The Epistemology of Identity

Samuel Z. Elgin

Abstract

The subject of this paper is the epistemology of identity: a general theory of knowledge, evidence and justification for the claim that one thing is identical to another. Although identity figures significantly in our epistemic lives, this is a topic that, to the best of my knowledge, has gone entirely unexplored. Initial attempts to integrate such an epistemology into existing theories of evidence—many of which are tailor-made for contingent propositions—are confounded by the necessity of identity. After discussing several proposals that are moderately successful (at best), I defend an account which relies upon Leibniz's Law and spatiotemporal coincidence.

1 Introduction

Identity gives rise to challenging questions which are not altogether easy to answer. It initially appears trivial, uninteresting, and a topic we already know a great deal about; it is that relation everything stands in to itself and to nothing else. Its logical attributes are widely known and entirely uncontroversial; it is transitive, reflexive and symmetric. In this respect, it seems utterly inconsequential and unworthy of philosophical investigation. Nevertheless, identity has generated a number of puzzles, some of which have troubled philosophers since the inception of the analytic tradition, and others of which have for much longer.

Perhaps the oldest of these puzzles concerns personal identity—what it is that makes a person at one time identical to a person at another.[1] This is especially pressing if, as some have thought, moral properties like responsibility and desert track identical persons. Another concerns the meanings of sentences that express identity claims—how it could be that 'Hesperus is Hesperus' differs in cognitive import from 'Hesperus is Phosphorus.'[2] While this is often discussed in light of the meanings of proper names, it directly concerns why some identity claims are substantive while others are trivial. Yet another concerns the modal profile of identity—both whether it is necessary or contingent and whether it holds between objects located in different possible worlds.[3] And some radical philosophers

[1] This literature is so large, and I suspect that most readers are so familiar, that it is probably pointless to provide references. However, see, e.g., Locke (1690); Hume (1738); Reid (1785); Parfit (1971); Lewis (1976); Shoemaker (1984); Swinburne (1997).

[2] See, canonically, Frege (1892).

[3] See Kripke (1980); Lewis (1986).

have even argued that identity is contextually dependent: i.e., that sentences of the form 'a is identical to b' really mean 'a is the same F as b' where F is a contextually-determined count noun.[4]

Given the sustained attention devoted to such an (apparently) trivial relation, some may suspect that we have exhausted the subject-matter of identity. Existing puzzles have seen genuine progress—for example, it was a substantial advance when philosophers recognized that identity holds necessarily.[5] Perhaps we will continue to develop new answers to old questions and perhaps existing topics will be framed in novel and illuminating ways, but, by now, it is reasonable to think that we have discovered all of the important debates.

I maintain that this is false. There is a philosophically important aspect of identity that has gone entirely overlooked: its epistemology. How can we come to know identity claims?[6] What constitutes evidence for the conclusion that one thing is identical to another? When is the inference to identity epistemically warranted?[7] As far as I know, this is a topic that philosophy has not yet broached. Nevertheless, if identity figures significantly in our cognitive lives (and it seems clear to me that it does), it behooves us to better understand its epistemology.

I aim to systematically investigate this epistemology of identity—to provide the first account of when the inference to identity is warranted. Leibniz's Law and spatiotemporal overlap figure centrally in my positive proposal, although I spend substantial time discussing why other natural suggestions are ultimately unsuccessful. Section 2 contains the bulk of this negative account, and raises the skeptical possibility that there is no unified epistemology to be found. I respond to this challenge in section 3 by providing such an epistemology, and conclude by discussing some of its shortcomings in section 4. While I hope to convince readers of my account, I am even more concerned with developing interest in this topic. If I pique others' curiosity enough to propose rival views, I will take myself to have succeeded.

It is my hope that this project's significance is manifest. However, I recognize that some may harbor motivational concerns and suspect that identity is so insignificant that its epistemology is as well. I urge such disinterested readers to reconsider the philosophical roles identity plays; it is not nearly as trivial as it first appears.

[4] See Geach (1967).

[5] These developments were largely due to Kripke (1980). However, for a recent defense of contingent identity, see Kocurek (2018).

[6] A point of terminological clarification: by an 'identity claim' I mean a sentence (or proposition) of the form '$a = b$,' where both a and b are singular terms. Often, a and b are proper names, but I am also concerned with sentences of the form 'a is the F' and 'The F is the G,' so definite descriptions figure centrally as well. And while I primarily discuss knowledge of the identity of objects, I am also concerned with the identity of properties and relations. In addition, by an 'identity belief' I mean a belief whose content takes the form of an identity.

[7] There are philosophically important distinctions between these questions. For example, not all those who have evidence of an identity know that identity if they fail to respond appropriately to their evidence. However, these questions all fall under the broad umbrella of the epistemology of identity, and I will largely leave the important distinctions between them for future debates.

For example, some have argued that identity is crucial to our understanding of quantification.[8] In particular, the concept of identity may be covertly invoked whenever a variable occurs multiple times within the scope of a quantifier. Consider the difference between '$\exists xFx \wedge \exists yGy$' and '$\exists x(Fx \wedge Gx)$.' While the first sentence asserts that something is F and something (quite possibly something else) is G, while the second sentence asserts that *the very same thing* is both F and G. As Hawthorne says,"Without mastery of the concept of identity it is not clear how we would understand the significance of the recurrence of a variable within the scope of a quantifier" (Hawthorne, 2003, pg. 100). If this is correct, then identity is important, at least in part, because it figures in our understanding of quantification.

