conclude that these works are characterized by a "whodunnit" format—viz. throughout the story readers are offered clues as to the identity of the perpetrator, who is only revealed at the story's conclusion. As a descriptive proposal, the theory must be judged on its ability to capture the way we think about mystery novels. And on those grounds, the whodunnit theory of mystery novels runs into some trouble when faced with the many sub-genres of mystery that are *not* characterized by a whodunnit structure: courtroom dramas and legal thrillers, cozy mysteries, historical mysteries, inverted detective stories, police procedurals, serial killer mysteries, etc.[197] The theory appears extensionally inadequate. Alternately, it might begin by considering the category of mystery novels, and find that there is an essential connection here to the whodunnit format. As a revisionary proposal, the theory must be judged on its added value, on the tidiness it brings to our theoretical apparatus. On these grounds, it owes us an explanation of the apparent variety of mystery sub-genres, which it might explain away as variations on the theme of

---

[195] Strawson (2003 [1959]: 9).
[196] See also Kania (2008).
[197] My understanding of the historical relationship here is that whodunnits helped to establish mystery as a literary genre (although the genre's roots reach back to the 19th century). Other subgenres of mystery novels arose largely as the result of attempts to subvert the genre conventions established by whodunnits in the golden age of detective fiction, resulting in the dramatic expansion of the genre.

whodunnits. In this way, the theory would have consolidated several different structures into a single one, thus tidying up our ontology of mysteriana.

I will have a great deal more to say about hybrid and revisionary methodologies in the next chapter. For the time being, however, I will focus my attention on the suggestion that our collective thoughts about the artworld reflect a nascent ontology whose content it is our job to explain. In this way, our intuitions give the rule to our metaphysics and constrain the results of our ontological investigations. This strand of descriptivism has been most carefully defended by Amie Thomasson, for whom the motivation seems to be the fact that 'art' is not a natural kind with a real essence, but a socially-constructed kind (perhaps with a nominal essence). For Thomasson, the reference of art-kind terms is fixed by the ontological properties that competent users of the terms explicitly or implicitly ascribe to them.[198] In this way, the reference of any art-kind term must be grounded in a "nascent" concept of the ontological kind involved.[199] "So," she tells us,

> it seems that in order to unambiguously ground the reference of a general term to name
> a kind of work of art, the grounder must not only have the idea that the reference of his
> or her term will be an art-kind, but must also have a background conception of what
> ontological sort of art-kind he or she means the term to refer to, establishing existence
> conditions and identity conditions for works of that kind.[200]

Grounders need not have a fine-grained analysis of the ontological sort in mind, but some concept of the kind in question is required in order to situate it against a background of practices which are already in place.[201] So, for instance, if we want to find out what paintings are, we must look to the ways in which competent users of the term use it—namely, to refer to instances of the

---

[198] Thomasson (2007: 189-90). Thomasson's argument is actually about kind-terms in general, not just art-kinds. I have taken the liberty, here, of adapting her framework and arguments to the particular case of 'art.'
[199] Thomasson (2005: 223).
[200] Thomasson (2005: 225).
[201] Thomasson (2005: 225).

higher-level sortals 'artwork' or 'visual artwork' involving the application of pigment to certain surfaces (esp. canvas, plaster, and wood).

On this model, the ontology of art is just an exercise in conceptual analysis, in discovering the term's proper use and the situations to which the term is properly applied (its application conditions[202]) as well as those in which the term can be properly applied again to one and the same entity (its co-application conditions[203]). Metaphysical investigation of our kind-terms—including art-kind terms—*cannot* revise the conceptual content that ordinary users of a term confer upon it, because they must begin their investigations with a concept already in hand. We cannot err substantially about the ontological properties of artworks, since in doing ontology we are pulling a rabbit out of the very hat we put it into in the first place. The task of metaphysics is just to articulate the content of these concepts, and to unravel the commitments they enjoin us to adopt. The rest of this chapter is devoted to building a case against this kind of descriptivism on the grounds that the history of art actually supplies us with good reasons to think that the nascent ontological concepts we deploy are often incomplete or wrong.

5.3 – The problem of conceptual instability

The first problem with this kind of descriptivism is what I call *the problem of conceptual instability*: the anthropological and historical evidence indicates that the concepts underpinning the various practices we regroup under the banner of 'art' are not stable across time. Our artistic practices have changed a great deal over time, not just in terms of techniques but also with respect to their functions and the place they occupy in our societies. This spells trouble for the

---

[202] Thomasson (2007: 39-40).
[203] Thomasson (2007: 41).

concept-driven approach to ontology because it proliferates the sources of nascent ontologies, and provides no guarantees that these will be consistent with one another.

The case for 'art's' instability begins with two hugely influential essays by Paul Oskar Kristeller, who argued that the modern "system of the arts" emerged in the mid-eighteenth century, through the work of Charles Batteux.[204] According to Kristeller's story, the ancients had concepts of *techné* and *ars*, but these did not specifically denote the "fine arts" and instead included all manner of craft or scientific activities.[205] The mediaevals likewise had no concept of the fine arts and persisted in applying 'art' to crafts such as shoemaking, juggling, and arithmetic.[206] This conflation persisted through the Renaissance, until the modern writers began to group some artifact-making activities (especially painting, sculpture, and architecture) together and distinguishing them from crafts. Still, it was Batteux who first "correctly" grouped together the arts that define our system of the arts today: poetry (including literature), painting, sculpture, music, and dance. These he grouped together on the grounds that their primary goal is pleasure, and from them he distinguished the "mechanical" arts, which serve practical needs such as food and shelter, and a third category combining the two.[207]

Although Kristeller's is still the more or less orthodox view among art historians and philosophers of art, there is good reason to suspect that it is just another "Just So" story. For one thing, as James I. Porter has pointed out, alternative analyses of the eighteenth century are

---

[204] Kristeller (1951) and (1952). Kristeller is often accused—rightly or wrongly—of claiming that this absence of a concept means that there was no 'art' prior to the eighteenth century. This is a textbook case of the kind of concept-dependence I rejected in §2.3. A more charitable version of this reading of Kristeller claims instead that the kinds of appreciative and cultural practices we currently associate with fine art as a whole first developed in the eighteenth century, so that the practices that came before were different in kind from those that followed. The truth of such a claim, however, is independent of the question of concept-dependence.
[205] Kristeller (1951: 498).
[206] Kristeller (1951: 508-9).
[207] Kristeller (1952: 20-1). The references here are to Batteux's *Les Beaux arts réduits à un même principe*, nouvelle edition (Leiden: Elie Luzac, 1753), 12.

available.[208] More importantly, much of the historical "evidence" that Kristeller adduces in support of his claim that, prior to Batteux, there was no concept of "fine art" or any systematic distinction between fine art and craft, is simply false. So, for example, Kristeller accuses Plato (in *Cratylus* 423c) of numbering the imitation of animal noises among the arts.[209] But, as James Young has pointed out,[210] the discussion of animal noises in the *Cratylus* has nothing at all to do with *art*, let alone fine art: Plato is simply arguing that one does not *name* animals by imitating their noises! Young's careful work reveals that Kristeller routinely mistakes discussions of imitation and imitative practices for discussions of artistic practices. Worse still, Young observes that both Aristotle[211] and Plato[212] *did* identify music, dance, poetry, painting, and sculpture as arts and distinguish them from other art-forms, and that Batteux himself explicitly credits Aristotle with this grouping. Kristeller's contention that pre-modern peoples did not share our concept of art is thus grounded in a systematic conflation of pre-modern talk of "the arts" with the fine arts.[213] To this list of errors we might also add that Batteux did not actually isolate the fine arts from considerations of morality and utility.[214] The result is that there is no sound historical reason to believe that our ancestors had a wildly different concept of the fine arts; what they may have had, instead, is a broader notion of the arts in general—certainly one that is broader than our notion of fine art. So much the better for concept-driven ontologies.

Kristeller was not wrong on every count, however. For one thing, the class of the *fine* arts is not an entirely stable one, and it never has been: art-kinds pass in and out of it, depending on

---

[208] Clement Greenberg, for example, thought the artworlds of the seventeenth and eighteenth centuries were dominated by literature (Porter, 2009: 4).
[209] Kristeller (1951: 504).
[210] Young (2015: 7).
[211] Young (2015: 4-5).
[212] Young (2015: 5-6).
[213] Young (2015: 10).
[214] See Porter (2009: 13) and Young (2015: 3).

prevailing cultural attitudes.[215] Contemporary culture, for instance, has added photography and film, while gardening exists only at the periphery of fine art today (in the West, at least). More importantly, Kristeller was right to focus his attention on eighteenth-century Europe because the practice of art-making and the cultural attitudes surrounding art underwent rapid changes at the time. As Peter Kivy has put it, "the eighteenth century witnessed a veritable explosion of separate, self-contained works on the subjects of the fine arts, beauty, sublimity, taste, criticism, the standard of taste, and much more—works unprecedented in Western intellectual history."[216] It is no coincidence that it is in the eighteenth century that aesthetics was elevated to the status of a canonical philosophical discipline, and that aesthetics and the philosophy of art were first distinguished from one another. Kristeller's cardinal sin was to identify the eighteenth century with the first moment in history when a consensus over 'art's' conceptual content was possible;[217] what actually happened in eighteenth century Europe is that philosophers first started doing the philosophy of art in the full knowledge that that was what they were doing.[218] They began to deploy theories of art in order to reflect upon the value of art and artistic practices.

More recent historical evidence and philosophical argument suggests that Kristeller misidentified the conceptual change that dominated the eighteenth century. What was new was not which kinds of things people classified as 'art', but rather *how they thought* about those things. The eighteenth century saw the crystallization of the concept of an art*work* as a special kind of entity.

An example will serve to make this clear. According to Lydia Goehr (1992), the eighteenth century saw a sea change in musical practice. By the close of the century, the

---

[215] Kristeller contends that it stabilizes after the eighteenth century, which is not quite right, but he may be right that it acquired *more* stability after that point.

[216] Kivy (2012: 70).

[217] See Porter (2009: 14), and Kivy (2012: 64).

[218] Kivy (2012: 73).

development of the concepts of accurate notations, composers, perfect compliance, performance-of-a-work, and scores, among others, resulted in a new concept of musical 'works.'[219] Prior to the eighteenth century, composers had very little creative freedom[220] and a composer's identity mattered less than the occasions for which or persons for whom works were commissioned.[221] Music was not typically the focus of attention—it was treated instead as a background accompaniment to other activities,[222] and false starts were common.[223] Passages were frequently reused to suit similar occasions,[224] and there was a great deal of "creative" overlap between composers and pieces.[225] The absence of complete and institutionalized systems of notation also meant that performers were given a great deal of leeway to complete or interpret passages, and that music was not typically made or expected to outlast its performances.[226] It was only towards the end of the eighteenth century that people began thinking of individual compositions as self-sufficient entities, as 'works' to be published and enjoyed on demand (as it were).[227]

The development and adoption of a complete system of notation alone is responsible for many of the changes wrought to the late eighteenth-century conception of music. This is not to say that there were no functioning concepts of composition, performance, or notation prior to the eighteenth century; rather, it is their significance which changed.[228] For Goehr, these changes reflect the establishment of a concept of musical *works* as distinct entities subject to individual attention, appreciation, and repetition. As Goehr puts it, "The claim is that given certain changes

---

[219] Goehr (1992: 103).
[220] Goehr (1992: 176). Note that Goehr's claim is actually that *musicians* only gained creative freedom at the end of the eighteenth century. Given her subsequent remarks about the freedom performers enjoyed and the underdeveloped nature of scores, however, it seems clear that she actually has *composers* in mind.
[221] Goehr (1992: 179).
[222] Goehr (1992: 192).
[223] Goehr (1992: 194).
[224] Goehr (1992: 181).
[225] Goehr (1992: 183).
[226] Goehr (1992: 186).
[227] Goehr (1992: 202).
[228] Goehr (1992: 108).

in the late eighteenth century, persons who thought, spoke about, or produced music were able for the first time to comprehend and treat the activity of producing music as one primarily involving the composition and performance of works."[229] As this work-concept crystallized in the eighteenth century, it began to develop regulative force; that is to say, it started to determine the normative content of subsidiary concepts (e.g. 'audience,' 'composer,' 'performance,' etc.) and practices.[230] The work-concept thus guides the development of what we can call a 'music-world,' helping it to coalesce into a stable cultural institution.

Larry Shiner finds parallels to Goehr's remarks on the musical work-concept in the development of the other arts. In the old system of art, he argues, a "work of art" was just the product of some art, a construct rather than a fixed creation.[231] Tiffany Stern (2000) has likewise observed, of theatrical productions, that their practice prior to the late eighteenth century departs radically from what we have come to expect. Consider Shakespeare's plays, for which there was, contrary to our current practice, no ensemble revision. Actors practised individual rehearsal ("study"), but received no direction on how to behave when not speaking (there were no producers or directors in the modern sense). Nor were actors given much indication of the play's content: there was no complete typescript to which they could refer—they were only given their cues and lines. Finally, the play's opening night would have been the actors' first opportunity to get a sense of the endeavour as a whole, but it was also a trial run of the play's plot and writing, which were subsequently subject to major rewrites.[232] The result is that, contrary to our ordinary intuitions, Elizabethan theatre was not a text-based art-kind, and neither were *any* theatrical

---

[229] Goehr (1992: 113).
[230] Goehr (1992: 103).
[231] Shiner (2001: 123).
[232] These observations come primarily from Stern (2000: 5-8, Ch. 3, and Ch. 6).

productions prior to the crystallization of a work-concept in the nineteenth century.[233] This amounts to a significant conceptual change between the Elizabethan era and our own—but a change in terms of how we think about theatrical "works", not whether we classify them as art.

The lesson here is that because artworld conventions develop incrementally over time and are historically contingent, the differences between an old and a new artistic practice are not always readily apparent. As Michael Baxandall has observed, our historical understanding of particular artworks is "an analytical construct about [the artist's] ends and means *as we infer them*."[234] Consequently, the history of art, insofar as it is composed of the history of individual artists and works, is likewise just an analytic reconstruction. The inferences we draw about works are based on our understanding of the work's historical context and the interests driving artists at the time, and these are in turn based upon our understanding of our own historical situation and interests. The trick is to discover when and where the analogies to our own time break down, and to revise our constructions in light of those facts—all while remaining aware that it is reconstructions all the way down.[235] In other words, there is no fully independent standard to which we can appeal, because our reconstruction of the historical facts is a reconstruction of something that was itself a construct, an idealization of a culture's and a time's kaleidoscope of different practices.

5.4 – The problem of conceptual imperialism

In the last section I argued that 'art's' conceptual content is unstable across time, and that this fact should caution us against top-down approaches to art's ontology which begin by taking

---

[233] Some, such as David Osipovich (2006) and James Hamilton (2007), claim that theatre has *never* been a text-based art-kind.
[234] Baxandall (1985: 109); the emphasis is mine.
[235] For more on the pragmatic value of Baxandall's heuristic principles and his 'inferential' approach to criticism, see D. Davies (2016).

some concept in hand and then squint at its entrails in the hope of divining something about art's nature. For this strategy to have any hope of working, we would have to repeat it for as many of the historically-situated concepts of art as we can discern. In this section, I will argue that even if we were to do so, we would end up with an inadequately representative ontology. The concepts of 'art' and art-kinds that competent users deploy are frequently wrong because we have a tendency to assume that objects and practices with which we are unfamiliar fit into the conceptual categories that are most familiar to us. I call this tendency *the problem of conceptual imperialism*, and it is most obvious in the way we talk about the artistic practices of other cultures.

By way of illustration, consider Bill Holm's seminal *Northwest Coast Indian Art: An Analysis of Form* (1970 [1965]). In this landmark monograph, Holm presents an analysis of the formal characteristics of more than four hundred objects created by the Kwakiutl First Nation of the Pacific Northwest and tries to derive the aesthetic and artistic principles underpinning their creation. Even so, his preface contains this startling admission:

> Ideally, a study of this sort should lean heavily on information from Indian artists trained in the tradition that fostered the art. Unfortunately, I was unable to locate a qualified informant from the area covered, i.e., the coastal region from Bella Coola to Yakutat Bay. That there may be some still living is not questioned, but contemporary work seen from the area reveals a lack of understanding by Indian craftsmen of the principles that are the subject of this study.[236]

One possible explanation for Holm's inability to find a suitably informed practitioner of Kwakiutl art is just the one he gives: the pure Kwakiutl artistic tradition is endangered, its practitioners either infected by other artistic traditions or insufficiently acculturated to prove reliable guides to its conventions. The more likely explanation, however, is that Holm approached his study with inaccurate or inflexible assumptions about what he should find, and where to find it.

---

[236] Holm (1970 [1965]: vii).