Another example concerns the connection between identity and cardinality. As Russell noted, it is possible to formalize facts about number with identity; the sentence 'There is one F' might be represented as: '$\exists x(Fx \wedge \forall y(Fy \rightarrow x = y))$.' For this reason, some philosophers maintain that the numerical claims we regularly employ are properly analyzed as statements involving identity (e.g., Hofweber (2005)), and some psychologists have offered empirical evidence that children employ identity when developing the concept of number (e.g., Carey (2009); Carey and Xu (1996)). Numerical issues are central to our daily lives (just consider the difference between the claim that there was one gunman from the claim that there were two), so perhaps identity derives significance from its connection to cardinality.[9]

My own interest in this topic arose from a different source. Recently, there has been renewed discussion of generalized identities—sentences of the form 'To be F is to be G,' which philosophers often use to express analyses, and which have occupied a central role in our discipline since antiquity.[10] Notable examples include: 'To be morally right is to maximize utility,' 'To be water is to be the chemical compound H_2O,' and 'To be even is to be a natural number divisible by two without remainder.' Some suspect that the 'is' of 'To be F is to be G' literally is the 'is' of identity; sentences of this form hold just in case the property of being F is identical to the property of being G.[11] If so, the epistemology of

[8] For arguments along this line, see, e.g., Quine (1964); Hawthorne (2003); McGinn (2000). For a critical response, see, e.g., Burgess (2018).

[9] Of course, it is also possible to represent number without employing the identity sign. Natural languages like English typically have numerical adjectives that work perfectly well. And it is technically possible to introduce additional quantifiers that perform the work attributed to identity. We could, for example, introduce '$\exists!$' as a primitive quantifier expressing unique existence. These tactics are certainly available—I do not mean to suggest that we *must* invoke identity—but those tempted by Russell's initial suggestion hold that identity performs important theoretical work.

[10] See, e.g., Dorr (2016); Correia (2017); Elgin (Forthcoming); Rayo (2013); Linnebo (2014).

[11] Others claim that the 'is' of 'To be F is to be G' merely *closely resembles* the 'is' of identity in at least its logical and modal profiles without its ontological commitments, preferring a theory that is compatible with nominalism. There is textual evidence that such philosophers also hold that these sentences epistemically resemble identity claims. Consider, e.g., "There is something odd about questions like 'Why is Hesperus Phosphorus?'. Unless this is understood as a request to be reminded of the reasons for *believing* that Hesperus is Phosphorus, it is hard to know what would count as a satisfying answer...and this also applies to questions like 'Why is it that to be a vixen is to be a female fox?' Once we set aside the "remind me of

identity may hold the key to the epistemology of analysis. Such an account would provide conditions in which the inference to 'To be F is to be G' is warranted.

I do not place inordinate weight on any one of these proposals. What matters is that identity plays some philosophically important role or other, and so its epistemology is significant. As we shall see, natural attempts to fold such an epistemology into existing theories of evidence and justification face insurmountable obstacles. The quest for an epistemology of identity is far more daunting than it initially appears.

2 Cause for Pessimism and Skepticism

This topic is new enough that it is worth briefly discussing criteria for success. What ought the epistemology of identity accomplish? Under what conditions can an account be dismissed? The constraints I have in mind are hopefully uncontroversial, but nevertheless eliminate numerous possibilities.

The first constraint is that an account ought to be *extensionally adequate*. All and only the right things ought to count as evidence for identity claims. This may seem obvious or even trivial. And because we currently lack an epistemology of identity, we may occasionally be unsure whether an account is extensionally adequate after all. Nevertheless, some cases are clear. While an observation that the Patriots won the Superbowl is no evidence that Hesperus is Phosphorus (at least in normal circumstances), appropriate astronomical observations are. Proposals which admit the wrong kinds of things as evidence and those which exclude legitimate evidence can both be dismissed on the grounds that they are extensionally inadequate.

The second constraint is that an account ought to be *suitably general*. I am relatively unconcerned with particular identical claims. A theory which provides conditions for something to constitute evidence for the claim that Cicero is Tully but applies to no other cases is relatively uninteresting—even if it is extensionally adequate. Admittedly, the notion of suitable generality is somewhat vague. A single application is surely not enough, but is two? How many cases need an account apply to in order for it to count as suitably general? Of course, an account with universal applications would be ideal, but perhaps we will find ourselves in a position to accept some limitations should one not arise.

This constraint imposes restrictions upon the epistemology of identity from the outset. In particular, an account ought to apply both to *a priori* and *a posteriori* cases. While the claim that the number two is the only even prime is knowable through introspection alone, the claim that water is the chemical compound H_2O requires empirical support. Suitably general accounts apply to both kinds of cases, and so take no stand on whether evidence for identity is attainable *a priori* or *a posteriori*.

The third constraint is that an account ought to be *psychologically accessible*. Ordinary people regularly make inferences to identity claims. Although some of these inferences are

reasons to believe" reading...it is hard to see what an answer would even look like."(Dorr, 2016, pg. 41)

made in error, others are felicitous; at least sometimes, we respond reasonably to available evidence. Accounts which place all evidence of identity beyond our reach fail to respect the role identity plays in our epistemic lives.

This requirement is not intended to be overly psychologistic. I leave open the possibility that people regularly—perhaps even systematically—make mistakes in this area. And I allow for accounts which humans approximate, but never perfectly achieve. Nevertheless, I also assume that we have some access to the relevant evidence. An account according to which all evidence of identity is in principle beyond our reach is objectionably psychologically inaccessible.

These three constraints (extensional adequacy, suitable generality and psychological accessibility) provide a sufficient framework to begin. What follows is an overview of candidate epistemologies of identity. None are successful, raising a worry that there is no unified epistemology to be found. Nevertheless, the ways in which these suggestions fail are illuminating, and indicate what an adequate epistemology of identity ought to do.

2.1 Epistemic Possibility and Our Location in Modal Space

One suggestion is boring, unoriginal, and initially plausible—the epistemology of identity can be straightforwardly folded into existing theories of evidence and justification. There are several reasons to suspect this is so. First, it promises a unified conception of evidence. We need not adopt a separate theory for each type of proposition; rather a general account determines what constitutes evidence for any claim. Second, this explains why it is no one has bothered to construct an epistemology of identity. There is no special problem of identity, the thought goes, we simply need to 'plug in' identity claims to existing theories.