The point here is not to fault Holm's intentions, which I think were beyond reproach; nor is it to take issue with the aesthetic patterns he discerns in Kwakiutl works, which are there to be seen (or not). The point, rather, is to observe that even this peerless scholar of First Nations artistic traditions fell into the methodological trap of beginning his investigation by assuming that his own categories offered an infallible guide to his investigation. In fact, part of the problem here seems to be his insistence on the cultural "purity" of the practices he aims to study, rather than pinpointing the conventions governing current Kwakiutl artistic practices and situating them relative to those which appear to have governed their past practices. The result ultimately denies the Kwakiutl's ability to develop their practices syncretistically, and to master new symbols and techniques.[237] The fact that Holm's search for "informed" practitioners came up short should come as no surprise, since he seems to have been looking for his subject in No True Scotland.[238] Larry Shiner (1994) tells a similar story, according to which an Alaskan State Arts Council representative finds himself in the position of "constantly explain[ing] to Alaskan [First Nations] that ivory carving and beadwork can be supported 'as art,' but kayak or harpoon making cannot."[239] Who is the State Arts Council to deny the close relationship between carving and kayak-making that Alaskan First Nations report over and over to their representatives? If the Alaskans' reports about this relationship are to be believed, then the State Arts Council should re-evaluate either their ideas about art's extension, or their designation of carving and beadwork as 'art'.

---

[237] Keil (1979: 250) is quoted making this point with respect to African artistic traditions in Kasfir (1992: 46). The effect is, ultimately, a silencing of other cultures and a refusal to acknowledge them as participants in the artworld.
[238] 'No True Scotsman' is the name of an informal fallacy in which one attempts an *ad hoc* rescue of a generalization which has already been refuted. So, e.g., when the generalization "All Scotsmen can toss a caber" is met with the counterexample "My brother Duncan can't toss a caber," the fallacious response would be a rejoinder to the effect that "Fine, but no *true* Scotsman can't toss a caber."
[239] Shiner (1994: 225-6).

A similar tendency characterizes the way in which we present works from other cultures, including works from the history of our own culture. As Susan Feagin has observed, "Altarpieces don't transform spaces the same way when they are hung on walls in museums. Neither do paintings originally produced for chapels in churches and cathedrals when they are also hung in museums."[240] Once hung in the museum, they are "stripped of their power to transform the spaces in which they reside into spaces where a viewer is to use them as props in something other than visual games of make-believe."[241] Because our engagement with art is so often mediated by the institutions of the museum and the gallery, as well as by the coffee-table book and textbook, our experience of these works is flat and artificial. We learn about them as two-dimensional caricatures deprived of their cultural and historical significance, or as curios broken off from the larger works that give them their cultural and symbolic power. The obvious example here is the case of the Elgin Marbles, which were broken off from the Parthenon and put on display as separate works of sculpture in the British Museum. Feagin's point is that this kind of violation of a work's context and integrity happens far more routinely than we might think; museums and textbooks are essentially predicated on severing works from their contexts in the interest of accessibility. What we lose is access to the conventions and practices which produced these works, and the result is a flattened ontology with little to show for itself other than an individual work's perceptual properties.

But we can also err too far on the side of caution and exaggerate the differences between cultural practices. This is the moral of Denis Dutton's critique of Lynn M. Hart's (1995) infamous analysis of the *jyonthi* paintings of Uttar Pradesh. Hart argued that the "aesthetic principles" of *jyonthi* painting are different in kind from those of the Western world because the

---

[240] Feagin (1995: 265).
[241] Feagin (1995: 266).

images and patterns are based on religion, ritual, and myth, and the paintings are produced primarily for use in religious contexts. As Dutton observes, however, Hart errs in two respects: first, she ignores the fact that all of the same observations are applicable to much of the Western artistic tradition; second, she assumes that because these works are painted their proper Western analogues are paintings. When we actually compare the *practices* at issue, however, it becomes clear that the appropriate comparison classes are the domestic and dowry arts.[242]

Sidney Kasfir (1992) has argued that this kind of misplaced concern actually betrays a simple double standard: we insist on interpreting the works of other cultures primarily in terms of their political, religious, and social functions, all the while conveniently ignoring the fact that our own artistic practices share parallel histories of political significance, religious patronage, and social utility.[243] When we try to identify which of another culture's artifactual practices are artistic, Baxandall argues that that our first task is to "[work] through to a realization of quite how alien [they] and the mind that made [them] are."[244] The lesson we should take away from Hart's and Holm's mistakes is that this first step must be balanced against a consideration of the ways in which the practice in question resembles practices in our own culture, and the aspects of its history which we call 'artistic'.

One last point of friction between Western and non-Western cultural practices deserves to be mentioned in connection with the imperialistic tendencies observed above: the double standard concerning what ultimately counts for a culture's artistic output. Western artists influenced by "foreign" notions can expect to enjoy acclaim, while non-Western (especially First Nations) artists influenced by Western styles can expect to be criticized for "inauthenticity" (as

---

[242] Dutton (2000: 218-20).
[243] See Kasfir (1992) and Shiner (1994: 231). Similarly, S. Davies's (2007) illustrates the interdependence of Balinese artistic practices and Bali's tourist industry.
[244] Baxandall (1985: 109).

we saw with the Kwakiutl above), or for betraying their cultural heritage and bowing to commercial interests.[245] Non-Western artistic heritage is expected to be static, self-contained, and unresponsive to outside influences. At its extreme, Shiner observes that this imperialistic tendency can have the absurd result that "carvings *intended* to be Art in our sense, i.e., made to be appreciated solely for their appearance, are called 'fakes' and are reduced to the status of mere commercial craft."[246] Similarly, Kasfir notes that "when a contemporary [African] carver from another ethnic group (or 'tribal style area') intentionally takes up this same style, the resulting object is said to be a fake because, it is claimed, there is conscious intent to deceive."[247] White (and male) artists like Gauguin and Picasso are free to explore whatever media, subjects, and techniques catch their fancy but Aleut, Baoulé, and Haida (and female) artists must stick to their "traditional" iconography and kinds—namely, figural and totemistic carvings, masks, and pictures of animals, as though their artistic practices were static and timeless.

The upshot here is just that the concepts of even competent users of a term like 'art' can be seriously flawed thanks to our tendency towards conceptual imperialism. This fact, in turn, stymies any ambitions a descriptivist may harbour towards a universally-applicable theory of art. Yet the descriptivist might still maintain that she is actually describing a concept with a limited range of application—*our* concept, as she might characterize it. In fact, this is exactly how Thomasson conceives of her project. To put the point in her terms, the descriptivist can respond that what we are interested in are the application and co-application conditions of 'art' explicit or implicit in the categorial intentions of the competent users of the term who ground its reference in English (or in our culture). Yet even so, the project stumbles on the facts that "our" concepts are routinely applied far beyond "our" practices, and that "our" concept of art has spread to

---

[245] Kasfir (1992, esp. 41-3, 45) and Shiner (1994: 230).
[246] Shiner (1994: 226-7).
[247] Kasfir (1992: 45).

cultures all over the globe. It no longer makes much sense to speak of "our" artworld as though the pronoun could do more than limit the range of its temporal application. The result is that otherwise competent users of 'art' and art-kind terms frequently apply them improperly to kinds and practices whose ontology they muddy rather than clarify. Any viable descriptive proposal must be careful to take into account the cultural and historical variability of our artistic practices.

5.5 – The problem of conceptual inclination

The third problem facing concept-driven ontologies of art is what I call *the problem of conceptual inclination*. The idea is just that our concepts of art and art-kinds reflect content that is *conventionally salient*, but not necessarily ontologically complete or even correct. If our use of 'art' and art-kind terms does reflect a nascent ontology, then that ontology is heavily influenced by the weight of artworld precedents, and this in turn colours the space of ontological possibilities we are prepared to entertain.

To see how this might be the case, we can begin by looking at women's role in art history. Since the example set by Linda Nochlin's groundbreaking essay *Why have there been no great women artists?* (1988 [1971]), feminist scholarship has argued that art history and art theory are fundamentally flawed. This is because they purport to be neutral and universal, but are in fact based on biased and narrowly applicable criteria.[248] The result is a narrow conception of artistic value—embodied in the art-historical canon—which in turn circumscribes the folk conception of art's ontology.

From the outset, women's access to the artworld was constrained by the traditional system of transmitting professions from father to son.[249] In order for a woman to receive even

---

[248] See, e.g., Brand (2000: 184).
[249] Nochlin (1988 [1971]: 156).

basic artistic training, her father would have had to take an exceptional interest in his daughter's education. Even so, marriage might easily disrupt a fortunate woman's career path; this is because until recently women were expected to abandon the amateurish pursuits of youth to concentrate on child-bearing, -rearing, and household maintenance, or to help their husbands in their business endeavours. While this social attitude is especially well-documented for the nineteenth century, when modest proficiency in several art-forms was considered a sign of a well-educated woman, it is one based on long historical precedent.[250] Aspiring female artists effectively faced a choice between career and love.[251]

Any woman who successfully negotiated these obstacles faced additional stumbling blocks. Her work might, for instance, be widely attributed to her father, mentor, or other male contemporaries, thereby erasing her from art history (as was the case with, e.g., Artemisia Gentileschi and Marie-Denise Villers). Similarly, female artists were often restricted in the subject matter it was culturally acceptable for them to depict (viz. still lifes and scenes of animals); unfortunately, these subjects were considered the purview of amateurs, not of masters.

The case of Rosa Bonheur is the exception that proves the rule. Like virtually all other female artists before her, Bonheur was the child of a family of artists. Although her father initially attempted to have her apprenticed as a seamstress, he eventually consented to train her as an artist. And while she attained renown, it is worth observing that Bonheur's elevated status owes as much to changing standards of evaluation as it does to her skill. As Nochlin puts it:

> We might say that Bonheur picked a fortunate time to become an artist if she was, at the same time, to have the disadvantage of being a woman: she came into her own in the middle of the nineteenth century, a time in which the struggle between traditional history painting as opposed to the less pretentious and more freewheeling genre painting, landscape and still-life was won by the latter group hands down. A major change in the social and institutional support for art itself was well under way: with the

---

[250] Nochlin (1988 [1971]: 164-6).
[251] Loosely speaking, since until the seventeenth century "love" had very little to do with a woman's matrimonial fate.

rise of the bourgeoisie and the fall of the cultivated aristocracy, smaller paintings, generally of everyday subjects, rather than grandiose mythological or religious scenes were much in demand.[252]

Bonheur enjoyed a successful career and much contemporary critical acclaim, but she was lucky to paint her animals and pastoral scenes at a time when these kinds of scenes were increasingly accepted among the subject matter of high art.

Even so, Bonheur's choice of subject matter was not quite as fortuitous as it might seem. History painting was the great genre of the nineteenth century, and it required the depiction of male and female nudes. Female students, however, were barred from attending life drawing classes in the various academies, lest they encounter naked humans. Later, when women *were* allowed to attend such classes, the models had to be partially draped.[253] Prevented from developing their skills in the most prized kinds of depiction, women had to seek their subjects elsewhere, in the lesser fields of genre painting, portraiture, and still life: the very genres at which Bonheur excelled, and which became fashionable at the end of the century. Finally, and most notoriously, in order to study her animal subjects more closely (i.e. in the abattoirs and fairs) Bonheur was forced to ask the prefect of police for permission to wear "masculine clothing" (viz. trousers) in these environments.[254] The point is just this: despite our pretensions to the contrary, quality of workmanship actually has very little to do with artistic "greatness" and an artist's or work's inclusion in the art-historical canon. Instead, it is the partiality of social context which exerts the most influence.

I have focused my attention on the case of women in order to illustrate the kinds of structural factors that have erased them from art's history, resulting in the cultural transmission of a skewed canon of art. This erasure is not necessarily total, since women and other extra-

---

[252] Nochlin (1988 [1971]: 170).
[253] Nochlin (1988 [1971]: 158-9).
[254] Bonheur, cited in Nochlin (1988 [1971]: 174).

canonical agents could still exert some influence over the precedents underpinning the conventional attitudes and practices of which the canon is composed, in the ways I outlined in §4.4. So, for example, women's amateur productions, expressions of judgements of taste, occasional patronage, and their initiation and reinforcement of their children's artworld participation all contributed to the concretization of their culture's artworld concepts. This is so despite the fact that the cultural roles women occupied made it difficult for them to exert such influence more directly (by, e.g., deliberately producing works that challenged the *status quo*). The transmission of counter-canonical practices and ideas is hardly impossible but, proceeding as it does largely by weight of precedent, it *is* hampered by the fact that the canonical narrative exerts its own considerable precedential weight.

Considered by itself, the erasure of some groups' contributions to the artworld does not yet pose an ontological problem. After all, the fact that paintings by Renaissance women were under-appreciated until recently does not mean that they were perceived as non-artworks. To get an ontological problem, we must look at the role that notions like the art-historical canon play in the development of artworld conventions. To that end, let us begin by considering the art-historical canon, which we can stipulatively define as follows:

> *Art-historical canon*
> The collected body of artworks judged by a culture to be of the highest quality in a given art-kind, based on criteria that go beyond mere 'interest,' which helped to define or to develop the art-kind's history.[255]

The canon is necessarily selective; only those works which significantly contributed to an art-kind's development are included, although it is understood that more than just these works are to be classed as art. The canon does not actually exist in any official capacity, but unlike most other

---

[255] Note that the majority of art-historians are agreed that there is no one canon of art; there are many different canons, each fitted to a particular art-kind, culture, thinker, or time. My observations are sufficiently general to apply equally to all of these canons; for simplicity of exposition, therefore, I have chosen to refer to the collection of all these disparate canons as 'the' art-historical canon.

folk concepts it is vouchsafed by art historians, curators, critics, galleries, and other official and semi-official artworld agents, who develop and perpetuate its influence by their economic and social activity. In fact, it so suffuses our culture that virtually any suitably acculturated person can name at least some of its major works and figures, with only minor quibbles at the margins. The result is that these are the practices we explicitly imitate when we set out to make art.

The canon occupies an important role at the heart of the artworld: it is not just a compendium of greatest hits, it also supplies the primary text for the artistic education of a culture (experts and folk alike). We need only crack open any art history textbook to see the extent of the canon's influence: it gives the academic discipline of art history the bulk of its subject matter, kindling student interest in particular artists, styles, traditions, and works, and plays a central role in disseminating the history of art to the broader public (especially through the intercession of critics, galleries, and museums). It furnishes criticism with a reference point for the evaluation of new works, and helps to define the parameters of genre and style. In a word, the art-historical canon selects from among the field of artworks those which are deemed worthy of special attention, and serves the function of turning critical, economic, historical, and even popular interest upon those works, and others like them.[256]

These functions are all instrumental to the development and concretization of the artworld as a self-perpetuating system of conventions. The trouble where ontology is concerned is that the weight of precedent disposes the system to preserve and to reproduce existing conventions. Conventions are by their nature conservative, and the canon gives us a necessarily incomplete picture of artworld practices. The relative absence of women from the canon's ranks shows it to be a reflection of the historically-situated preferences of a privileged few. As Peggy Brand has put it, what we have inherited is "an artworld whose conventions have been

---

[256] Locher (2012: 33-4).

established and perpetuated by a relatively elite group. [...] What has come down to us is an art of exclusion."[257] This means that the judgement of whether a new practice belongs to an art-kind depends, to a large extent, on whether other practices like it have been so judged in the past. The result is that perfectly good candidate art-kinds (e.g. calligraphy, gardening, needlework, textile arts, etc.) are left out of the canon entirely, and this absence affects the space of possibilities we are willing to entertain.[258] Even when new art-kinds are proposed (or discovered, take your pick), their plausibility as art-kinds hinges on their relation to the conventions and traditions embodied by the canon. Consider, for instance, the early efforts to justify photography's inclusion among the arts, which tended to make the case on the grounds that photography shares its goals and methods with the other pictorial arts and that, contrary to appearances, it requires the skilled manipulation of its vehicular medium.[259] These efforts ultimately succeeded in extending the cardinality of art-kinds by one, but the work required to perform this extension also served to legitimate and to reinforce existing criteria for art-kind membership.[260]

To illustrate this point, imagine that Hastings is familiar with all of the great masters of our canon but that his knowledge of artistic practice ends with the likes of Caravaggio, Cézanne, Dürer, van Gogh, el Greco, Leonardo, Michelangelo, Monet, Picasso, Raphael, Rembrandt, Titian, Vermeer, etc. We could hardly fault Hastings for thinking that the pinnacle of human artistic achievement occurred between the fifteenth and twentieth centuries, that art is primarily painted, though maybe also sculpted and printed, or even that art-making is a uniquely European (and male) phenomenon. The more central a place canonical value judgements occupy in one's

---

[257] Brand (2000: 177).

[258] The fact that they are not valued as highly as canonical art-kinds does not mean that they are non-art, even for cultures subscribing to the canonical judgements in question. But, as I shall shortly argue, their absence from the canon directly influences their plausibility as art in the eyes of members of the culture in question.

[259] See, e.g., Emerson (1974 [1889], esp. 328) and Stieglitz (1899).

[260] That is to say, the criteria which were generally *believed* to be required for art-kind membership.

artistic education, the more salient the ontological properties they enshrine become for identifying new artistic practices. In fact, this seems to be borne out by evidence from social psychology. In 1968, Robert Zajonc found that mere exposure to a stimulus is sufficient to enhance a subject's attitude towards that stimulus. A subsequent meta-analysis of 208 studies showed that Zajonc's exposure effect is both robust and reliable.[261] In a follow-up study, James Cutting (2003) likewise found that subjects' preferences for certain artworks is a function of their familiarity with the work in question: aesthetic preference is a function of frequency of appearance, not canonicity, prototypicality, or the subject's expertise (except insofar as these contribute to the frequency of the subject's exposure to the work in question).[262] To say, as Griselda Pollock did, that art history creates its own objects (viz. art and artists) is really just to observe this tendency in action.[263] The canon's influence populates the realm of ontological possibilities.

There is no doubt that our art-historical canons presuppose some general understanding of what makes an artifactual practice an art-kind. My point in this section has just been to argue that we should not look to an emergent entity like the art-historical canon to supply our ontologies with the kinds they must describe, since the canon does not paint a full picture of the variety of artistic practices. Likewise, although there is no doubt that our use of 'art' and art-kind terms presupposes a general understanding of what makes something art, we should not make the mistake of thinking that our uses of these terms are a reliable guide to their ontology. Our uses of these terms just reflect the historically-situated and conventionally-reinforced preferences of a

---

[261] Bornstein (1989: 268). A recent study by Aaron Meskin, Mark Phelan, Margaret Moore, and Matthew Kieran (2013), however, found that mere exposure did not increase liking for *bad* art.
[262] I am indebted here to Lopes (2016), where I was alerted to Zajonc's and Cutting's work.
[263] Pollock (1999: 27). For Pollock, the point of feminist interventions in art history is not just to highlight women's contributions to established art forms, but to expand the sphere of what counts as 'art' in the first place (so that it includes, e.g., the textile arts).

privileged few. What is needed instead is a new account of 'art' that, as Brand describes it, "[looks] seriously at the way the roles of the institution have been meted out to particular subpopulations across the centuries."[264] In other words, concepts are red herrings: we must look for our ontology among our artistic practices as a whole.

5.6 – Descriptivism as practice-driven ontology

None of the arguments I have offered so far against the concept-driven approach to ontology are fatal to the descriptivist project. What the problems of conceptual instability, imperialism, and inclination problematize is the view that the ontology of art is an exercise in bare conceptual analysis—they challenge the legitimacy of the concepts used as data by the concept-driven approach. All this means, however, is that descriptivists must be sensitive to the possibility that their data are misleading. In other words, our "folk" beliefs about art may well be substantially mistaken. This point is neatly illustrated by empirical work in the psychology of text processing, which indicates that readers do not encode quite as much evidence (during the act of reading) as we might think they do. Gail McKoon and Roger Ratcliffe (1992), for instance, found that readers are not overly concerned to encode causal relations: they set aside consideration of the rationale for a character's actions, and do not make inferences about how those actions are performed. They are not even much concerned for the coherence of the text so long as inconsistencies are more than a few sentences apart. The upshot, as Derek Matravers (2014) has argued, is that our theories of literature should not place too much emphasis on readers drawing important conceptual distinctions during the act of reading itself.[265]

---

[264] Brand (2000: 191).
[265] Matravers's target is the fiction/non-fiction distinction, which he argues is a product of convention, not theoretical necessity. Note, however, that the evidence from the psychology of text processing concerns only the occurrent act of reading itself, and not our reflective critical understanding of our literary practices. The evidence

In fact, many of the other descriptivisms on offer today have opted to treat compliance with our intuitions as just one *desideratum* among others. Berys Gaut, for example, cites three methodological constraints on an account of any concept (but of 'art' in particular): intuitive adequacy, normative adequacy, and heuristic utility.[266] An ontological explanation can diverge from the descriptive "facts," provided it does so with good reason. In making the case for a buck-passing theory of art, Dominic McIver Lopes (2014) likewise stipulates his *desiderata* for any good theory of art: viability, informativeness (i.e. the ability to ground empirical research in the arts), and extensional adequacy (especially with respect to the hard cases).[267] Because our folk concepts of 'art' and art-kinds are so disordered, however, Lopes allows for a measure of revisionism: the more disorderly the folk-concept of some art-kind, the more revisionary our theory can be. Folk concepts and intuitions have no privileged status in Lopes's preferred methodology; what does enjoy such a status are the properties implicit in our best empirical understanding of works belonging to the art-kind under investigation.[268]

In this final section, I want to point to an alternative descriptive strategy which I think is especially promising given the kinds of considerations I have raised so far in this monograph. In contrast to the concept-driven approach to ontology, we can characterize this strategy as *practice*-driven. The practice-driven approach to art's ontology takes a different view of what supplies our ontological data: rather than concepts or intuitions, it is our artworld practices which should constrain our ontological investigations. This approach concedes that our concepts and

---

should thus be consistent with (or explained by) our theories of literature, even if it need not characterize them in the same way as it does the occurrent act of reading. Causal relations and global coherence might not matter much for reading, but they do matter a great deal when it comes to making sense of our literary practices!

[266] Gaut (2000: 30-1). Intuitive adequacy names the requirement that the theory agree with our intuitions about actual and counterfactual cases, while normative adequacy simply stipulates that deviations from intuitive adequacy require explanation. Heuristic utility just ties theoretical adequacy to whatever is true of art, or what we think is plausibly true, upon reflection.

[267] Lopes (2014: 15-6).

[268] Lopes (2014: 128).

intuitions may well be substantially wrong, embracing an epistemic humility absent from its concept-driven counterpart.

This is not by any means a novel suggestion, although its appeal has, until recently, been largely sidelined by the concern for extensional adequacy that has dominated definitional projects for the last sixty years. As early as 1977, Timothy Binkley argued that the reference point for our art-ontological investigations should be our artistic practices.[269] Julius Moravcsik (1993) was the first to explicitly suggest that the cross-cultural identification of art should proceed by a detailed comparison of the candidate practice to a virtual catalogue of all the properties we regularly attribute to the things we call 'art', a catalogue which we constantly revise in light of the insights gained through the process of comparison.[270] Noël Carroll has likewise argued that "a comprehensive theory of art must accommodate the facts as [the theorist] finds them revealed in our practices."[271]

The most influential of these proposals, however, comes from David Davies (2004, 2009c, and forthcoming a), who has argued for what he calls the "pragmatic constraint" on art ontology:

> Artworks must be entities that can bear the sorts of properties rightly ascribed to what are termed 'works' in our reflective critical and appreciative practice; that are individuated in the way such 'works' are or would be individuated, and that have the modal properties that are reasonably ascribed to 'works,' in that practice.[272]

In a nutshell, the idea is just that the ontology of art should be circumscribed by a critical assessment of a practice's goals and values in order to determine what its *actual* practices are like, and to what they commit us. So, although our untutored beliefs provide the starting point for

---

[269] Binkley (1977: 271).
[270] Moravcsik (1993: 431).
[271] Carroll (1999: 182). Carroll makes a similar point with respect to film theory in particular when he argues that "Successful films should test the [film] theory, not vice-versa" (2003: 4-5). In other words, theories are tailored to the phenomena they explain. A theory should not invite us to revise our classification of *X*s; new *X*s should prompt us to revise our theory.
[272] D. Davies (2004: 18).

ontological investigations, they are not its end point: we must remain open to the possibility that our investigations will ultimately mandate that we revise either our practices or our pre-existing metaphysical categories.[273]

To see what a practice-driven approach to ontology might look like in action, let us take Balinese *mabarung* as our case study and ask whether it should be classified as art and, if so, whether it belongs to the art-kind we call 'music'. The *mabarung* is a kind of competitive concert offered, among other things, to propitiate the Balinese gods.[274] Stephen Davies describes it as follows:

> Sometimes, Balinese art competitions are relatively informal. *Jegog* is a form of gamelan in which all the instruments are made of bamboo. *Mabarung* between side-by-side *jegog* groups involves the simultaneous playing of different pieces, with each ensemble trying to drown out and outlast the other.[275]

From here, S. Davies picks up on Minagawa Koichi's observations of a *mabarung*:

> Shortly after one of them begins to play, the music becomes highly animated, and suddenly the other group enters into the midst of the musical argument. Both groups seem to attempt to destroy the music of the other by interfering. The result is something quite at odds with our normal concept of 'music.' Rather than music, this is closer to sports.[276]

The strongest point in favour of identifying Balinese *mabarung* with the Western art-kind 'music' is that both are clearly (intentionally ordered) sound-events.[277] On the other hand, the *mabarung* seems to differ from Western music in significant respects. The intended audience, for example, is the Balinese gods, not their fellow humans; indeed, its devotional aspect seems to be the *mabarung*'s primary intended purpose. Just as in a Western sports arena, the human audience

---

[273] D. Davies (forthcoming a).
[274] For more on the devotional aspect of Balinese artistic practices, see S. Davies (2007).
[275] S. Davies (2007: 24).
[276] Koichi (1992), cited in S. Davies (2007: 24).
[277] At least, it seems to be so according to our folk conception of music. We should nevertheless be sensitive to the possibility that it only looks that way for parochial reasons, such as the use of what look like musical instruments. I am bracketing, for the time being, the question of whether our folk conceptions of art's ontology can be substantially wrong. I will return to that problem in Ch. 6.

cheers particularly well-executed passages and jeers mistakes, and tries to support its 'side' in whatever ways it can.[278]

So: how should we go about determining whether *mabarung* is an artistic practice, let alone its art-kind membership? The answer, I think is quite simple: we need to compare *mabarung*'s properties to those of all the practices which we already classify as 'art', and consider the extent to which they overlap. We must then make a decision based on how well it fits into this list and—this is the crucial point—revise our list in light of the new data points offered by our recent investigation and decision. If we decide, for example, that *mabarung*'s devotional significance undermines its art-status, then we must be prepared to excise liturgical music and religious iconography from our starting list, or to offer countervailing reasons for their inclusion. If, on the other hand, we decide to number *mabarung* with the musical arts, then we must be prepared to reconsider the extent to which and the ways in which audiences can be involved in musical performances.

When we properly perform this comparison for *mabarung*, it should be clear that it is not all that far removed from Western musical traditions. Western music has a very long liturgical tradition of its own, after all, and Western audiences regularly cheer, clap, and jeer at concerts. Even the competitive element of *mabarung* is reflected in the European tradition, where musical competitions are relatively commonplace. There is no question that the *mabarung* is different from, say, eighteenth-century European pure music—but then, so is melodic death metal. Carefully considered, the differences between these practices look more like differences of emphasis and inflection—perhaps of *genre*—than of kind. So I think it is safe enough to classify

---

[278] S. Davies (2007: 25). This support can even take the form of trying to make the other 'side' err.

mabarung as an artistic practice, and even as a form of music, so long as our reasoning is made explicit and left open to revision.[279]

In fact, this is more or less what we already do. Just recall Hart's analysis of *jyonthi* painting (§5.3), which proceeded by comparing *jyonthi* painting's "aesthetic principles" to those at work in Western painting. Or consider Holm's analysis of Kwakiutl art (§5.3), which began with a derivation of "aesthetic principles" from a list of the characteristics of more than 400 Kwakiutl artifacts. Hart and Holm were on the right track, insofar as the classification of artifactual practices ought to be based in careful observation of the practices at issue. It is natural enough to assume that the practices of another culture work like our own, especially when they resemble one another closely, but that assumption must be tempered by observation and remain revisable in light of evidence (including testimony from practitioners).[280] Hart and Holm run into trouble not just because they prejudge the issue, but because they do not stop to reflect on their initial judgements in light of the facts they later discover. Hart does this by identifying the wrong comparison class, and Holm by failing to recognize that Kwakiutl culture is not static and timeless, and thus failing to situate his analysis of their past practices within the framework of the conventions governing their current artistic practices.

The strategy may seem somewhat underwhelming, since it does not offer us the comfort of the certainty that we have correctly identified an artistic practice and its corresponding art-kind. Yet recall Ch. 4, where I argued that the artworld is built up out of an arbitrary and contingent network of conventions; it should come as no surprise, then, that the artworld has no

---

[279] This is not to say that differences of emphasis cannot add up to differences of kind: just recall the discussion of guernicas from §4.2. I do not think the case of *mabarung* presents as stark a difference in conventions as that of guernicas, although I would not choose to die on that particular hill. The point is just that classification ultimately requires a *decision* on our part. The power of my proposal is that it allows for just such disagreements, and forces us to make explicit the reasons underpinning them.

[280] On this subject, see Dutton (2000) and (2009), and S. Davies (2016).

neat ontological joints for us to carve. Nor are our aesthetic tastes and interests based on reasoning from metaphysical first principles. As Dutton puts it, they are "rather more like a haphazard concatenation of adaptations, extensions of adaptations, and vestigial attractions and preferences."[281] Different historical accidents result in different preferences and interests, so that different populations are likely to have different artistic conventions—at least until they come into contact with one another and begin the process of cultural transmission.

To be sure, there is a problem of vagueness here but it only serves to illustrate the point that art's basis in conventions precludes metaphysical tidiness and calls instead for a (defeasible) decision on our part. Suppose we were faced with classifying a new practice that had grown out of artworld practices (thereby giving us a good but defeasible reason to class it as art), but which now looked frightfully non-artistic. Imagine, for instance, a *mabarung** several centuries hence which grew out of the *mabarung* as we know it (a practice which we just saw maps fairly neatly onto certain kinds of Western musical practices). The *mabarung** would thus have the weight of precedence behind it, motivating its classification as art. But now suppose that the competitive element gains so much sway in centuries to come that it spawns a deadly practice of sonic warfare reminiscent of the Aztec ritual of flower wars.[282] At that point, it would probably be safe to say that *mabarung**, though rooted in an artistic practice, is no longer an art. But again, that decision would have to be rooted in careful consideration of *mabarung**'s cultural role, and in a comparison to the roles assessors already ascribe to artistic practices in their own cultures.

---

[281] Dutton (2009: 219).
[282] The similarity here is between two specialized forms of waging war that differ in significant respects from their ancestor practices. I do not mean to suggest that Aztec flower wars were sonic in nature or that they grew out of an artistic practice. The evidence, rather, suggests that flower wars developed either out of a religious ritual (perhaps intended to relieve a famine), for combat training, or, more likely, as a sophisticated combination of propaganda, attrition, and proxy warfare (see, e.g., Hicks 1979, Isaac 1983, and Hassig 1988).

5.7 – Conclusion

Readers will recall that I argued, in Ch. 2, that art-making proceeds either deliberately (i.e. with a direct intention to make art) or incidentally (i.e. with an indirect intention to make art). The justification for this disjunction just lay in the fact that so much of the art around us appears to have been made under conditions that would not satisfy direct intention-dependence. The choice, then, was between accepting the possibility of incidental art, and excising the bulk of what we regularly think of as art. The grounds I cited in defense of IID were considerations of descriptive adequacy.

This chapter has, in effect, proven to be a more detailed defense of just that position. By presenting the case against concept-dependence in Ch. 2 as based on descriptive adequacy, I invoked descriptivism as my preferred methodology. And yet a prominent strand of descriptivism maintains that ontology is just an exercise in conceptual analysis, so the hydra grew its head anew. Unless I was willing to concede defeat, I owed readers an explanation of just how an account of art that champions indirect intention-dependence might still lay claim to a kind of descriptivism.

So, I argued that the concept-driven approach to ontology faces three main challenges to its claims to descriptive adequacy: (1) that concepts of art appear to be variable across cultures and times, (2) that our approach to identifying art in other cultures systematically ignores conceptual input from practitioners, and has a tendency to assimilate their categories to our own, and (3) that our ways of thinking about art are heavily influenced by established artworld practices and values (as embodied in the art-historical canon). Our intuitions and concepts thus underdetermine our artworld ontologies. The result is that by opting for a concept-driven ontology, we run the risk of excising from 'art's' extension most of the artifacts from our own

128

culture's past, the artifacts and practices of other cultures, and the kinds of practices, especially by minority populations, which do not enjoy widespread recognition. These considerations, I have argued, mandate an epistemic humility that is simply not available to the concept-driven approach.

We should not be too quick to dismiss descriptivism, however. 'Art' is a social kind, after all, not a natural kind, and that fact seems to make some kind of difference. I have argued that a practice-driven ontology of art offers better prospects for explaining art's status as a social kind rooted in a historically-contingent network of conventions that replicate by weight of precedent. The price is that our ascriptions of art-status and our kind-classifications lose their veneer of objectivity and metaphysical necessity. And yet, if I am right that art is rooted in Millikanian conventions, then this is exactly the right result. The identification and classification of artistic practices depends primarily upon a decision on the part of assessors; the best we can hope to offer is a methodology that accurately describes our practices, and which embraces epistemic humility.