A benefit of this suggestion is that it is easy to test. Accounts of evidence and justification are readily available, so it is straightforward to select from the bevy of options. I do not canvass every theory of evidence and justification here. This is partially due to restrictions on space, and partially because some accounts determine little (if anything) about what the epistemology of identity consists of. According to reliabilist theories of justification, for example, a belief is justified just in case the process by which it is formed is reliable—i.e., leads to true beliefs sufficiently frequently.[12] Absent an understanding of what it is that leads to the belief that one thing is identical to another, this says nothing in particular about the epistemology of identity. So it is difficult to determine what, precisely, justifies the inference to identity from this account alone.

Fortunately, other views are easier to evaluate. One such view relies upon the notion of *epistemic possibility*—a kind of modality relating to how the world might be for all that one knows. Epistemic modality is standardly contrasted with subjunctive modality—a kind not indexed to agents' epistemic states. A prominent conception of epistemic modality concerns our ability to locate oneself within modal space.[13] Roughly, the picture is this:

[12] For defenses of this account of justification, see, e.g., Goldman (1979, 1986).

[13] There are other prominent accounts of epistemic modality. For example, (Hacking, 1967, pg. 153)

there are a vast number of possible world. In some, Donald Trump is President, while in others Hillary Clinton is President; in some the United States never entered World War II, while in others it is currently engaged in World War III; in some intelligent life never evolved, while in others it is widespread throughout the universe. Presumably, no one has uniquely identified which possible world is actual, at least in a non-indexical way. While someone may truly say 'This world is actual,' she does not know whether or not she lives in a world with an even number of blades of grass. Nevertheless, some restrictions are at hand. I know that I am not in a world entirely devoid of intelligent life, for example, on the grounds that I exist.

Some maintain that evidence serves to eliminate candidate possible worlds. The more evidence I have, the narrower the scope of possible worlds that I might occupy. Thus, for example, Hawthorne (2004) states "It is possible that p for S at t (there is a chance that p for S at t) iff p is consistent with what S knows at t" (pg. 26. See also Hawthorne (2012)), and Stanley (2005) claims "It is possible$_{\text{A}}$ that p is true if what A knows does not, in a manner that is obvious to A, entail not-p" (pg. 128).[14]

Along these lines, we might say that e is *conclusive evidence* that p just in case p is epistemically necessary for all of those in possession of e (i.e., just in case for every agent in possession of e, all epistemically possible worlds are worlds in which p is true). For example, my phenomenal experience is conclusive evidence for the existence of intelligent life, because there is no possible world in which someone possesses my phenomenal experience and intelligent life does not exist. Inconclusive evidence might be defined in various ways: we might say that e is inconclusive evidence that p just in case there are more possible worlds in which both e and p are true than worlds in which e is true and p is false, or perhaps that all (or most) of the relevantly close worlds in which e is true are also worlds in which p is also true, or that the counterfactuals 'If e were true then p would be true' and 'If e were false then p would be false' hold. A newspaper report that Edmund Hillary climbed Mount Everest is inconclusive evidence that Hillary climbed Everest just in case the number of possible worlds in which the newspaper report exists and Hillary climbed Everest than possible worlds in which the report does not exist and Hillary did not climb Everest (or one of the other candidate theories of inconclusive evidence).

On this view, something constitutes evidence for an identity claim just in case it narrows which possible worlds we might occupy in the appropriate way. e is conclusive evidence that $a = b$ just in case it epistemically necessary for all those in possession of e that $a = b$—i.e., all epistemically possible worlds are worlds in which $a = b$ (and similarly so for inconclusive evidence). So, in which possible worlds does $a = b$?

Well, *all of them* or *none of them*, depending on whether '$a = b$' is true. Since at least the 1970's the received view has been that identity holds necessarily. If Hesperus is

defends the view that "'it is possible that p' means that p is not easily known to be false, nor would practicable investigations establish that it is false." For related views, see, e.g., Teller (1972); DeRose (1991); MacFarlane (2011).

[14] For a more formal development of this sort of view, see, e.g., Yalcin (2007).

identical to Phosphorus in the actual world, then Hesperus is identical Phosphorus in every possible world; there is no possible way for them to be distinct. The necessity of identity receives support from multiple fronts; Marcus (1947) first offered a proof from modest modal assumptions, and Kripke (1980) provided additional linguistic support. If this is correct, then the set of worlds in which $a = b$ is either the set of all possible worlds (if it is true) or else the null set (if it is false). On this view, *everything* constitutes evidence of (true) identity claims, as it is epistemically necessary for everyone that the identity holds, regardless of what evidence is at hand. This proposal thus fails the test of extensional adequacy, for it inaccurately entails that everything whatsoever is conclusive evidence of identity.

Perhaps some are tempted to appeal to metalinguistic ignorance in this kind of case.[15] Although every possible world is one in which Hesperus is Phosphorus, not every world is one in which 'Hesperus' and 'Phosphorus' co-refer. Perhaps anything that narrows the possible worlds to ones in which 'Hesperus' and 'Phosphorus' denote one and the same object counts as evidence that Hesperus is Phosphorus.

This suggestion remains extensionally inadequate. Consider, for example, a person who has no knowledge of what either 'Hesperus' or 'Phosphorus' refer to—they could denote planets, people, countries, etc.. Such a person might have it on good authority that 'Hesperus' and 'Phosphorus' co-refer—i.e., they might hear expert testimony that the two names denote the same object, whatever that object is. In this case, the person has evidence that 'Hesperus' refers to the same object as 'Phosphorus,' but she lacks evidence that Hesperus is Phosphorus.

So too, it seems, someone might have evidence that Hesperus is Phosphorus while lacking evidence that 'Hesperus' refers to the same object as 'Phosphorus.' Perhaps an astronomer, ignorant of the Ancient Greek names and out-of-touch with contemporary philosophy, has no idea whether or not 'Hesperus' and 'Phosphorus' co-denote. Such an astronomer could, with the appropriate astronomical evidence and theories, come to know that Hesperus is Phosphorus, without ever hearing either name. This astronomer has evidence that Hesperus is Phosphorus, but lacks evidence that 'Hesperus' denotes the same object as 'Phosphorus.' So it is possible to have evidence that 'Hesperus' co-refers with 'Phosphorus' without having evidence that Hesperus is Phosphorus, and it is possible to have evidence that Hesperus is Phosphorus without having evidence that 'Hesperus' co-refers with 'Phosphorus.'