Although I have now presented the case for a practice-driven approach to descriptive ontology, I have not yet argued for why this approach is preferable to the revisionist alternative. It is this task which will occupy me in the next chapter, Ch. 6.

Chapter 6 – Adjudicating the Dispute

6.1 – Introduction

Throughout this monograph, my arguments—for art's indirect intention-dependence, for an attempt-theory of failed-art, for an analysis of the artworld in terms of Millikanian conventions, and for the desirability of a practice-driven approach to ontology—have been motivated in large part by a concern to ensure that our theories of art fit the artworld data. But what guarantees do we have that, in deferring to the artworld data, we are deferring to a body of knowledge that is substantially *correct*?

Readers will have noticed that in the course of arguing for art's indirect intention-dependence I have left behind me a trail of debris which might be used to motivate just this kind of skepticism. The possibility of incidental art introduced in Ch. 2, for example, indicates that we might not even be aware that some of our practices are artistic in nature, while Ch. 3's investigation of failed-art showed that non-art may sometimes be mistaken for art, especially if the reason for the work's failure is not immediately accessible to us. Ch. 4 characterized the development of the artworld as a series of mutually-reinforcing historical accidents, while Ch. 5 argued that our judgements about art and art-kinds are susceptible to several different sources of error. These considerations all seem to motivate what Julian Dodd has called "folk-theoretic modesty" (FM), the principle that our common-sense ontological views are susceptible to radical error. Indeed, the sheer number of potential pitfalls which I have described might lead us to wonder whether revisionism might not in fact be preferable to descriptivism. FM, however, leaves us with a serious problem: if descriptive adequacy is a fool's errand, then how are we to choose between the many different ontologies of art which are on offer? What is more, if I was

right to argue in Ch. 4 that art-making is a fundamentally conventional practice, then how is it possible for us to be substantively wrong about the content we attribute to those conventions in the first place?

This chapter will argue that all is not lost for the practice-driven descriptivist: we can reconcile our inclination towards deference to the artworld data with the epistemic humility embodied by FM once we recognize that artworks are social, not natural, kinds. I shall begin, in §6.2, by giving the case for revisionary ontologies of art, focusing in particular on Julian Dodd's landmark argument for folk-theoretic modesty as a consequence of metaontological realism.§6.3 will explore the reference of natural-kind terms in an effort to provide additional motivation for Dodd's observations by showing that our beliefs about natural kinds have no bearing upon their ontology. In §6.4 I will extend this analysis of the reference of natural-kind terms to social kinds, showing that 'art' and art-kinds, at least, have access to the same kinds of reference-fixing tools as natural kinds. In §6.5 I will return to the issue of metaontological realism, arguing that while MR sets plausible constraints on natural kinds, it is too strict for social kinds. I follow David Davies in arguing that a term's ability to play a particular explanatory role in a theory of 'art' should be cashed out in terms of its referent's capacity to play certain kinds of roles in our practices. The upshot will be that although our artistic practices and beliefs are susceptible to error, we cannot err with respect to specifying the *subject* of our inquiries. Finally, in §6.6 I will argue that we must take care, in specifying that subject, not to pre-judge the issue in favour of our beliefs and practices. FM enjoins us to adopt epistemic humility; we cannot pre-reflectively single out some of our beliefs as constraining our inquiries. We must first consider the set of our beliefs and practices as a whole, and subject them all to philosophical scrutiny. It is only what is left at the end of this process that should be taken to constrain our theories.

6.2 – Revisionism

As I noted in §5.2, revisionary metaphysics aims to improve the structure of our thinking about the world.[283] Rather than hold its explanations to the standards set by our thinking, revisionary metaphysics holds them to external standards set by, e.g., the principles of science, logic, or general metaphysics. The motivation for revisionary metaphysics goes back at least as far as Whitehead, who thought of our conceptual schemes as accidents of history rather than robust intuitions worthy of a privileged status in philosophical investigations.[284] Revisionists thus advocate deference to general metaphysical and scientific principles rather than to folk beliefs about some concept or thing. So far as the ontology of art is concerned, revisionist accounts sacrifice descriptive adequacy for the sake of other theoretical *desiderata*.

Before tackling the issue of ontological revisionism, a few words are in order concerning another kind of revisionism which might seem relevant here. Art-historical revisionism is a practice that aims to reinterpret the historical record about some piece of art history, to challenge and replace the orthodox views surrounding it. In the immediate aftermath of the publication of Nochlin's *Why have there been no great women artists?*, for example, there was a rush to "add women and stir,"[285] so that the art-historical canon might be more representative of women's contributions. Indeed, the increased interest in women's works led to the discovery of several artworks which had been misattributed to male contemporaries of women artists. Legitimate historical revisionism employs appropriate methods to re-examine the historical record so as to ensure that it reflects new discoveries of facts and evidence, and to purge bias from existing interpretations. Illegitimate revisionism, on the other hand, aims to distort the historical record,

---

[283] Once again, see Strawson (2003 [1959]: 9).
[284] Haack (1979: 363-4). Notice that, if Ch. 4 is correct, this is exactly the situation we are in with respect to 'art' and the artworld.
[285] This expression is borrowed from Peggy Zeglin Brand (2000: 189).

often by using otherwise inadmissible historical or rhetorical techniques (e.g. introducing forged documents as genuine, misattributing conclusions to sources, manipulating statistics, deliberately mis-translating texts, etc.).[286] Consider, for example, Donald Kuspit's (2004) proposal that we excise post-Impressionism and its legacy from the canon because it led to the demise of the aesthetic in favour of the theoretic. Because his argument is based in personal preference and would distort the art-historical record, it amounts to a kind of historical negationism[287] or denialism—in other words, it is illegitimately revisionary.

I am quite sympathetic to the kind of revisionist cottage-industry that followed on the heels of Nochlin's essay; it seems necessary to undo some of the damage wrought by the legacies of conceptual imperialism and inclination. My concern here, however, is with revisionism about the *ontology* of art, not its history. And while it seems clear that some kinds of ontological revisionism may be legitimate, and others not, I do not have the resources here to do this distinction justice. Suffice it to say that the division between legitimate revisionism and outright denialism is not as obvious for the ontology of art as it is for historical claims. Doubtless this is because the facts and methods of ontology are neither as firmly settled nor as indisputably established as their historical analogues; they are responsive to arguments rather than to physical evidence. Therein lies both their great strength, and their great weakness: strength, insofar as it is a simple enough task to eliminate illegitimate arguments on the grounds of fallacy or invalidity, and weakness inasmuch as intuitions, partiality, and vogue often decide between arguments.

Let us turn now to the case for a revisionary ontology of art. Perhaps the most prominent motivation for this kind of revisionism comes from Julian Dodd's work on the ontology of

---

[286] For more on historical revisionism, see Michael Shermer and Alex Grobman's *Denying History: Who Says the Holocaust Never Happened and Why Do They Say It?* (2009 [2002]).
[287] The term is Shermer and Grobman's (2009 [2000]).

music.[288] In a nutshell, Dodd argues that attempts to determine the ontological status of musical works ought to defer to the results of metaphysics rather than to folk intuitions about music. When we do so, he thinks, we discover that musical works are eternally existing and uncreated *types* rather than created works.[289] Although he initially restricts his remarks to the case of works of pure music,[290] Dodd eventually[291] extends his arguments for musical revisionism to a defence of revisionary ontologies of art more generally.[292]

Dodd bases his argument on a metaphysical principle he thinks we should all endorse:

> *Meta-ontological realism* (MR)
> The correct answers to first-order ontological questions are in no way determined by what we say or think about these questions.[293]

According to MR, the answers to questions concerning an entity's ontological status and its individuation conditions have nothing whatsoever to do with our beliefs or intuitions about its ontological status and its individuation conditions. Our beliefs and intuitions might well have a great deal to do with the *structure* of those practices, with how their objects are *treated*, but they have nothing at all to do with whether they exist, or what shape that existence takes. To beat a dead horse, the correct answer to what water is has nothing to do with what we say or think about it, but rather with its chemical and physical properties: water is H_2O.

The real work, however, is done by a different principle which follows from MR:

> *Folk-theoretic modesty* (FM)
> Our common-sense art ontological views might be substantially mistaken.[294]

---

[288] Especially Dodd (2007), (2008), and (2013).
[289] For a more detailed exposition and analysis of Dodd's argument, see §6.5.
[290] In Dodd (2007).
[291] In Dodd (2013).
[292] Although perhaps not of revisionary ontologies in general. This is because 'art' is a social kind without a real essence. Its essence is dependent on cultural practices and social attitudes to an extent that is not necessarily generalizable to all other social kinds (e.g. 'weed'—see S. Davies [2003]).
[293] Dodd (2013: 1048).
[294] Ibid.

Dodd takes FM to follow straightforwardly from MR: so long as we concede that an entity *E*'s ontological nature is independent of what we think about *E*s, it follows that *E*s might turn out to be very different from the way we think they are.[295] Once MR is granted, FM has been secured. Folk-theoretic modesty thus presents us with a straightforward background principle: we must bear in mind that common ways of talking and thinking about artworks may be misleading or just plain wrong when it comes to their ontological nature. We saw in Ch. 5 that our art-ontological frameworks call for a dose of epistemic humility; FM just enshrines that principle.

6.3 – Reference and natural kinds

The case for FM is borne out by closer consideration of the meaning of kind-terms—at least as far as natural kinds are concerned (I will deal with social kinds in §6.4). The question we have to ask ourselves is whether our beliefs about natural kinds play any role at all in determining what it is that these entities are—and, if so, just what that role is. One possibility, explored in §5.2, is that it does: these concepts play an ineliminable role in fixing the reference of the kind-terms at issue. We saw, there, that this means that competent users of a kind-term cannot be mistaken about the basic or categorial properties of the kind in question (i.e. their categorial concepts cannot fail to apply to members of the relevant kinds). But, as we saw in §5.3, §5.4, and §5.5, we have good reason to be skeptical of this kind of epistemic privilege.

A more plausible alternative comes from Hilary Putnam (1975, 1990), who suggested that a natural-kind term's reference is fixed by the world itself, not by our theories of the world. For Putnam, natural-kind reference is determined by the causal powers of the entity we pick out when we introduce our natural-kind term. The fact that the referents of a term have their causal powers in common (as evidenced by the sameness of their microstructural properties) will then

---

[295] Dodd (2013: 1049).

determine which empirical generalizations (theories) those terms can enter into (rather than the other way 'round, as descriptivism would have it). Natural kinds are thus those kinds which can play an explanatory role with respect to various phenomena in virtue of their entering into these kinds of empirical generalizations. Consider 'gold', which has many different macro- and microstructural properties: e.g. it is the precious reddish-yellow element that is the most malleable of metals and has atomic weight 196.967, atomic number 79, and specific gravity 19.3 at 20ºC. We can certainly *describe* gold in terms of these properties, but its reference is not fixed by that description. If it were, 'gold' would change its reference every time we refined that description or proposed an entirely new one to take its place. In fact, it would turn out that for most of human history, dragons, thieves, and warlords have amassed hoards of some nameless stuff! Allow me to explain.

Pre-scientific peoples knew nothing at all about gold's atomic structure, let alone its specific gravity. Yet their ignorance of gold's microstructural properties (and of its actual causal powers) did not prevent them from caring deeply about whether their brooches were made of gold or fool's gold; only one of the two is valuable, after all. Yet the fact that they were ignorant of gold's microstructural properties (and of its actual causal powers) does not mean that their uses of the term failed to refer, or that they referred to something else entirely. In fact, pre-scientific peoples are not all that different from most of us today, at least as far as a thing's microstructural properties are concerned. Comparatively few people today, after all, know much about gold's microstructural properties beyond its being an element; its atomic weight, atomic number, and specific gravity are beyond most of our abilities to recite (let alone properly comprehend). Surely this does not mean that most contemporary users of the term 'gold' are incompetent, and fail to refer when they deploy the term.

Certainly not! Language is a cooperative endeavour: individuals do not need to reinvent the wheel for every term in their lexicon. Hastings need not be a pedologist to acquire or make competent use of the term 'muskeg'; he can glean at least some of its conditions of application from pictures and novels set in the Arctic.[296] Putnam's insight was that in using terms like 'gold' or 'muskeg', we rely on a *division of linguistic labour*. That is to say, we do not ourselves need to be in possession of the conditions needed to reliably distinguish between an instance of the kind and other things. All we need is to be in possession of a conventional idea of what members of the kind typically look like, act like, or are like (a "stereotype";[297] e.g. gold is a yellow metal, and muskeg is a mossy northern bog), and we need to stand in the right kind of causal-historical relation to the referent of the term.[298] The important thing is that the linguistic community as a whole is in possession of more refined means of individuating the kind in question. In other words, when it comes time to distinguish gold from fool's gold and our folk concept or stereotype lets us down, we can defer judgement to some group of experts who can reliably distinguish between the two.[299] In this way, no one individual needs to have a full grasp of some term's conditions of application, so long as these are present elsewhere in the linguistic community.

So far as gold is concerned, our competent use of the term today relies on the fact that experts elsewhere in our linguistic communities have a reliable method (based on gold's microstructural properties) of recognizing which things are gold and which are not. But what about pre-scientific peoples, none of whom knew anything at all about gold's atomic structure

---

[296] What he refers to, however, will be *muskeg*, not whatever entities satisfy his conception of muskeg (if there are any).
[297] Putnam (1975: 230). Stereotypes are like the descriptions associated with proper names according to the causal theory of naming: they give us one way of picking out an object, but they do not give us necessary and sufficient conditions for falling into the name's (term's) extension. It is the world around us that supplies its extension.
[298] So that, e.g., if Hastings keeps trying to apply 'muskeg' to rabbits, he will be all out of luck.
[299] Putnam (1975: 227-9).

and other microstructural properties, and whose ideas about its causal powers (e.g. chrysopoeia) were substantially mistaken? It turns out that even they had access to a linguistic division of labour which produced more or less reliable results, thanks to (among other things) what is known as a 'streak test'. Streak tests work on the (largely correct) assumption that metals dragged on a touchstone leave behind a coloured streak; gold leaves a golden streak whose hue matches that of the object being tested. Streak tests are not foolproof (modern chemistry can defeat them, as can plating the object—provided the assessor does not attempt to test a cross-section!), but for the most part they are reliable. The result is that pre-scientific linguistic communities had access to a reliable method for identifying gold and other metals. Some of their beliefs about 'gold's' extension turned out to be false, but not nearly so many as if they had been relying on a simple folk theory to guide their use of the term.

So pre-scientific peoples were not incompetent users of natural-kind terms, and neither are we. But did the *meaning* of 'gold' change once we finally gained access to knowledge of gold's microstructural properties? The answer hinges on just what we mean by 'meaning'. The *sense* (intension) of 'gold' certainly changed, but its *reference* (extension) did not.[300] People still cared about the same *stuff*; what they stopped caring about was the description of that stuff in terms of 'things that streak gold', since they discovered that the description was insufficiently individuative. They shed one *description*, not one stuff, for another. Put another way, the thing pre-scientific dragons cared to hoard was some particular stuff which they individuated by means of the colour of its streak (or perhaps its susceptibility to bite marks); modern-day dragons care to hoard the very same stuff, viz. gold-streaky stuff *with the right microstructure*, which they learned about from ancient dragons who stood in the right kind of causal-historical relationship to the term's referent. The only difference is that modern-day dragons have more sophisticated

---

[300] Putnam (1975, 1990: 60).

methods for individuating gold: they no longer make the mistake of thinking that, say, *schmold* (which streaks gold) is gold. When Sigurðr came for Fáfnir's gold, he had to slay the dragon before the hoard could be his. If a modern-day Fáfnir were to discover that half his hoard was *schmold*, however, he would gladly give it up to Sigurðr's modern-day counterpart who, in turn, would presumably refuse it. Saga tales aside, the moral of this story is just that a natural-kin term's reference is dependent on its actual nature, not an individual's psychological states or a linguistic community's beliefs. Our collective interests have a role to play in determining the kinds of explanatory *frameworks* in which natural-kind terms occur (e.g. scientific vs. folk discourse[301]), but not in determining that term's extension. *That* much is supplied by the world around us. A term's extension is determined by its reference, not its description, and its reference is secured by the way the world is.