2.2 E=K

Perhaps this trouble is peculiar to this particular account of evidence. Perhaps other theories are better equipped to accommodate identity. And, perhaps, the ability to incorporate the epistemology of identity is a mark in favor of alternate views.

[15] See, e.g., Frege (1879); Stalnaker (1976) for this sort of appeal.

To that end, I turn to Williamson (2000)'s recent proposal that evidence is knowledge. He defends this view while outlining important theoretical roles knowledge can play even if it is irreducible to anything more fundamental. In arguing that E=K, Williamson claims that a person's total evidence is identical to what they know. His argument is the following:

1. All evidence is propositional.
2. All propositional evidence is knowledge.
3. All knowledge is evidence.
4. Therefore, all and only evidence is knowledge.

This argument is incontrovertibly valid, but each of the premises is contentious. Many prosecutors, for example, would enter a bloody knife as evidence in a murder trial, but a bloody knife is not a proposition.

Williamson defends these premises in multiple ways. In support of premise 1, for example, he notes that sentences of the form 'A because B' are ungrammatical if the terms taking the place of A and B were not assertoric expressions; for example, the statement 'Albania because ... ' is ungrammatical. If evidence is to be understood as that which underlies explanatory assertions, it may be that only truth-evaluable entities count as evidence. In addition, particular evidence is sometimes said to be inconsistent with one hypothesis or another, and it is unclear how something non-propositional is capable of being either consistent or inconsistent with anything at all.[16]

The majority of Williamson's discussion concerns agents' total bodies of evidence. An exhaustive theory of evidence surely requires more. In addition to accounting for the total evidence someone has available, it ought, minimally, to determine when evidence lends support to one claim rather than another. According to Williamson, my knowledge that whales are mammals counts among my evidence, but which propositions does it support in particular?

To this end, Williamson provides the following account:

> EV: e is evidence that h for S if and only if S's evidence includes e and $\mathrm{P}(h|e) > \mathrm{P}(h)$.

Whether S's evidence includes that e is determined by S's knowledge, and whether or not it is evidence for a particular proposition h is determined by whether the probability of h given e is greater than the probability of h without e.

This proposal faces the opposite problem plaguing epistemic possibility accounts. While they entailed that everything is evidence of identity, Williamson's proposal entails that nothing is. On the standard view, identity claims are either necessarily true or necessarily

[16]The claim that E=K has come under attack on multiple fronts. For example, Brueckner (2005) argues that Williamson cannot accommodate perceptual knowledge; I may know that a cup is red based on my perception, but my perception is not a proposition, and Joyce (2004) argues that it cannot accommodate the fact that the evidential status of many propositions is a matter of degree.

false. So the probability of any identity claim is either 1 or 0. According to standard probability theory, it is impossible for any evidence to raise the probability of propositions whose initial probabilities are either. Nothing raises the probability of claims with a probability of 1; that it the highest probability anything can possibly have. And nothing impossible becomes more likely on the basis of new evidence; that which is impossible remains impossible regardless of what evidence is obtained.

Evidence is defined in terms of knowledge which raises probability, and nothing raises the probability of identity claims, so nothing is evidence of identity. This account thus fails the test of extensional adequacy, for I assume that there is evidence of identity or other. Accounts that entail that there is no evidence of identity (like the proposal that E=K) misdiagnose the relevant cases.

Neither sections 2.1 nor 2.2 are intended to be knockdown assaults on the respective theories of evidence. Each may elucidate the nature of evidence for other kinds of propositions. However, they are unable to explicate evidence of propositions that are either necessarily true or necessarily false. Because these include identity claims, they cannot be account for the epistemology of identity.

2.3 Direct Perception

Perhaps, instead, we have unmediated access to identity.[17] That is to say, we directly perceive an object's identity just as we directly perceive its color, shape and size. The epistemology of identity is of a kind with the epistemology of perception; I am justified in believing that a thing is identical to itself because I perceive that it is so.

I am skeptical of this suggestion. Through sight, I perceive color; through hearing, I perceive sound; through touch, I perceive size and shape. But I know of no perceptual mechanism that allows me to perceive identity. Rather, it is something I conclude—perhaps often unconsciously—on the basis of what I do perceive. However, my initial skepticism is no argument, and this proposal warrants a substantive response.

Suppose that I know identical twins, Bert and Ernie, who are indistinguishable to me; my perception of one is the same as my perception of the other. Of course, Bert and Ernie can easily distinguishable themselves 'from the inside'—Bert has no confusion over whether or not he is Ernie. But I lack this internal access, so my perceptions of them are the same.

Suppose that I observe Bert enter a room—or, rather, suppose that I observe someone phenomenally indistinguishable from Bert enter the room who unbeknownst to me is Bert instead of Ernie. The person looks and sounds precisely as Bert looks and sounds. Of course, he also looks and sounds exactly as Ernie looks and sounds—after all, they look and sound precisely the same. What facts about identity do I perceive?

Four possibilities could obtain:

1. I perceive that the person is identical to Bert and not to Ernie.

[17] My thanks to Ralph Wedgwood for this suggestion.

2. I perceive that the person is identical to Ernie and not to Bert.
3. I perceive that the person is identical both to Bert and to Ernie.
4. I neither perceive that the person is identical to Bert nor to Ernie.

The first two possibilities can be straightforwardly dismissed. Bert and Ernie are indistinguishable to me, so my perceptions of one of them entering the room is identical to my perception of the other entering the room. I have no special sense that allows me to perceive that the person is identical to Bert and not Ernie (nor do I have a faulty mechanism causing me to perceive that the person is Ernie and not Bert). There could be no reason for me to observe that the person entering is identical to one rather than the other.