6.4 – Reference and social kinds

Trouble starts to brew when we consider social, rather than natural, kinds. A natural kind's hidden structure determines its kind-membership, but social kinds have no hidden structure in the first place. Consider Hastings, who is an Englishman—let us say he belongs to the kind *English citizen*. As a human being, Hastings[302] himself has microstructural properties: he is made of carbon and other elements, acts through the transmission of action potentials across synapses, etc. But none of these microstructural facts about him make an *English citizen*. So while an

---

[301] See, e.g., the brief discussion of 'jade' in §6.6.
[302] Assuming, for the sake of argument, that Hastings is a real person and not just a fictional character. It remains an open question whether or in what sense the fictional character is a human being (see, e.g., Thomasson 1999).

individual instance of a social kind has microstructural properties, the kind itself does not and so the instance's microstructure cannot figure in an explanation of its kind-membership.[303]

Yet even though there is no microstructure for experts to discover and appeal to in their explanations of citizenship, or for them to use to distinguish genuine citizens from frauds and imposters, the fact of the matter is that there *does* exist an objective measure of citizenship. This measure is given by the complex network of conventions that figure in explanations of the existence and behaviour of nation-states. That network of conventions, in turn, determines which individuals living within a certain geographic area owe allegiance to the area's government and are entitled to its protection (along with determining the character of the allegiance and the type of protection). In this case, we can simplify things somewhat by saying that Hastings possesses (or is entitled to possess) an English passport.

Notice that the meaning of 'citizenship' is not determined by our folk theory of citizenship, even if citizenship is a social kind; its meaning is codified in a series of legal conditions and documents. Determining which individuals are citizens and which are not is thus not merely a matter of what we think, but a matter of determining which individuals satisfy the conditions laid out by English law, such as being born in the right places or under the right conditions, having filed the appropriate paperwork, etc. *Citizenship* may have no microstructural properties, but this fact does not prevent us from dividing our linguistic labours and relying on the relevant "experts" (in this case, bureaucrats).[304]

---

[303] So, too, for 'art'—or at least, for visual art like paintings and sculptures, since these are obviously composed of some matter. The important point is just that the particular arrangement of matter is not what makes these entities *artworks*, or *paintings*, or *sculptures* (see Ch. 4). Note also that although Hastings instantiates (belongs to) a social kind, the man himself is not a social kind.

[304] The conditions for citizenship are somewhat arbitrary, to be sure. This is because citizenship, like art, is a matter of convention—see Ch. 4.

Contrast this with a term like 'chair', which Putnam thinks is not subject to the division of linguistic labour: when a speaker acquires the term 'chair', she also thereby acquires something that contributes to fixing its extension (e.g. a functional role which chairs occupy).[305] This is just because chairs are neither natural kinds, nor sufficiently regimented in our social practices to require any kind of division of labour (much like 'water' or 'gold' for pre-scientific communities). Yet even so, it is not the chair-speaker's individual psychological states which do the work of fixing 'chair's' extension; her use is embedded in, and contributes to, the general sociolinguistic state of her linguistic community, which in turn picks out whatever entities satisfy the particular functional role that chairs perform.[306] In other words, the term's extension depends on the entities that actually serve the relevant functional role.

So where does all this leave artworks? The first thing to notice is that 'art' is not a natural kind-term. Most theories of art class it as an artifactual kind-term like 'chair' or 'pencil' and, thus, as a social kind; likewise, on the account I gave in Ch. 4, 'art' is a social kind. So what determines membership in the kind *art*? The answer hinges on whether 'art' exhibits a linguistic division of labour, and this is where things get tricky, because the evidence is mixed. On the one hand, no art-experts are capable of telling us definitively whether a particular entity is a work of art or not.[307] We certainly have a great many art experts—artists, art historians, critics, curators, philosophers of art, etc.—and we do often ask them to supply reasons for a work's putative art-status.[308] But these experts do not have access to a privileged method of recognizing art, since art has no properties that could ground these kinds of empirical generalizations. There is no art-

---

[305] Putnam (1975: 227-8).
[306] Putnam (1975: 228).
[307] Certain theories of art do confer such a power upon art-experts. Dickie's version of the institutional theory, for instance, sees their verdicts as constitutive of art, and riven through with illocutionary force. Or consider Donald Judd's infamous dictum, "if [an artist] says it's art, it's art" (Judd, quoted in de Duve 1990: 272).
[308] Consider, e.g., the various expert witnesses called to testify before the U.S. Congress on behalf of the National Endowment for the Arts after the controversy over Robert Mapplethorpe's photographs.

ontological expertise akin to the expertise an entomologist wields with respect to insects, or that an astronomer has with respect to the formation of gas giants. Indeed, many of an art-expert's judgements may well be mistaken, for the reasons I outlined in Ch. 5. Nor is it clear that we regularly defer to the judgements of art-experts; just consider the perennial controversies over works of public art such as Richard Serra's *Tilted Arc* (1981). In fact, *Tilted Arc*'s is a case where expert judgements were defeated by public outcry, and the work was destroyed.

On the other hand, we do sometimes defer to expert judgements when it comes to membership in relatively well-defined art-kinds, such as fugues, film noir, odes, neo-classical ballets, sonnets, etc. Indeed, if Walton (1970) and Lopes (2014) are correct in thinking that categorization according to an art-kind or genre is a necessary first step in the identification of artworks, then it seems that this kind of division of labour plays a very important role in our linguistic community.[309]

My own view is that, all things considered, 'art' is rather more like 'citizen' than 'chair'.[310] While it is true that we do not *defer* to the judgements of art-experts for determinations of art-status in general, we *do* consult them regularly when a work's art-status is in doubt, and we demand that putative experts supply us with reasons for their judgements. And they do so; not in terms of an objective measure of art-status, but by highlighting the kinds of conventions that govern—or have governed—our artistic practices, and showing how the work under consideration fits into them. Determining which entities are artworks is not merely a matter of canvassing public opinion, it is a matter of comparing that opinion to the kinds of conventions that govern our artistic practices, and which figure in our best theories of those practices.

---

[309] I am indebted to Eric Murphy for suggesting this line of thought.
[310] Although as I explained above, even 'chair's' meaning is inextricably tied to its use by a linguistic community rather than an individual's intentions.

The artworld supplies us with an alternative to folk theory in the form of a network of conventions that replicate by weight of precedent. There *is* a "method of recognizing" present in our community, considered collectively; it just depends on an assembly-line model of the division of linguistic labour rather than the craftsman model. Through their participation in the artworld, and through their applications of the term 'art', speakers help to reinforce existing conventions by increasing the salience of various precedents. Unlike craftsmen, speakers are not typically in possession of a complete picture of how we go about distinguishing art from non-art; they are alienated from the fruits of their labour. That is to say, they may not be fully aware of the ways in which they participate in the artworld and reinforce its conventions.[311]

6.5 – The possibility of error

We are now equipped to return to the issues of revisionism and folk-theoretic modesty that were raised in §6.2. As we saw in §6.2, folk-theoretic modesty (FM) follows straightforwardly from metaontological realism (MR),[312] the view that the correct answers to first-order ontological questions are mind-independent. Accordingly, a great deal in the revisionist argument hinges on our conceding MR. The considerations I adduced in §6.3 and §6.4 all help to motivate FM for 'art'; the question, however, is whether 'art' is also constrained by MR, and it is here that my path diverges from Dodd's.

While it is intuitively obvious that MR applies to natural kinds, it is not clear that MR is also true of social kinds like 'art'. Ian Hacking (1995, 1999) and Sally Haslanger (1995, 2012), for example, have argued that the very act of classifying something in a social context can change the object of classification, or even establish and reinforce an entirely new classificatory

---

[311] See, e.g., Ch. 4 (esp. §4.3 and §4.4) and Ch. 5 (esp. §5.3, §5.4, and §5.5).
[312] Dodd (2013: 1048).

scheme. Hacking calls this phenomenon the *looping effect*.[313] Just consider the category 'queer',

whose persistent pejorative use led Anglophone LGBTQIA communities to reclaim the term in

the late 1980s. Today, it has largely lost its pejorative connotation and is instead predominantly

used to refer to that community in a neutral manner by both its members and the general public.

The possibility of this kind of interaction with a classificatory scheme indicates that the

concomitant kind concepts can also change their extensions. Haslanger identifies socially

constructed kinds as those for which the conditions of kind membership include social properties

and relations; in other words, they are kinds whose nature and extension depends on just such a

feedback loop.[314] Natural kinds, by contrast, are 'indifferent': they resist interaction.[315]

If Hacking and Haslanger are right, then the looping effect lends credence to the

hypothesis that at least some kinds are mind-dependent (at least as far as our collective thoughts

are considered). This, in turn, supports the conclusion that MR does not apply to social kinds:

what we collectively do (where our agency reflects our thinking) can and does influence the

answers to *some* first-order ontological questions—namely, those concerning social kinds. What

is more, this suggests an avenue of response to Dodd's revisionism which marries descriptivism

to FM by way of rejecting MR: Dodd's remarks on the objects of our deference are appropriate

to natural kinds but they do not apply to social kinds.[316]

But we are getting ahead of ourselves: what *are* these mysterious ways in which thoughts

can determine ontology? Ian Hacking (1995, 1999) and Muhammad Khalidi (2010) have argued

that the mind-dependence of social kinds follows from the fact that their *existence* depends on

---

[313] Hacking (1995: 369).
[314] Haslanger (1995: 97-9).
[315] Hacking (1999: 104-9).
[316] This possibility was first suggested by S. Davies in his (2003); Dodd comments on it in his (2013: 1064-6). Note that this is not quite the solution for which I argue here, according to which our theories may well be revisionary of our standard linguistic use of a term. On my account, MR is wrong because it is *too strict*, not because metaphysics imposes no constraints on social ontology.

human minds: without human intervention, they would not exist in the first place. From this fact, they infer that some of these kinds' *properties* are also mind-dependent. Haslanger, on the other hand, distinguishes between social kinds that are strongly and weakly "pragmatically constructed": only the existence of the former is substantively mind-dependent.[317]

While it is transparently true that the kinds 'food,' 'queer,' and 'weed' all depend on human minds for their existence, Dodd rightly objects that we are not yet operating at the first order of ontology.[318] Although these *kinds* all depend on human minds for their existence and grouping, the underlying *entities* do not. The observation that some social kinds owe their existence to human minds is nothing more than the observation that we determine which properties count towards kind-membership relative to our own interests, or that some things only come into existence as a result of human thought and action. Our proper focus, instead, is on the *objects* of those kind-terms, namely grains (etc.), people, the pathological condition, and the pesky plants themselves. All of *these* objects clearly exist and have their essential properties independently of any minds or thoughts about them.

There are two sets of countervailing considerations here. The first is that the boundary between natural and social kinds is not especially well marked. We already saw in §6.4 that it cannot be drawn simply in terms of what does and does not call for a linguistic division of labour, since some social kinds *do* call for such a division.[319] The second is that this frontier is also populated by liminal kinds that owe their existence to human conceptions and actions but

---

[317] The difference is just this: a kind is weakly pragmatically socially constructed if social factors only partly determine our use of it, and strongly pragmatically socially-constructed if they wholly determine our use of it (Haslanger 1995: 100). Note, however, that Haslanger later qualifies her position: the natural/social contrast "is better viewed as a spectrum from the non-social to the social within the natural" (2012: 213). Social kinds are ineluctably embedded in the natural world; the only properly "mind-dependent" things are thoughts and their ilk (see Haslanger 2012, esp. Ch. 6).

[318] Dodd (2013: 1059).

[319] To this fact, we can add Putnam's later skepticism (1990) that all natural kinds are amenable to definition in terms of their microstructural properties (where those properties reflect the causal powers that entitle them to figure in powerful empirical generalizations).

which, once instantiated, resist further interaction with the (bare) world of thought. Haslanger, for example, has suggested that 'food' is one such kind, since "opinions about what is appropriate for humans to eat and so about what counts as 'food' have had a huge causal impact on the size, distribution, and behavior of animal populations."[320] To this we can add plant populations, too. Stephen Davies[321] has offered the concept 'weed' which, although it fails to capture a natural division, does depend on the mind-independent category of 'plant'. Finally, both S. Davies and Khalidi have observed that even in chemistry, there may exist some reactionary, short-lived elements or compounds that can only exist in the laboratory, as a result of human conception and intervention.[322]

These observations go some way towards reinforcing the view (suggested in §6.4) that natural and social kinds alike exist along a continuum, so that many putatively natural kinds have significant nominal properties not captured by definitions in terms of their microstructural properties, just as some putatively social kinds may also have some kind of real essence which can be captured by a definition. The important thing to notice is just that many social kinds, 'art' and art-kinds among them, seem to make essential reference to the social factors that constitute them. If these observations are correct, then we should reject MR—not because deference to metaphysics is inappropriate for social kinds, but because MR stipulates conditions that are too strict: the way we think about *some* questions really does have *some* effect upon the answers to *some* first-order ontological questions.

To see what this influence amounts to, let us turn back to the case of 'art' and art-kinds. I have already argued (in Ch. 5) that the things we say and think about artistic practices are susceptible to massive error. But I have also argued (in Ch. 4 and Ch. 5) that our collective and

---

[320] Haslanger (1995: 104).
[321] S. Davies (2003: 6).
[322] S. Davies (2003: 5) and Khalidi (2010: 353-4).

reflective thinking about artistic practices informs and sustains the system of conventions that governs their art-status. The upshot, I think, is that we can be wrong about our theories, but not about the *subject* of our inquiries. A theory of 'gold' which does not allow us to discriminate between *Au* and *FeS₂* is not a very good theory of gold, but it *is* a theory *of gold*: the object of our interest is the kind of entity that plays a particular explanatory role, and we treat this role as rigid across possible worlds.[323] My contention is simply that parallel considerations apply to the cases of 'art' and art-kinds. Consider music: if a theory of music is so radically revisionary that it has the result that musical works are not actually the kinds of things that can play the kinds of cultural roles conventionally attributed to them, then it is not clear that what we have is a theory of *music* in the first place, rather than a theory of *schmusic*.

Compare Hilary Putnam's (1975) remarks on a similar hypothetical scenario in which we discover that the "pencils" on a twinned Earth are actually organisms:

> When we discovered this, we would not say: 'some pencils are organisms'. We would be far more likely to say: 'the things on Twin Earth that pass for pencils aren't really pencils. They're really a species of organism.'[324]

The reason for this is just that our use of 'pencil' is rigid: we have in mind not some particular description of pencils, but any and all things that share their nature (whatever it may be) with the things that, in our world, play the pencil-role. It would be more than passing strange to think that we could collectively be wrong about the cultural role that pencils occupy—speaking, at least, of our considered judgements rather than our bare intuitions. The cultural role which pencils occupy is just a matter of the uses to which we put them, after all. A theory of *pencils*—that is, a theory for which 'pencil' designates rigidly—whose result was that pencils do not or cannot occupy this role would be absurd, since the result undermines the theory's ability to refer in the first place.

---

[323] It is worth noting that I think the prospects for giving a definition of 'art' in terms of a Ramsified description of its socio-cultural role(s) are bright, although it lies well beyond the scope of the present monograph to make such a case here. D. Davies paves the way for this kind of analysis of 'art' in his (2009) and (2016), however.
[324] Putnam (1975: 243).

We saw, in §5.6, that this is what D. Davies has in mind with his "pragmatic constraint" on art's ontology. D. Davies, like Putnam before him, rejects the view that our untutored intuitions give us access to the metaphysics of the world around us; instead, they provide an entirely defeasible *starting point* for those investigations. This is not because our thoughts determine the ontological properties of the entities in question. It is because our practices do the work of specifying *which kinds* of entities it is whose ontology we are interested in to begin with: those that play *this* kind of role, or *that*, in our practices. When considering the ontology of a social kind like art, D. Davies argues, we must begin with certain "topic-specific constraints that identify entities in terms of the roles they have in our practices."[325] The result is just an account that fixes the reference of social-kind terms in the same way as natural-kind terms.

6.6 – Dodd's pudding

It should come as no surprise that Dodd has argued, on the contrary, that an artifact's ontological nature supervenes on neither its function nor the network of practices in which it is produced and embedded. He tells us, for instance, that "Pencils [...] could still be used for writing whether or not they turned out to be enduring entities, spacetime worms, or instantaneous temporal stages."[326] On the endurantist model that supplies the default 'folk-theoretic' view, pencils are wholly present at every moment of their existence. On the perdurantist model, however, they are not: instead, they have a temporal part at every instant in which they exist, and all of these parts can be strung together into a single spacetime worm. Talk of pencils would thus be loose talk for either individual temporal parts of pencils, or the mereological sum that is a wormy pencil.

---

[325] D. Davies (forthcoming a).
[326] Dodd (2013: 1059).

Dodd's point here is just that the functional or cultural roles of pencils can be satisfied by many different ontological pictures; there is no close tie between the two.

That much is true enough, but Dodd misrepresents the character of his example. The fact that pencils can perform pencil-functions or play pencil-roles on either an endurantist or a perdurantist model might just as easily count as evidence in favour of the view that these ontologies are not substantially different, in Dodd's sense of 'substantial'. Both enduring-pencils and perduring-pencils are the same kinds of things—writing implements, artifacts, etc.—and each model leaves most of our ordinary ways of thinking and talking about pencils untouched. It is not as though pencils turned out to be *non-artifactual*, as in Putnam's example (where they are organisms). An example should serve to make this clear.