The third possibility has one advantage. Minimally, it does not entail that my perceptions of identity are arbitrary. My perceptions when either Bert or Ernie enter the room are precisely the same; in either case, I perceive that the person is identical to Bert and that the person is identical to Ernie. But the advantages of the third possibility end there—it requires that my perception is *incoherent*. After all, I know that Bert is not identical to Ernie, and if I perceive that a person entering a room is identical to them both, then what I perceive is incompatible with what I know to be true. But, surely, my observation of one of the two entering the room is not incoherent; it does not conflict with the fact that I know there are two distinct people.

What remains is option 4. I neither perceive that the person is identical to Bert nor to Ernie. Because these are the only two people I could perceive that this person is identical to, I do not perceive identity. Of course, this case is somewhat idiosyncratic—not everyone has an identical twin. But, presumably, my perception when Bert enters the room is in no way dependent upon whether or not another phenomenally indistinguishable person exists. And so I do not directly perceive facts about identity, regardless of whether the object of my perception has a phenomenally identical twin.

2.4 The Logic of Identity and Inference from Identity Claims

The proposals discussed so far were fruitless in a general way. They do not provide a plausible account of evidence for any identity claims, and the obstacles they encounter are, it seems to me, insurmountable. Other natural suggestions are moderately successful; there are sources of evidence for identity claims that are beyond dispute. Despite their relative success, I also interpret these proposals as grounds for pessimism. They apply in only restricted cases, rather than constituting a general account. This raises a worry that there is no unified epistemology of identity to be found; it may ultimately be an ugly, motley and disjointed affair.

One source of evidence arises from the meaning of the identity sign. All those who know what '=' means know that every object is identical to itself.[18] Knowing, as we do,

[18] For an argument that it lies in the essence of identity for everything to be identical to itself, see Dasgupta (2015).

that all objects are identical to themselves, we rightfully infer that Socrates is identical to Socrates and that water is identical to water. What justifies my belief that a particular object is self-identical is my knowledge that all objects are self-identical, and the ensuing inference concerning that particular object.

I do not mean to suggest that this process never occurs, nor that beliefs formed in this way do not constitute knowledge. However, this method applies to a small subset of cases. When a victim looks at a police lineup and says 'He is the mugger,' they presumably do not arrive at that conclusion by instantiating the claim that everything is self-identical. Outside of the context of a logic classroom, it is unclear how often, if at all, reasoners employ that sort of inference.

Perhaps one resource can be supplemented with others; after all, we know of other logical attributes of identity. Minimally, identity is symmetric and transitive. We can appeal to these features to generate knowledge of additional identity claims. For example, if I know that Superman is the unique alien on Earth, and I know that Superman is Clark Kent, I can appeal to the transitivity of identity to conclude that Clark Kent is the unique alien on Earth.

This appeal to the logic of identity is, arguably, an instance of a more general inference-type: the inference from one identity claim to another. While logic provides one resource, others are also available. A subtler example concerns identity claims containing a definite description, i.e., ones that take the form 'a is the F' (or 'The F is a'). Natural examples include 'Barack Obama is the 44^{th} president' and 'The number 2 is the even prime.' Quite plausibly, someone is justified in concluding that 'a is the F' if they know that $F(a)$ and that there is a unique F. So it is epistemically permissible to infer that Obama is the 44^{th} president on the basis of knowing that he is a 44^{th} president and that there is a unique 44^{th} president. Similarly, if someone knows that some object or other is both F and G and knows that there is both a unique F and a unique G, they are presumably justified in concluding that the F is the G. This type of inference has limited applications; it does not account for identity claims without definite descriptions.[19]

One part of the justification (i.e., the claim that a bears the property F) need not involve identity—in order to conclude that Barack Obama is a 44th president, I need to invoke knowledge of history and election outcomes, rather than identity claims. However, the other part, which concerns uniqueness, does appeal to identity. The claim that there is a unique F amounts to the claim that all objects that are F are identical to one another. The claim that there is a unique 44^{th} president amounts to the claim that all 44^{th} presidents are identical. So, in concluding that a is the F (partially) on the basis of knowing that there is a unique F, one thereby infers one identity claim from another.

I do not dispute that the inferences discussed are extensionally adequate. Those who infer that $a = c$ on the basis of knowing that $a = b$ and $b = c$ make no mistake—nor do

[19] I make the hopefully uncontroversial assumption that proper names are not disguised definite descriptions.

those who conclude that a =the F on the basis of knowing that Fa and that there is a unique F. Of course, these techniques are of no help to someone who is ignorant of any identity claims at all; at best, they expand some knowledge of identity into more. They do not explain how it is we come to know substantive identity facts to begin with—knowledge of trivial identity facts (i.e., ones that take the form $a = a$) are not enough. There are many inferences to identity that these suggestions do not apply to—those that neither take the form '$a = a$' nor are inferences from other identity claims. So, these proposals are not suitably general, and so cannot provide a foundation for the epistemology of identity.

2.5 The Satisfaction of Identity Conditions

I know that {Cicero} is identical to {Tully}. My justification relies on two claims; the fact that Cicero is Tully and the fact that sets with identical members are themselves identical. How I know that Cicero is Tully remains somewhat mysterious, but my second piece of justification comes from my understanding the identity conditions of sets. Of course, this is not peculiar to {Cicero} and {Tully}. One can always appeal to the identity conditions of sets in order to know that two sets are identical. Perhaps this can serve as a model for other cases. Perhaps we come to know that two objects are identical by realizing that their identity conditions are satisfied.

This proposal evades concerns that others faced. It is undoubtedly extensionally adequate: any time one knows that identity conditions are satisfied, no error will occur in inferring identity. And, plausibly, it has a wide scope. If we assume that everything has some identity conditions or other, this proposal provides conditions for coming to know any identity claim whatsoever.

Nevertheless, this proposal is psychologically inaccessible—at least most of the time. Set theory is something of an outlier. We usually have no idea what the identity conditions of a thing are. The conditions for personal identity has received more attention than many other cases, but no one knows what its true identity conditions are. And some particular proposals do not provide much of a guide. Take, for example, Locke (1690)'s claim that two people are identical just in case one person can remember acting as the other did. We standardly lack evidence about what people around us remember. So, even if we knew that personal identity is keyed to memory, we would be unable to appeal to satisfying the identity conditions of persons in order to come to know that people are identical.

In other areas we are even more ignorant; we presumably have no idea what identity conditions for tables, planets or chairs are. It is safe to say that cases in which we come to infer identity because they satisfy identity conditions are extremely rare. Perhaps other, more knowledgeable creatures can appeal to this approach. Someone who knew the identity conditions of everything whatsoever might appeal to them when inferring identity claims. But it is beyond the reach of creatures like us.

Five proposals have been canvassed that ranged from the completely hopeless to the moderately successful. In light of the persistent failings, one might suspect that there is no

satisfactory epistemology of identity to be found. However, I hope to convince the reader that this pessimism is premature.

3 A Hybrid View

If psychological accessibility is a legitimate theoretical criterion, perhaps we ought to take our cue from inferences that we actually make. Take a perfectly banal example: last week a friend of mine visited from out of town, and this morning I drove him to the airport. Throughout this process, I inferred numerous identity claims. I believed that the person I was driving to the airport was identical to my friend, for example, and that the car I drove was identical to my car. These beliefs were implicit at the time, but it mattered very much that they were true. Had I driven someone else to the airport, my friend would have missed his flight; had I driven some other car, my chances of being arrested for grand-theft auto would have been greatly increased.

What evidence did I rely upon in forming these beliefs? Presumably, there was a multitude of factors. The person I drove looked and sounded very much like I remember my friend looking and sounding; his conversations with me indicated that he recalled our past shared activities; he exited the spare bedroom that I saw my friend enter the night before. My car looked qualitatively identical to mine, down to the precise license-plate; it was in the parking spot where I left mine; the seat was positioned correctly for someone of my height; my keys opened the door and started the engine. Some of this evidence could be construed in terms of other identity claims. For example, I believed that the bedroom I saw my friend exit was identical to the one I saw him enter the night before, and that the keys I found next to my bed were identical to my keys. Many of these inferences concerned identity over time. However, I do not want to place undue weight on temporal considerations—some of my identity beliefs (such as the belief that his departure terminal is identical to terminal 2) had nothing to do with time at all.

Similarities emerge from my various pieces of evidence. Much of it involves the preservation of properties.[20] Had the person I observed exiting the spare bedroom looked or sounded differently than my friend looks and sounds, I would not conclude that they were identical. I do not confuse differently colored cars for mine. With an expansive enough conception of properties, it may even be that I relied upon the fact that both my car and the car I drove bore the property *having been located at the same parking spot.* Perhaps all evidence of identity is similar—it concerns the preservation of properties.

Care must be taken when articulating this conception of evidence. It cannot be that something constitutes evidence that $a = b$ just in case it is evidence that a and b share some property or other. Take, for example, a pair of boarding tickets indicating that Sarah

[20] I do not want to weigh into the debate between realism and nominalism here. My use of 'properties' is heuristic—I suspect nominalists could accept my proposal by replacing 'properties' with whatever it is they take to play the theoretical roles that realists claim properties play.

flew to Budapest and Cara flew to Vienna. The tickets are plausibly evidence that Sarah and Cara have some properties in common—for example that they both flew somewhere in central Europe. However, they are not evidence that Sarah is identical to Cara. A more plausible suggestion is the following:

> **Epistemology through the Identity of Indiscernibles (EII):**
> Something counts as evidence that $a = b$ just in case it counts as evidence that a and b bear all of the same properties.

This proposal is deliberately vague in some respects. In stating that 'something' counts as evidence, I aim to avoid debates over what things count as evidence in general. If it turns out that all evidence is propositional, then evidence for identity claims must itself be propositional. If perceptions and objects can count as evidence, then they too can count as evidence for identity claims.

Presumably, this view is coupled with an epistemology of distinctness. Something counts as evidence that $a \neq b$ just in case it counts as evidence either that a bears some property that b does not, or that b bears some property that a does not. For example, the fact that a car is the same model as mine might count as evidence that the two are the same, while the fact that a car is a different model from mine is evidence that they are not.

According to EII, the epistemology of identity is tightly connected to the identity of indiscernibles—the metaphysical claim that two objects are identical just in case they share all of the same properties. This may well be seen as a mark in its favor; many suspect that the identity of indiscernibles is central to our conception of identity. On this view, it does not simply underlie the metaphysics of identity, but also underlies the epistemology of identity. The way that we come to know two things are identical is by realizing that they bear all of the same properties.

However, problems that plague the metaphysics of the Identity of Indiscernibles also challenge its epistemic application. Consider, for example, Black (1952)'s example of a world consisting only of two qualitatively identical spheres.[21] By stipulation, we can be assured that the spheres share every property with one another. Yet knowing that the spheres share every property does not allow us to know that they are identical. Those who countenance Black's sort of example ought to find EII extensionally inadequate; it inaccurately suggests that evidence that the spheres share every property is evidence that they are identical.

[21] There have been numerous discussions of Black's example over the years. For example, Hacking (1975) notes that the same scenario could be described as a single sphere in non-Euclidean space, and Della Rocca (2005) objects that abandoning the identity of indiscernibles on the basis of Black's example requires adopting a primitive or brute conception of identity. In defense of Black's scenario, Adams (1976) notes that everyone countenances scenarios in which two objects are nearly indiscernible—e.g., a world consisting of two spheres that are identical except for the existence of a tiny scratch on one. It is difficult to see why the presence of a scratch could render the scenario intelligible if the absence of a scratch is not.

I do not mean to accept Black's conclusion prematurely. While it seems to me that the case he describes is perfectly coherent and poses problems for EII, others might reasonably disagree—perhaps qualitatively identical objects are numerically identical after all. For the moment, I assume that Black's world undermines EII, but I should emphasize that this assumption is temporary. The view I ultimately defend is one on which this assumption can be discharged.

Perhaps problems arose because Black's spheres are located in different places. If so, accounts in terms of spatiotemporal overlap may fare better.