In science, a term's ability to play a particular explanatory role is cashed out in terms of the referent's possession of certain causal powers in virtue of which the term can enter into powerful empirical generalizations about natural phenomena.[327] So, in order to determine 'gold's' reference it is not enough to know that *Au* stands in the right kind of causal-historical relationship to our introduction and use of the term 'gold'; we also need to know that *Au* is capable of playing the right kind of explanatory role in our theoretical framework. In this case, the framework at issue is a scientific one. Compare this to the case of 'jade', where both jadeite ($NaAlSi_2O_6$) and nephrite ($Ca_2(Mg, Fe)_5Si_8(OH)_2$) stand in the right kinds of causal-historical relationship to our introduction and use of the term 'jade'. Yet these are distinct mineral species, each with a different microstructure and, thus, different causal powers. As a *scientific* term 'jade' therefore fails to refer, since we cannot secure the uniqueness of its reference; it cannot play the right kind of explanatory role to figure in our empirical generalizations.[328] This is not to say that

---

[327] See Putnam (1975, 1990) and D. Davies (fothcoming a).
[328] Note, however, that considered *individually* $NaAlSi_2O_6$ and $Ca_2(Mg, Fe)_5Si_8(OH)_2$ *can* do so.

'jade' cannot be a kind-term at all, or that there is no jade. In its disjunctive form, 'jade' can serve perfectly well for ordinary, though not for scientific, purposes: it may well be a *phenomenal* kind, but it is not a *natural* kind.[329]

With respect to social kinds like 'art', the argument is that a term's ability to play a particular explanatory role should be cashed out in terms of its capacity to play certain kinds of roles in our practices. So, for example, if a theory of 'pencils' has the result that pencils are incapable of serving as writing implements (perhaps because they are organisms, as Putnam suggested, or because they are classes) then it is a bad theory of pencils. In fact, it is not clear that it is a theory of *pencils* at all, as opposed to some other phenomenon that resembles pencils. Similarly, a theory of 'gold' that picked out all and only FeS_2 would be a very bad theory of gold; we would be much better off treating it as a theory of *fool's gold* instead. The point is not that pencils' cultural role determines their ontological nature; as D. Davies puts it, the point, rather, is that our inquiries into specific entities are governed by certain topic-specific constraints.[330] For our scientific inquiries into the nature of natural kinds, those topic-specific constraints are determined by the referent's causal powers; but because a social kind's explanatory value is not exhausted by its microstructural properties, the constraints must instead come from the role the referent plays in our practices.[331]

The requisite proof, here, can be found in Dodd's own pudding. Dodd's ontology of music sets out to answer two questions about instrumental works of pure music: the *categorial* question (to which ontological category do these works belong?) and the *individuation* question (what are the identity conditions of musical works?).[332] To answer these kinds of questions, he

---

[329]At least, so long as natural kinds are the kinds that are supposed to serve an explanatory role *in science*.
[330] D. Davies (forthcoming a).
[331] D. Davies (forthcoming a).
[332] Dodd (2007: 1).

thinks, we must look to metaphysics. The problem, however, is that 'works of instrumental pure music' are not the kinds of entities metaphysics usually concerns itself with: the category is too broad. So we must first determine what *kind* of entity a work of instrumental pure music is, and this means distilling such works to their essential properties, namely audibility and repeatability.[333] Dodd thus begins his inquiry by looking for the metaphysical kinds that support audibility and repeatability:

> Plausibly, musical works are in themselves both repeatable and audible: more precisely, such works can have multiple occurrences (e.g., performances), and we can listen to a work by listening to an occurrence of it. [...] What kind of entity must a work of music be, given that it can have multiple occurrences? And how, given that an occurrence of a work is distinct from the work itself, is it possible to listen to the work in listening to an occurrence of it? An ontological proposal for works of music should try to answer these questions.[334]

Dodd identifies types as the appropriate metaphysical kind; all that remains is to choose the most appropriate kind of type from the metaphysical menu (viz. norm-types[335]).

My concern here is not to dispute Dodd's musical ontology. Instead, I wish only to observe that Dodd's methodology stands at odds with his commitment to folk-theoretic modesty. Recall that FM stipulates that our common-sense art-ontological views might be substantially mistaken.[336] Because they are essential properties of musical works, audibility and repeatability hold the key to answering first-order ontological questions about musical works. But why should we believe that these properties are more essential to music than, say, concreteness,[337] creatability,[338] modal or temporal flexibility,[339] or unrepeatability[340]? Dodd's answer, I am

---

[333] Dodd (2007: 3-5) and (2013: 1053). 'Audibility' is meant to capture the fact that we experience musical works by listening to them; repeatability, the fact that musical works can be instantiated at many different times and in many different ways, such as through performances.
[334] Dodd (2013: 1053).
[335] Dodd (2007: 3).
[336] Dodd (2013: 1048).
[337] e.g. Mag Uidhir (2013: Ch. 5).
[338] e.g. Levinson (1980) and (2011).
[339] e.g. Rohrbaugh (2003).

afraid, is that both audibility and repeatability are common-sense views about music that any adequate theory must explain.[341] But, as such, they are susceptible to radical error. The problem is that the selection of audibility and repeatability as the fundamental properties of musical works proceeds by an appeal to folk intuition, while folk intuitions are cited as the grounds for dismissing other plausible candidate properties. Given the fallibility of folk intuitions, compliance with those intuitions is not an appropriate criterion for pinpointing a kind's essential properties—at least not if we are being folk-theoretically modest.

To be clear, I am not arguing that we are wrong to think that musical works are audible and repeatable; in fact, I think that audibility and repeatability form part of the topic-specific constraints that an ontological investigation of musical works must observe. Dodd's mistake is to single them out from all other candidate topic-specific constraints *pre-theoretically*. In this way, he arbitrarily privileges one set of folk intuitions about musical works over others. As D. Davies puts it,

> There is, *pace* Dodd, no *non sequitur* in the idea that our philosophical interests play a constraining role in ontological inquiry, simply a recognition of the need to particularize any ontological inquiry to the things about which we are inquiring, and a further recognition that what particularizes our ontological inquiry into the nature of artwork-kinds is the explanatory roles that such kinds are intended to serve. The interests, in other words, determine *what it is whose ontological status is at issue*, not *what that ontological status is.*[342]

Appropriate topic-specific constraints must not presuppose the correctness of either our beliefs about artworks or our artistic practices. This means that the first step for an ontological investigation into a social kind like 'art' is to consider the *totality* of our beliefs about art and our artistic practices, and to subject these to philosophical scrutiny.[343] Some of these beliefs and

---

[340] All I mean here is that individual compositions might be tailored to specific occasions, and not meant to be repeated for different occasions and audiences. See, e.g., Goehr's (1992) remarks on attitudes towards musical 'works' prior to the 18th century (esp. Ch.7).

[341] Dodd (2007: 3).

[342] D. Davies (forthcoming a).

[343] Including, among other things, considerations of which metaphysical constraints are applicable.

practices will survive unrevised, while others will not. It is only once we have developed a reflective catalogue of these beliefs and practices that we will be in a position to begin asking and answering first-order ontological questions. We have no privileged epistemic access to the ontology of social kinds; our only privilege lies in our ability to determine the proper subject of our inquiries. The result is that, in the end, even Dodd must rely on our collective (and reflective) cultural practices to fix the topic-specific constraints governing ontological investigations of musical works.

6.7 – Conclusion

Other theories of the reference of artifactual terms are available, of course. A descriptivist like Thomasson might, for instance, maintain that the reference of artifact terms is fixed by a description of their functional roles implicit in speakers' categorial intentions—a view which at least sounds close to the one I endorse. I do not think that the prospects for such a view are very promising, however. First, because it is not clear just what art's function *is*—indeed, different theories of art have posited very different functions (including a characteristic *lack* of function!). Moreover, applied to art-kind terms this strategy would seem to throw us back into the clutches of the doctrine of medium specificity, since different art-kinds would have to have different functions (otherwise their referents would be the same). Second, I do not think that this strategy latches on to quite the right phenomenon: art's *functional* role(s), if it has any, is (are) just one part of what I am calling its *cultural* role, which explains the ways in which it enters into cultural explanations. Finally, it would still have difficulty with the kinds of cases we saw in Ch. 5, which showed that our artistic practices have changed a great deal over time, and can look very different across cultures. A descriptive theory would have to class these under distinct terms, or

find some means of denying that they showcase any significant changes in, e.g., the function of art or art-kinds.[344]

These same considerations could be seen as evidence of reference shift thanks to speakers' historically-compounded errors, and thus as motivation for a hybrid theory like that suggested by Gareth Evans (1973 and 1982). Just consider the theory of art for art's sake. In conversation with an English friend who was studying Kant's aesthetics under Schelling's tutelage, Benjamin Constant's imagination was captured by (his misunderstanding of) Kant's notion of disinterested attention, of which he wrote in his journal: "l'art pour l'art et sans but; tout but dénature l'art. Mais l'art atteint au but qu'il n'a pas."[345] For Kant, 'disinterest' characterizes a distinctive kind of pleasure which accompanies judgements of *beauty*, not a theory of art. Nevertheless, Constant's (or his friend Robinson's) error spread through Europe by way of the influential works of Victor Cousin and Théophile Gautier, as well as John Ruskin's criticism in England, and eventually came to dominate the way we talk about art's function and value, even today.[346]

This misunderstanding might well be taken to ground a shift in the reference of 'art' from, say, objects of craftsmanship intended to perform particular functions to works of fine art created for no particular purpose. Because the shift in question would be just one of many, according to the hybrid theory the reference of 'art' or of an art-kind term would be fixed by the dominant source of the body of descriptive information which speakers collectively associate with the term. I take it that this sort of strategy is generally friendly to the one I endorsed above. I would simply add that, once again, it is speakers' *reflective* judgements about the descriptive

---

[344] An especially tough sell given Feagin's (1995) observations about art's site-specificity, which I introduced in §5.4.

[345] "Art for art's sake and without aim; all aims pervert art. But art achieves the aims it lacks" (my translation; Constant 1952: 58; 11 February, 1804).

[346] John Wilcox traces the historical transmission of this mistake in his (1953).

content that they associate with a term (which has a particular explanatory role) that do the work of fixing its reference.

To return to the question set in §6.1: what guarantees do we have that, in deferring to the artworld data, we are deferring to a body of knowledge that is substantially correct? The answer is that we have precious few such guarantees. What I have tried to suggest in this chapter is that, *pace* Dodd, the appropriate response is not to treat artworld data as defeasible by the merest scrap of metaphysical inconvenience. The appropriate response, rather, is to ensure that our methodology is epistemically humble, so that it does not pre-reflectively privilege one set of intuitions over others. We should beware our intuitions and experts who come bearing canons, looking instead to our collective reflective understanding of our artistic practices to supply the data to be explained by our theories.

The artworld is not the result of rational deductions from a logical system; it is a slapdash amalgam of works and practices which have somehow or other captured our individual interests. As a result, some of the data points it supplies are bound to be contradictory or otherwise incoherent. Just consider the category of 'outsider art', which is supposed to denote artworks that are the result of a spontaneous exercise of creativity and emotion by artists working independently of the artworld.[347] In fact, it turns out that almost none of this description tracks reality[348]—yet that fact has not done much to dampen enthusiasm for the category, which drives exhibitions the world over. Our ontological investigations cannot pre-judge the issue in favour of one or another of the artworld's data points; they must first clear the ground, and determine which of these properly constrain our subject matter. The solution to the artworld's arbitrary and

---

[347] Dubuffet (1973: 105).
[348] See, e.g., Cardinal (1994), Elkins (2006), and D. Davies (2009),

contingent nature is neither widespread mind-dependence nor wholesale revisionism; it is a pinch of epistemic humility.

Chapter 7 – Concluding Remarks

This brings us to the end of my analysis of art's intention-dependence. My central aim in this monograph has been to give an account of our basic commitment to art's intention-dependence by thinking through what it requires of us and of our theories of art.

I began, in Chapter 2, by asking what we mean when we say that art is necessarily intentional. I argued that this commitment is best understood as a claim about the practice of art-making, which is clearly a kind of intentional action. I showed there that action theorists are largely agreed that actions are simply events which are intentional under some description. By applying the same insight to artworks, we get the result that an artwork is an entity which is intentional under some description. The operative question thus became: *what does it mean to be intentional under some description*? I argued that, properly understood, intentionality is a relatively weak criterion. It does not require a direct intention to *F*, but only what I characterized as an *indirect* intention to *F*; i.e. an intention to *G*, where *G*-ing also satisfies the conditions for *F*-ing. What is more, this result helped to show that intention-dependence is distinct from *concept*-dependence, and that concept-dependence is far too exclusive a condition for theories of art-making. Ultimately, I argued that indirect intention-dependence offers us the most plausible gloss on our commitment to art's intention-dependence.

Whereas Chapter 2 considered *successful* art-attempts, Chapter 3 focused on their *failure*. I began by distinguishing between two senses of 'failure', the conformative and the performative, and showed that these senses of failure attach to an attempt's direction (i.e. whether it is directly or indirectly intentional). Armed with this distinction I argued, against other treatments of artistic failure, that art's intention-dependence entails the possibility of failure *in*

general, not of failed-art in particular. I also argued that existing treatments of art-attempts and their failure have overemphasized the final product's conformity to the artist's intentions, resulting in a descriptively inadequate treatment of the process of art-making. I argued instead that art-making involves a complex hierarchy of intentions, and that the failure of lower-level intentions need not entail the failure of higher-level intentions. In fact, I argued that art-making presents, phenomenologically, as a series of low-level plans, obstacles, failures, and new plans to take care of the obstacles and failures that presented themselves. We must therefore postpone asking whether a particular art-attempt is successful until the agent ceases her activity, at which point it is plain to see that 'the manner intended' is negotiable at every step of the process. These considerations prompted me to offer an amended theory of failed art-attempts which I think is better equipped to handle the nuts and bolts of artistic practice.

Having explained the mechanics of art's intention-dependence, I then turned my attention towards extirpating the spectre of concept-dependence. Concept-dependence resurfaced in Chapter 4 as the suggestion that art-making requires a direct intention to make a work belonging to one of the arts (i.e. dance, literature, music, painting, sculpture, etc.). Simply put, the idea behind this suggestion is just that being 'art' is a question of belonging to an art-kind. Although this idea has gained currency in recent years in the form of a buck-passing theory of art, I argued that, properly considered, even these theories are committed to art's indirect intention-dependence. That is because, at the very least, the relevant actions must be analysed as *X*-attempts (e.g. music-attempts, painting-attempts, etc.), and once again the proper analysis of attempts requires indirect intention-dependence. What is more, such theories owe us an answer to the question of what makes a particular practice (or kind) an *artistic* practice (or *art*-kind) in the first place. Kinds and practices cannot be artistic simply in virtue of their associated physical

media, since these do not suffice to distinguish art from non-art, or to account for the origins of artistic practices. What is required, instead, is an appeal to *conventions* and, in particular, to conventions as normative tendencies that propagate by weight of historical precedent. But if kind-membership is a matter of convention, then the only way to explain the very first members of a kind is by reference to indirect intention-dependence.

Chapter 5 shifted the focus of my attention from art's ontology to its meta-ontology, to a consideration of the philosophy of art's proper methodology. My aim there was to show that the philosophy of art cannot simply be an exercise in conceptual analysis. I began by considering descriptivism as a concept-driven approach to art's ontology, and argued that we actually have very good reasons to doubt the correctness of our intuitions about art's ontology. I showed that art's conventional nature affects the salience of particular (art-ontological) possibilities, thus influencing the field of plausible possibilities irrespective of the ontological facts. I then offered an alternative model for art's ontology with good prospects for identifying instances of 'art' cross-culturally. This was a practice-driven approach to ontology according to which it is our artworld practices, not our concepts or intuitions, which constrain our ontological investigations. Our concepts and intuitions might well be misleading or wrong; our practices, however, are what they are.

Finally, Chapter 6 tackled the question of how it is that we can be substantially wrong about the nature of a social kind like 'art', whose essence is rooted in human conventions. I answered the challenge by explaining how it is that we can be substantially wrong about the nature of *natural* kinds, arguing that the reference of social kind-terms is fixed in much the same way. Just as our beliefs about natural kinds have no bearing on their ontological properties, neither do our beliefs about social kinds like 'art' do any work in fixing their ontological

properties or reference. What matters, for natural and social kinds alike, is the referent's capacity to play a certain kind of explanatory role in our practices. If the philosophy of art is to be an explanatory enterprise, then the result is just that we cannot be wrong about which entities it is that we specify as the subject of our inquiries—no more, and no less. In this way, we can reconcile our descriptivist leanings with the epistemic humility mandated by art's intention-dependence and its conventional origins. We have few, if any, guarantees that our artworld data are substantially accurate. The correct response, I argued, is not to treat that data as a set of conclusions to be explained by our theories, but rather as defeasible *starting points* for our investigations.