> **Epistemology through Spatiotemporal Overlap (ESO):**
> Something counts as evidence that $a = b$ just in case it counts as evidence that a and b are spatiotemporally coincident.

ESO easily accommodates Black's world. Because the spheres are spatiotemporally distinct by stipulation, we can be sure that they are not identical. And other cases can also be handled in a straightforward manner. For example, perhaps my observation that a car is in the parking spot where I left mine is evidence that it is my car because it is evidence that the two spatiotemporally overlap.

However, other cases are harder to diagnose. Consider that classic example of the statue and the clay. Many hold that they are coincident—the two objects occupy the very same location in space.[22] Nevertheless, we can be sure that the objects are distinct because they bear different properties. Minimally, their dispositional properties differ. A fragile statue is disposed to break when it is dropped, but a lump of clay would remain. Perhaps this disposition can be cashed out in terms of essence. While the statue is essentially shaped thus-and-so, the clay is only accidentally shaped thus-and-so. So while an alteration of the shape would destroy the statue, the clay would remain.

Some discussions of this example also allow the statue and clay to have varied in their past properties; the clay existed for a period of time before it constituted a statue. In this case, the two do not completely overlap throughout spacetime, they merely overlap through part of it. But we can consider versions of this case where both the clay and statue come into being and are destroyed simultaneously, and so vary in their dispositional properties but are always entirely coincident. Many philosophers thus contend that it is possible for two objects to occupy the same points throughout all of space-time. But why stop at two? How many objects are co-located with the statue?

The statue of David is currently located in Florence. It could be moved elsewhere—it might be displayed in an exhibition in Rome. Neither the statue nor the marble that compose it would be destroyed by this move; both only accidentally bear the property

[22] Those who defend objectual coincidence (in some form or other) include, but are not limited to Baker (1997, 2000); Fine (2003); Forbes (1987); Johnston (1992); Koslicki (2004); Kripke (1971); Lowe (1995); Oderberg (1996); Shoemaker (2003); Thomson (1983, 1998); Yablo (1987). For a defense of a conception of identity in terms of material constitution see e.g., Noonan (1993).

being located in Florance. But perhaps there is some other object which *is* destroyed. Perhaps something is coincident with David that is only accidentally composed of marble and accidentally shaped thus-and-so, but is essentially located in Florence. This sort of object is destroyed when it is transported to another exhibition, but would not be destroyed if David were smashed. Of course, people rarely pay attention to this sort of object; it does not impinge upon our lives. But some philosophers suspect that this reflects facts about us, rather than the world. Unlimited essentialists maintain that for anything that bears some collection of properties, there exists an object that bears each subset of those properties essentially and the rest accidentally.[23] The primary motivation for unlimited essentialism is that it is difficult to construct a principled reason for some properties should be essential to an object, while other properties are merely inessential to objects that bear them. Of course, we need not be unlimited essentialists in order to maintain that the statue is distinct from the clay; perhaps there are only two objects at that location.

If there are two (or more) coincident objects, challenges arise for ESO. Suppose someone were to tell me "The statue of David would be destroyed if that piece of marble were smashed." This testimony provides evidence that David is co-located with that particular piece of marble. David would survive if some other bit of marble were smashed instead. However the testimony is not evidence that David is identical to the marble—I may very well know that they remain distinct, yet coincident, objects. And so, those who countenance colocation would maintain that ESO fails the test of extensional adequacy. It maintains that I have evidence that the statue of David is identical to the marble that composes it, while I maintain that I do not.

As with EII, I do not want to prematurely weigh in on a contentious philosophical debate. After all, some philosophers deny that colocation is coherent.[24] Such philosophers may well take testimony that David is co-located with a lump of marble is evidence that David is identical to the lump. For the moment, I temporarily assume that such philosophers are mistaken. However, I will shortly demonstrate how this assumption can be discharged.

ESO could not accommodate cases in which objects shared every property and yet were located in different places. EII could not accommodate cases in which objects were collocated and yet did not share every property. The obvious solution to these problems is to take these accounts together.

A Hybrid View
Something constitutes evidence that $a = b$ just in case it is evidence that a and b bear all of the same properties and are spatiotemporally coincident.

[23] For defenses of unlimited essentialism, see Fine (1999); Johnston (2006); Koslicki (2008). For further discussion of unlimited essentialism, see Dasgupta (2018).

[24] Those who deny this possibility do so for a variety of reasons. For example, mereological nihilists like Unger (1979) and some mereological moderates like Van Inwagen (1990) deny that a lump of clay composes an object, and so deny that there is a lump of clay for a statue to be identical to.

This proposal easily accommodates the cases discussed so far. Those who know by stipulation that Black's example consists of two qualitatively identical spheres lack evidence that the spheres bear all of the same properties and are spatiotemporally coincidence, so the stipulation is not evidence that the two spheres are identical to one another. And those who have evidence that a statue is identical to the clay that compose it lack evidence that the two bear all of the same properties.

This proposal is available to those who deny that qualitatively identical objects are ever be distinct. After all, philosophers who deny Max Black scenarios presumably maintain that qualitatively identical spheres are numerically identical, so evidence that two spheres are qualitatively identical is also evidence that they are spatiotemporally coincident. Similarly, this proposal is available to those who deny that there are distinct yet spatiotemporally coincident objects. They maintain that statues are identical to the material that compose them, and so (given Leibniz's Law) bear all the same properties. Evidence that two objects are coincident is thus also evidence that they bear all of the same properties. And so this hybrid view is available to philosophers regardless of what stand they take on the Black's imagined world and the relationships between the statue and the clay.

It seems to me that this proposal is extensionally adequate. Minimally, it fares better than some of the suggestions previously considered that entailed either that everything is evidence of identity or that nothing is. I myself cannot think of cases it misdiagnose, and await patiently for a philosopher who can.