The intuition that art is intention-dependent is no exception. We might well be wrong that art is intention-dependent. In fact, we have seen that at least some treatments of intention-dependence—those which read it as direct intention-dependence—are almost certainly wrong. All things considered, however, it seems unlikely that art is not intention-dependent. This is because, as I have shown, intention-dependence is so closely intertwined with our artistic practices, theories, and the explanatory apparatus offered by analysing art's origins in terms of conventions. In a way, this entire monograph has been an exercise in illustrating the practice-driven approach to the ontology of art. I began by considering what could possibly be meant by the assertion that art is intention-dependent, analysed the constraints which the most plausible interpretation imposed upon theories of art-making, and then considered what our artistic practices had to say on the matter. Having done so, I showed that intention-dependence has the capacity to ground powerful explanatory generalizations about art. This may be the best we can do, but it is no mere consolation prize: it offers us the outline of an account of the proper role of intuitions in the philosophy of art.

Intuitions have recently come under a great deal of philosophical scrutiny. Often, they give us the starting point from which to begin some sort of conceptual analysis. This is how philosophers typically use thought-experiments; they describe a fictional scenario with tightly controlled parameters, and ask us to form a judgement about the likely outcome of the situation. The resulting judgements form the starting-point for investigations into the phenomenon at issue. Just as frequently, however, intuitions are cited as data used to test candidate theories for their plausibility—just as we have done for decades in the philosophy of art. So, for instance, we think that musical works are the kinds of things one *hears* by attending a concert or tuning in to a radio station, and any theory which tells us that they are the kinds of things we *see* hanging on gallery walls is just a bad theory of music, or a theory of something altogether different. Yet these are only *hunches*; why should we think they offer us ontological data which theories of art and art-kinds ought to respect?

As we have seen, a growing body of empirical art scholarship (from anthropology, art history, and social psychology) and philosophical argument indicates that our intuitions about art reflect entirely arbitrary historical interests. This monograph represents an argument to the effect that the ontology of art and art-kinds is a fundamentally explanatory enterprise. As such, the ontological properties at issue are not constrained by our intuitions, but rather by our best explanatory accounts of our artistic practices. In other words, a theory of art's plausibility does not depend upon the number of our intuitions which it captures, but on its ability to ground the explanatory hypotheses of empirical art scholarship.[349]

If my treatment of art's intention-dependence is to believed, then I hope to have established that at least one of our reflective commitments about art (viz. intention-dependence) is correct. It remains to be shown which of our other folk or reflective intuitions also merit

---

[349] Lopes makes just this point in his (2016).

assent. These might include, for example, the thesis that aesthetic properties are context-dependent,[350] which enjoys widespread philosophical support but which might seem suspicious to lay audiences. Alternately, it might include intuitions that are philosophically controversial but which are widely accepted by non-philosophers, such as the intuition that artworks are modally and temporally flexible,[351] or that certain kinds of works—paradigmatically, musical works—are created rather than discovered. But that, as they say, is another story.

---

[350] So that, e.g., in Jorge Luis Borges's *Pierre Menard, Author of the Quixote*, Menard does not just reproduce Cervantes's classic novel. Cervantes's is a novel about his own times and is written in his own idiom, while Menard's is a historical novel written in archaic prose.

[351] i.e. they could have been otherwise. So, for example, Tchaikovsky's *1812 Overture* could have called for *French* cannon fire to counter a fragment of *God Save the Tsar* rather than the other way around, or could have contained some different notes.

Works Cited

Andina, T. (2013). *The philosophy of Art: The Question of Definition - From Hegel to Post-Dantian Theories.* (N. Iacobelli, Trans.) New York: Bloomsbury.

Anscombe, G. (2000 [1957]). *Intention* (Second Edition ed.). Cambridge, Massachussetts: Harvard University Press.

Arnheim, R. (1957). A New Laocoön: Artistic Composites and the Talking Film. In *Film as Art* (pp. 199-230). Los Angeles: University of California Press.

Batteux, C. (1753). *Les Beaux Arts Reduits a un même principe.* Leide: Elie Luzac, Fils.

Baxandall, M. (1985). *Patterns of Intention: on the historical explanation of pictures.* New Haven: Yale University Press.

Baxandall, M. (1988 [1972]). *Painting and Experience in Fifteenth-Century Italy: A Primer in the Social History of Pictorial Style.* New York: Oxford University Press.

Beardsley, M. (1958). *Aesthetics: Readings in the Philosophy of Criticism.* New York: Harcourt, Brace.

Beardsley, M. (1976). Is Art Essentially Institutional? In L. Aagaard-Mogensen (Ed.), *Culture and Art.* Atlantic Highlands: Humanities Press.

Beardsley, M. (1982). *The Aesthetic Point of View.* Ithaca: Cornell University Press.

Beardsley, M. C. (1982). What Is Going on in a Dance? *Dance Research Journal, 15*(1), 31-6.

Bell, C. (1913). *Art.* New York: Frederick A. Stokes.

Bennett, D. (1965). Action, reason, and purpose. *Journal of Philosophy, 62*(4), 85-96.

Binkley, T. (1977). Piece: Contra Aesthetics. *The Journal of Aesthetics and Art Criticism, 35*(3), 265-277.

Bonk, E. (1995). Marcel Duchamp: The Woolworth Building as Readymade, January 1916 (An Approximation). *Grand Street*, 165-75.

Bornstein, R. F. (1989). Exposure and Affect: Overview and Meta-Analysis of Research, 1968-87. *Psychologial Bulletin, 106*(2), 265-89.

Brand, P. Z. (2000). Glaring Omissions in Traditional Theories of Art. In N. Carroll (Ed.), *Theories of Art Today* (pp. 175-83). Madison, Wisconsin: The University of Wisconsin Press.

Brandom, R. B. (1994). *Making It Explicit: Reasoning, Representing, & Discursive Commitment.* Cambridge, Massachusetts: Harvard University Press.

Bratman, M. E. (1987). *Intention, Plans, and Practical Reason.* Cambridge, Massachusets: Harvard University Press.

Bullough, E. (1957). *Aesthetics: Lectures and Essays.* (E. M. Wilkinson, Ed.) London: Bowes & Bowes.

Caplan, B., & Matheson, C. (2006). Defending Musical Perdurantism. *British Journal of Aesthetics, 46*(1), 59-69.

Cardinal, R. (1994). Toward an Outsider Aesthetic. In M. D. Hall, & W. W. Metcalf (Eds.), *The Artist Outsider: Creativity and the Boudnaries of Culture* (pp. 20-43). Washington: Smithsonian Institution Press.

Carney, J. (1991). The style theory of art. *Pacific Philosophical Quarterly, 72*(4), 272-89.

Carroll, N. (1988). Art, Practice, and Narrative. *The Monist, 71*(2), 140-56.

Carroll, N. (1993). Historical Narratives and the Philosophy of Art. *Journal of Aesthetics and Art Criticism, 51*, 313-26.

Carroll, N. (1999). *Philosophy of Art.* London: Routledge.

Carroll, N. (2003). Forget the Medium! In N. Carroll, *Engaging the Moving Image* (pp. 1-9). New Haven: Yale University Press.

Cheshire, T. (2013, August 29). *Hacking meets clubbing with the 'algorave'.* Retrieved May 7, 2016, from Wired: http://www.wired.co.uk/magazine/archive/2013/09/play/algorave

Chisholm, R. (1966). Freedom and Action. In Lehrer (Ed.), *Freedom and Determinism* (pp. 11-44). New York, New York: Random House.

Cleveland, T. (1997). *Trying Without Willing.* Aldershot: Ashgate Publishing.

Clowney, D. (2011). Definitions of Art and Fine Art's Historical Origins. *The Journal of Aesthetics and Art Criticism, 69*(3), 309-20.

Collingwood, R. (1938). *The Principles of Art.* Oxford: Clarendon Press.

Constant, B. (1952). *Journaux intimes.* (A. Roulin, & C. Roth, Eds.) Paris: Gallimard.

Croce, B. (1992 [1909]). *The aesthetic as the science of expression and of the linguistic in general.* (C. Lyas, Trans.) New York: Cambridge University Press.

Croce, B. (1992). *The aesthetic as the science of expression and of the linguistic in general.* (C. Lyas, Trans.) New York: Cambridge University Press.

Currie, G. (1989). *An Ontology of Art.* London: St. Martin's.

Currie, G. (2014). On getting out of the armchair to do aesthetics. In M. C. Haug (Ed.), *Philosophical Methodology: The Armchair or the Laboratory?* (pp. 435-50). New York: Routledge.

Cutting, J. (2003). Gustave Caillebotte, French Impressionism, and Mere Exposure. *Psychonomic Bulletin and Review, 10*, 319-43.

Danto, A. C. (1964, October 15). The Artworld. *The Journal of Philosophy, 61*(19), 571-84.

Danto, A. C. (1965). Basic Actions. *American Philosophical Quarterly, 2*(2), 141-8.

Danto, A. C. (1981). *The Transfiguration of the Commonplace.* Harvard University Press.

Davidson, D. (1963). Actions, Reasons, and Causes. *Journal of Philosophy, 60*(23), 685-700.

Davidson, D. (1991 [1977]). The Method of Truth in Metaphysics. In D. Davidson, *Inquiries Into Truth & Interpretation* (pp. 199-214). Toronto: Oxford University press.

Davidson, D. (2001 [1980]). *Essays on Actions and Events* (Second Edition ed.). Oxford: Clarendon Press.

Davidson, D. (2001). *Essays on Actions and Events* (Second Edition ed.). Oxford: Clarendon Press.

Davies, D. (2004). *Art as Performance.* Oxford: Blackwell.

Davies, D. (2009a). Dodd on the 'Audibility' of Musical Works. *British Journal of Aesthetics, 49*(2), 99-108.

Davies, D. (2009b). On the very idea of 'Outsider Art'. *British Journal of Aesthetics, 49*(1), 25-41.

Davies, D. (2009c). The Primacy of Practice in the Ontology of Art. *Journal of Aesthetics and Art Criticism, 67*(2), 159-71.

Davies, D. (2010a). Medium in Art. In J. Levinson (Ed.), *The Oxford Handbook of Aesthetics* (pp. 181-191). Oxford University Press.

Davies, D. (2010b). Multiple Instances and Multiple 'Instances'. *British Journal of Aesthetics, 50*(4), 411-26.

Davies, D. (2011). *Philosophy of the Performing Arts.* Malaysia: Wiley-Blackwell.

Davies, D. (2012). Enigmatic Variations. *The Monist, 9*(4), 644-63.

Davies, D. (2013). Pornography, Art, and the intended response of the receiver? In J. Levinson, & H. Maes (Eds.), *Pornographic Art and the Aesthetics and Pornography.* Oxford: Oxford University press.

Davies, D. (2015, September 04). *Dominic McIver Lopes: Beyond Art.* Retrieved July 28, 2016, from Notre Dame Philosophical Reviews: http://ndpr.nd.edu/news/60352-beyond-art/

Davies, D. (2016). The Function of Generalization in Art History: Understanding Art across Traditions. *Arts and Literary Studies*, 8-19.

Davies, D. (forthcoming a). Descriptivism and its Discontents. *Journal of Aesthetics and Art Criticism*.

Davies, D. (forthcoming b). Reflections on Mag Uidhir's Art and Art Attempts. *Journal of Aesthetic Education*.

Davies, S. (1991). *Definitions of Art.* Ithaca: Cornell University Press.

Davies, S. (2000). Non-Western Art and Art's Definition. In N. Carroll (Ed.), *Theories of Art Today* (pp. 217-40). Madison: University of Wisconsin Press.

Davies, S. (2001). *Musical Works and Performances: A Philosophical Exploration.* Oxford: Oxford University Press.

Davies, S. (2003). Essential Distinctions for Art Theorists. In S. Davies, & A. C. Sukla (Eds.), *Art and Essence* (pp. 3-16). Westport, Connecticut: Praeger.

Davies, S. (2007a). Balinese Aesthetics. *Journal of Aesthetics and Art Criticism, 65*(1), 21-9.

Davies, S. (2007b). Is Architecture Art? In S. Davies (Ed.), *Philosophical Perspectives on Art* (pp. 129-47). Oxford: Oxford University Press.

Davies, S. (2007c). *Philosophical Perspectives on Art.* New York: Oxford University Press.

Davies, S. (2008, Spring). Trying to Define Art as the Sum of the Arts. *Pazhouhesh Nameh-e Farhangestan-e Honar, 8*, 12-23.

Davies, S. (2016). Aesthetics in cross-cultural perspective. *Theoretical Studies in Literature and Art, 36*(1), 20-6.

Davis, W. (2003). *Meaning, Expression, and Thought.* New York: Cambridge University Press.

de Duve, T. (1990). The Monochrome and the Blank Canvas. In S. Guilbaut (Ed.), *Reconstructing Modernism: Art in New York, Paris, and Montreal, 1945–1964* (pp. 244-310). Cambridge: MIT Press.

de Duve, T. (1996). *Kant after Duchamp.* Cambridge: MIT Press.

de Smedt, J., & de Cruz, H. (2011). A Cognitive Approach to the Earliest Art. *The Journal of Aesthetics and Art Criticism, 69*(4), 379-389.

Dean, J. T. (2003). The Nature of Concepts and the Definition of Art. *Journal of Aesthetics and Art Criticism, 61*(1), 29-35.

Dickie, G. (1974). *Art and the Aesthetic: An Institutional Analysis.* Ithaca: Cornell University Press.

Dickie, G. (1984). *The Art Circle: A Theory of Art.* New York: Haven.

Dickie, G. (2000). The Institutional Theory of Art. In N. Carroll (Ed.), *Theories of Art Today* (pp. 93-108). Madison: The University of Wisconsin Press.

Dilworth, J. (2003). Medium, subject matter and representation. *Southern Journal of Philosophy, 41*, 45-62.

Dipert, R. (1980). The Composer's Intentions: An Examination of their Relevance forPerformance. *Musical Quarterly, 66*, 205-18.

Dissanayake, E. (1988). *What Is Art For?* Seattle: University of Washington Press.

Dodd, J. (2007). *Works of Music: An Essay in Ontology.* Oxford: Oxford University Press.

Dodd, J. (2008). Musical Works: Ontology and Meta-Ontology. *Philosophy Compass, 3*(6), 1113-34.

Dodd, J. (2013). Adventures in the Metaontology of Art: Local Descriptivism, Artefacts and Dreamcatchers. *Philosophical Studies, 165*(3), 1047-1068.

Dubuffet, J. (1973). Honneur aux valeurs sauvages. In *L'Homme du commun à l'ouvrage* (pp. 93-118). Paris: Gallimard.

Dutton, D. (1995). Mythologies of Tribal Art. *African Arts, 28*, 32-43.

Dutton, D. (2000). But They Don't Have Our Concept of Art. In N. Carroll (Ed.), *Theories of Art Today* (pp. 217-40). Madison: University of Wisconsin Press.

Dutton, D. (2009). *The Art Instinct: Beauty, Pleasure, & Human Evolution.* New York: Bloomsbury Press.

Edidin, A. (1991). Look What They've Done to My Song: "Historical Authenticity" and the Aesthetics of Musical Performance. *Midwest Studies in Philosophy, XVI*, 393-420.

Elkins, J. (2006). Naïfs, Faux-naïfs, Faux-faux naïfs, Would-be Faux-naïfs: There is No such Thing as Outsider Art. In J. Thompson (Ed.), *Inner Worlds Outside* (pp. 71-9). Dublin: Irish Museum of Modern Art.

Emerson, P. H. (1974 [1889]). *Naturalistic Photography.* London: Creative Camera.

Empiricus, S. (1996). *Outlines of Pyrronism.* (B. Yates, Trans.) New York: Oxford University Press.

Evans, G. (1973). The Causal Theory of Names. *Proceedings of the Aristotelian Society, Supplementary Volumes, 47*, 187-225.

Evans, G. (1982). *The Varieties of Reference.* (J. McDowell, Ed.) New York: Oxford University Press.

Feagin, S. L. (1995). Paintings and their places. *Australasian Journal of Philosophy, 73*(2), 260-8.

Freeland, C. (2001). *But Is It Art? An Introduction to Art Theory.* New York: Oxford University Press.

Gaut, B. (2000). "Art" as a Cluster Concept. In N. Carroll (Ed.), *Theories of Art Today* (pp. 25-44). Madison: University of Wisconsin Press.

Ginet, C. (1990). *On Action.* New York: Cambridge University Press.

Gisbourne, M. (1994). French Clinical Psychiatry and the Art of the Untrained Mentally Ill. In M. D. Hall, & E. W. Metcalf (Eds.), *The Artist Outsider: Creativity and the Boundaries of Culture* (pp. 228-51). Washington: Smithsonian Institution Press.

Goehr, L. (1992). *The Imaginary Museum of Musical Works: An Essay in the Philosophy of Music.* Oxford: Oxford University Press.

Goldman, A. I. (1970). *A Theory of Human Action.* Princeton, New Jersey: Princeton University press.

Gombrich, E. (1984 [1960]). *Art and Illusion: A Study in the Psychology of Pictorial Representation.* London: Phaidon Press.

Gombrich, E. (1985, July-September). Image and word in twentieth-century art. *Word & Image, 1*(3), 213-41.

Gombrich, E. H. (1971). *Meditations on a Hobby Horse and other essays on the theory of art.* New York: Phaidon.

Goodman, N. (1968). *Languages of art : an approach to a theory of symbols.* Indianapolis: Bobbs-Merrill.

Greenberg, C. (1960). *Modernist Painting.* Retrieved July 2013, from University of Chicago Council on Advanced Studies: http://cas.uchicago.edu/workshops/wittgenstein/files/2007/10/Greenbergmodpaint.pdf

Haack, S. (1979). Descriptive and Revisionary Metaphysics. *Philosophical Studies, 35*, 361-71.

Hacking, I. (1995). The Looping Effects of Human Kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal Cognition: a multi-disciplinary debate* (pp. 351-83). Oxford: Oxford University Press.

Hacking, I. (1999). *The Social Construction of What?* Cambridge, Massachusetts: Harvard University Press.

Hamilton, A. (2007). *Aesthetics and Music.* New York: Continuum.

Hamilton, J. R. (2007). *The Art of Theater.* Malden, MA: Blackwell Publishing Ltd.

Hanslick, E. (1891 [1885]). *The Beautiful in Music: A Contribution to the Revisal of Musical Aesthetics.* (G. Cohen, Trans.) London: Novello & Company, Ltd.

Hart, L. M. (1995). Three Walls: Regional Aesthetics and the International Art World. In G. E. Marcus, & F. R. Myers (Eds.), *The Traffic in Culture: Refiguring Art and Anthropology* (pp. 127-50). Berkeley: University of California Press.

Haslanger, S. (1995). Ontology and Social Construction. *Philosophical Topics, 23*(2), 95-125.

Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique.* New York: Oxford University Press.

Hassig, R. (1988). *Aztec warfare : imperial expansion and political control.* Norman: University of Oklahoma Press.

Hazlett, A. (2012). Against Repeatable Artworks. In C. Mag Uidhir (Ed.). Oxford: Oxford University Press.

Hegel, G. (2001 [1835]). *Hegel's Aesthetcs: Lectures on Fine Art* (Vol. 1). (T. Knox, Trans.) Oxford: Clarendon Press.

Hicks, F. (1979). "Flowery War" in Aztec History. *American Ethnologist, 6*(1), 87-92.

Higginbotham, J., Pianesi, F., & Varzi, A. C. (Eds.). (2000). *Speaking of Events.* New York: Oxford University Press.

Hilpinen, R. (1992). On artifacts and works of art. *Theoria, 52*(1), 58-82.

Hirsch, E. (1967). In Defense of the Author. In E. Hirsch, *Validity in Interpretation* (pp. 1-23). New Haven: Yale University Press.

Holm, B. (1970 [1965]). *Northwest Coast Indian Art: An Analysis of Form.* Seattle: University of Washington Press.

Hornsby, J. (1980). *Actions.* Boston: Routledge & Kegan Paul.

Isaac, B. L. (1983). The Aztec "Flowery War": A Geopolitical Explanation. *Journal of Anthropological Research, 39*(4), 415-32.

Irvin, S. (2012). Installation Art and Performance: A Shared Ontology. In C. Mag Uidhir (Ed.), *Art & Abstract Objects* (pp. 242-62). New York: Oxford University Press.

Iseminger, G. (1992). An Intentional Demonstration? In G. Iseminger (Ed.), *Intention & Interpretation* (pp. 76-96). Philadelphia: Temple University Press.

Kania, A. (2008). The Methodology of Musical Ontology: Descriptivism and its Implications. *British Journal of Aesthetics, 48*(4), 426-44.

Kania, A. (2011). Definition. In T. Gracyk, & A. Kania (Eds.), *The Routledge Companion to Philosophy and Music* (pp. 1-13). New York: Routledge.

Kant, I. (1987 [1790]). *Critique of Judgment.* (W. S. Pluhar, Trans.) Indianapolis: Hackett Publishing Company.

Kasfir, S. L. (1992, April). African Art and Authenticity: A Text with a Shadow. *African Arts, 25*(2), 40-53+96-97.

Keil, C. (1979). *Tiv Song.* Chicago: University of Chicago Press.

Khalidi, M. A. (2010). Interactive Kinds. *British Journal for the Philosophy of Science, 61*, 335-60.

Kieran, M. (2001). Pornographic Art. *Philosophy and Literature,, 25*, 31-45.

Kim, J. (1976). Events as Property Exemplifications. In M. Brand, & D. Walton (Eds.), *Action Theory* (pp. 310-26). Dordrecht: D. Reidel.

Kim, J. (1989). Mechanism, Purpose, and Explanatory Exclusion. *Philosophical Perspectives, 3*, 77-108.

Kivy, P. (1983). Platonism in Music. *Grazer Philosophische Studien, 19*, 109-29.

Kivy, P. (2012, January). What Really Happened in the Eighteenth Century: The 'Modern System' Re-examined (Again). *British Journal of Aesthetics, 52*(1), 61-74.

Koichi, M. (1992). Liner Notes. *Jegog of Negara*. (G. Groemer, Trans.) World Music Library KICC-5157.

Kristeller, P. O. (1951, October). The Modern System of the Arts: A Study in the History of Aesthetics Part I. *Journal of the History of Ideas, 12*(4), 496-527.

Kristeller, P. O. (1952, January). The Modern System of the Arts: A Study in the History of Aesthetics (II). *Journal of the History of Ideas, 13*(1), 17-46.

Kuspit, D. (2004). *The End of Art.* New York: Cambridge University Press.

Lessing, G. E. (1905 [1766]). *Laocoön.* London: George Routledge & Sons, ltd.

Levinson, J. (1979). Defining Art Historically. *British Journal of Aesthetics, 19*(3), 232-50.

Levinson, J. (1980). What A Musical Work Is. *Journal of Philosophy, 77*(1), 5-28.

Levinson, J. (1990). The Concept of Music. In J. Levinson, *Music, Art, and Metaphysics* (pp. 267-78). Ithaca, NY: Cornell University Press.

Levinson, J. (1996). The Work of Visual Art. In J. Levinson, *The Pleasures of Aesthetics* (pp. 129-37). Ithaca: Cornell University Press.

Levinson, J. (2005). Erotic Art and Pornographic Pictures. *Philosophy and Literature, 29*(1), 228-40.

Levinson, J. (2011). What a Musical Work Is, Again. In J. Levinson, *Music, Art and Metaphysics* (pp. 215-63). New York: Oxford University Press.

Lewis, D. K. (1986). *On the Plurality of Worlds.* Oxford: Basil Blackwell.

Lewis, D. K. (2002 [1969]). *Convention: A Philosophical Study.* Malden: Blackwell Publishers.

Locher, H. (2012). The Idea of the Canon and Canon Formation in Art History. In M. Rampley, T. Lenain, H. Locher, A. Pinotti, C. Schoell-Glass, & K. Zijlmans (Eds.), *Art History and*

Visual Studies in Europe: Transnational Discourses and National Frameworks (Vol. 4, pp. 29-40). Boston: Brill.

Lopes, D. M. (2007, January). Art Without 'Art'. *British Journal of Aesthetics, 47*(1), 1-15.

Lopes, D. M. (2008, March). Nobody Needs a Theory of Art. *The Journal of Philosophy, 105*(3), 109-27.

Lopes, D. M. (2014). *Beyond Art.* New York: Oxford University Press.

Lopes, D. M. (2016). Aesthetics and Philosophy of Art. In H. Cappelen, T. S. Gendler, & J. Hawthorne (Eds.), *The Oxford Handbook of Philosophical Methodology* (pp. 657-70). New York: Oxford University Press.

Mag Uidhir, C. (2009). Why Pornography Can't Be Art. *Philosophy and Literature*, 193-203.

Mag Uidhir, C. (2010). Failed-Art and Failed Art-Theory. *Australasian Journal of Philosophy, 88*(3), 381-400.

Mag Uidhir, C. (2013). *Art and Art-Attempts.* Oxford: Oxford University Press.

Mandelbaum, M. (1965, July). Family Resemblances and Generalization concerning the Arts. *American Philosophical Quarterly, 2*(3), 219-28.

Margolis, J. (1980). *Art and Philosophy.* Atlantic Highlands, N.J.: Humanities Press.

McCann, H. (1974). Volition and Basic Action. *Philosophical Review, 83*, 451-73.

McKoon, G., & Ratcliffe, R. (1992). Inference During Reading. *Psychological Review, 99*(3), 440-66.

Mele, A., & Moser, P. K. (1994, March). Intentional Action. *Noûs, 28*(1), 39-68.

Meskin, A., Phelan, M., Moore, M., & Kieran, M. (2013). Mere Exposure to Bad Art. *British Journal of Aesthetics*, 139-64.

Millikan, R. G. (2001 [1984]). *Language, Thought, and Other Biological Categories: New Foundations for Realism.* Cambridge: MIT Press.

Moravcsik, J. M. (1993). Why Philosophy of Art in Cross-Cultural Perspective? *Journal of Aesthetics and Art Criticism, 51*(3), 425-35.

Moya, C. J. (1990). *The Philosophy of Action: An Introduction.* Oxford: Polity Press.

Nagel, T. (1979). *Mortal Questions.* New York: Cambridge University Press.

Naumann, F.M. & Obalk, H. (Eds.) (2000). *Affectionately, Marcel: The Selected Correspondence of Marcel Duchamp.* Ghent: Ludion Press.

Nochlin, L. (1988 [1977]). Why have there been no great women artists? In L. Nochlin, *Women, art, and power : and other essays* (pp. 145-78). New York: Harper & Row.

Novitz, D. (1998). Art by Another Name. *British Journal of Aesthetics, 38*, 19-32.

O'Shaughnessy, B. (1973). Trying (as the Mental 'Pineal Gland'). *Journal of Philosophy, 70*, 365-86.

Osipovich, D. (2006). What is a Theatrical Performance? *The Journal of Aesthetics and Art Criticism, 64*(4), 461-70.

Passmore, J. A. (1951, July). The Dreariness of Aesthetics. *Mind, New Series, 60*(239), 318-35.

Paul, S. K. (2014). Embarking on a Crime. In E. V. Villanueva (Ed.), *Law and the Philosophy of Action* (pp. 101-24). Rodopi.

Pollock, G. (1999). *Differencing the canon: feminist desire and the writing of art's histories.* New York: Routledge.

Porter, J. I. (2009, January). Is Art Modern? Kristeller's 'Modern System of the Arts' Reconsidered. *British Journal of Aesthetics, 49*(1), 1-24.

Putnam, H. (1975). The meaning of 'meaning'. In *Mind, Language, and Reality: Philosophical Papers* (Vol. 2, pp. 215-71). Cambridge: Cambridge University Press.

Putnam, H. (1990). Is Water Necessarily H20? In H. Putnam, & J. Conant (Ed.), *Realism With a Human Face* (pp. 54-79). Cambridge, MA: Harvard University Press.

Raffman, D. (1992). *Language, Music, and Mind.* Boston: MIT Press/Bradford.

Rawls, J. (1955). Two Concepts of Rules. *Philosophical Review, 3*, 3-32.

Rohrbaugh, G. (2003). Artworks as historical individuals - lessons from photography. *European Journal of Philosophy, 11*(2), 177-205.

Rudder Baker, L. (2004). The ontology of artifacts. *Philosophical Explorations, 7*(2), 99-111.

Santayana, G. (1896). *The sense of beauty : being the outlines of aesthetic theory.* New York: Charles Scribner's Sons.

Savile, A. (1969). The Place of Intention in the Concept of Art. *Proceedings of the Aristotelian Society, New Series, 69*, 101-24.

Schopenhauer, A. (1966 [1818]). *The World as Will and Representation, volume II.* (E. Payne, Trans.) Mineola: Dover.

Schopenhauer, A. (2010 [1818]). *The World as Will and Representation, volume I.* (C. Janaway, Ed., J. Norman, & A. Welchman, Trans.) Cambridge: Cambridge University Press.

Shermer, M., & Grobman, A. (2009 [2000]). *Denying History: Who Says the Holocaust Never Happened, and Why Do They Say It?* Berkeley: University of California Press.

Shiner, L. (1994). "Primitive Fakes," "Tourist Art," and the Ideology of Authenticity. *Journal of Aesthetics and Art Criticism, 52*(2), 225-34.

Shiner, L. (2001). *The Invention of Art: A Cultural History.* Chicago: The University of Chicago Press.

Shiner, L. (2003). Western and Non-Western Concepts of Art: Universality and Authenticity. In S. Davies, & A. C. Sukla (Eds.), *Art and Essence* (pp. 143-56). Westport, Connecticut: Praeger.

Stecker, R. (1994). Historical functionalism or the four factor theory. *British Journal of Aesthetics, 34*(3), 255-65.

Stecker, R. (1997). *Artworks: Definition, Meaning, Value.* University Park, PA: The Pennsylvania State University Press.

Stecker, R. (2010 [2005]). *Aesthetics and the Philosophy of Art: An Introduction* (Second Edition ed.). Toronto: Rowman & Littlefield.

Stern, T. (2000). *Rehearsal from Shakespeare to Sheridan.* New York: Oxford University Press.

Stieglitz, A. (1899, November). Pictorial Photography. *Scribner's Magazine*, pp. 528-37.

Stock, K. (2003). Historical Definitions of Art. In S. Davies, & A. C. Sukla (Eds.), *Art and Essence* (pp. 159-76). Westport, Connecticut: Praeger.

Strawson, P. (2003 [1959]). *Individuals: An Essay in Descriptive Metaphysics.* New York, New York: Routledge.

Strawson, P. F. (2008 [1974]). Aesthetic Appraisal and Works of Art. In P. F. Strawson, *Freedom and Resentment and Other Essays* (pp. 196-207). New York: Routledge.

Thomasson, A. L. (1999). *Fiction and Metaphysics.* New York: Cambridge University Press.

Thomasson, A. L. (2005). The Ontology of Art and Knowledge in Aesthetics. *Journal of Aesthetics and Art Criticism, 63*(3), 221-9.

Thomasson, A. L. (2006). Debates about the Ontology of Art - What Are We Doing Here. *Philosophy Compass, 1*(3), 245-55.

Thomasson, A. L. (2007). *Ordinary Objects.* New York: Oxford University Press.

Tolstoy, L. (1960 [1899]). *What is art?* (A. Mauder, Trans.) New York: Macmillan Pub. Co.

Vogel, S. M. (1997). *Baule: African Art, Western Eyes.* New Haven: Yale University Press.

Walton, K. (1970). Categories of Art. *The Philosophical Review, 79*, 334-67.

Weitz, M. (1950). *Philosophy of the Arts.* Cambridge, MA: Harvard University Press.

Weitz, M. (1956, September). The Role of Theory in Aesthetics. *The Journal of Aesthetics and Art Criticism, 15*(1), 27-35.

Wilcox, J. (1953). The beginnings of l'art pour l'art. *Journal of Aesthetics and Art Criticism, 11*(4), 360-377.

Williams, B. (1981). *Moral Luck.* Cambridge: Cambridge University Press.

Wilson, G. (1989). *The Intentionality of Human Action.* Stanford: Stanford University Press.

Wimsatt, W. K., & Beardsley, M. (1946). The Intentional Fallacy. *Sewanee Review, 54*, 468-88.

Wollheim, R. (1980 [1968]). *Art and Its Objects* (Second Edition ed.). Cambridge: Cambridge University Press.

Wolterstorff, N. (1980). *Works and Worlds of Art.* Oxford: Clarendon Press.

Yaffe, G. (2010). *Attempts: In the Philosophy of Action and the Criminal Law.* New York: Oxford University Press.

Young, J. O. (2015). The Ancient and Modern System of the Arts. *British Journal of Aesthetics, 55*(1), 1-17.

Young, J. O. (2016). Review of Dominic McIver Lopes, Beyond Art. *Philosophy in Review, 36*(2), 83-5.

Zajonc, R. B. (1968). Attitudinal Effects of Mere Exposure. *Journal of Personality and Social Psychology Monographs, 9*, 1-27.

Zangwill, N. (1995, October). Groundrules in the Philosophy of Art. *Philosophy, 70*(274), 533-44.

Zangwill, N. (2007). *Aesthetic Creation.* Oxford: Oxford University Press.

Index

176