I also maintain that this satisfies the criterion of suitable generality. It applies to many different kinds of things; to cars, houses, people, planets and many more besides. One potential complication concerns the identity of abstract objects. How could this condition apply to the number 2 or the property of *being a bachelor*? To the extent that I can, I remain agnostic about the existence of abstract objects. However, if there are any such things, I maintain that the condition of spatiotemporal coincidence is vacuously satisfied. Because these objects are not physically, all of them vacuously occupy the same region of spacetime—none at all. Evidence that two abstract objects are identical is evidence that they bear all of the same properties.

While some philosophers countenance Max Black type situations for physical objects, I know of none who do for abstract objects. No one maintains that there is a multiplicity of entities all of which bear the same properties as the number two, or that there are numerically distinct properties indistinguishable from *being red*.[25] And if there are no

[25]Quidditists come close. They maintain that properties are not defined in terms of their causal roles (see, e.g., Armstrong (1989); Lewis (2009); Bird (2007); Schaffer (2005)). It is a contingent matter of fact that the property which actually playing the mass-role plays the role that it does, and contingent that the property which actually plays the charge-role plays the role that it does. In principle, the two could 'switch roles' such that the property that had played the mass role could come to play the charge role and vice versa. But while quidditists maintain that properties have all of the same essential characteristics, they do not maintain that they have all the same accidental characteristics. While the quiddity playing the mass role is essentially indistinguishable from the quiddity playing the charge role, it alone bears the property *plays the mass role* (For a discussion of quidditism along these lines, see Dasgupta (2015)).

numerically distinct yet indistinguishable abstract objects, then evidence that two abstract objects are indistinguishable is also evidence that they are identical. It has recently been suggested that properties are identical just in case they are higher-order indistinguishable (see Dorr (2016)). I take no stand on this claim as a metaphysical suggestion. However, if I am right, it underlies the epistemology of the identity of properties.

More worrying is the criterion of psychological accessibility. How often do we ever have evidence that two objects bear all of the same properties? We don't have time to check every one. And once we consider silly enough properties, we always possess evidence that they have *some* property or other in common. Both a chair and a car, for example, bear the property *being either a chair or a car*, but surely this does not lend support to the claim that the two are identical.

This criterion, I believe, is one that humans can approximate but never perfectly achieve. We are never in a position to check every property that an object bears—we lack the time and the ability to do so. However, we can (and do) continue to verify that objects share a range of properties. The more we verify, the higher the likelihood that the objects are identical. Of course, we can always do a bad job—it is perfectly possible for us to check 'silly' properties that all (or nearly all) objects have in common. We do not gain much evidence from verifying that two objects are both not entirely composed of hydrogen, or entirely composed of helium, etc.. However, we are readily capable of distinguishing the silly properties from substantive ones. That two objects share a property nearly everything does is slim evidence that they are identical; that two objects share an uncommon property is weightier. Although are rarely in a position to verify either that two objects have every property in common or that they occupy the same region of spacetime, we can make progress; continuing to verify more than we had before.

The hybrid view thus succeeds where other accounts have failed; it satisfies the three minimal conditions mentioned at the outset. These conditions are minimal; there may be further virtues that this proposal lacks. Nevertheless, it fares better than contenders that were unable to do so.

4 Conclusion

In spite of the advantages the hybrid view has over competitors, intellectual integrity requires being forthright about its shortcomings. There is a potential problem that I have been unable to resolve.

On one conception, the hybrid view is a covert method of inferring one identity claim from others, rather than a tool allowing philosophers to infer identity claims from evidence that do not concern identity. Take overlap. We might say that two objects completely overlap in spacetime just in case they occupy *the same* region of time and space; i.e., the locations that they occupy are identical. And so an aspect of the hybrid account—that concerning overlap—itself draws upon identity claims.

So too, the identity of indiscernibles might be seen to subtlely appeal to identity claims. In a second-order language, I might state 'a and b share all of the same properties; as '$\forall F(Fa \leftrightarrow Fb)$.' The identity sign does not appear in this sentence. However, as mentioned in the introduction, some philosophers maintain that variables that reoccur within the scope of a quantifier covertly invoke identity. In inferring that a and b bear the same property, I conclude that the property borne by a is identical to the property borne by b. Of course, this is controversial; not all philosophers believe quantification covertly invokes identity. Those who do, however, may suspect that the hybrid view appeals to identity in that evidence both relies upon evidence of spatiotemporal overlap and relies upon evidence of indiscernibility.

It may be, then, that the hybrid view provides conditions to infer one identity from others, without providing conditions to infer an identity claim from no identity claims at all. And—as discussed in section 2.4—it may be that we thus lack a foundation for the epistemology of identity. Conditions for inferring some identity claims from others are well and good, but it is difficult to see how such a view could 'get off the ground' without being supplemented by conditions for evidence of some identity claims that do not appeal to others.

I see no solution to this worry, except to deny that either condition covertly invokes identity. I ask the reader not to prematurely decide against the hybrid view without attempting to construct a more viable alternative. I suspect that this attempt will, at minimum, demonstrate that whatever the correct theory is, it is not as simple as we might initially expect.

References

Adams, Robert Merrihew. 1976. "Primitive Thisness and Primitive Identity." *Journal of Philosophy* 76:5–26.

Armstrong, David. 1989. *A Combinatorial Theory of Possibility*. Cambridge University Press.

Baker, Lynn. 1997. "Why Constitution is not Identity." *The Journal of Philosophy* 94:599–621.

Baker, Lynn. 2000. *Persons and Bodies: A Constitution View*. Cambridge University Press.

Bird, Alexander. 2007. *Nature's Metaphysics*. Claredon press.

Black, Max. 1952. "The Identity of Indiscernibles." *Mind* pp. 153–64.

Brueckner, Anthony. 2005. "Knowledge, Evidence and Skepticism According to Williamson." *Philosophy and Phenomenological Research* 70(2):436–43.

Burgess, Alexis. 2018. "The Things we do with Identity." *Mind* 127(505):105–28.

Carey, Susan. 2009. *The Origin of Concepts*. Oxford University Press.

Carey, Susan and Fei Xu. 1996. "'Infants' Metaphysics: The Case of Numerical Identity." *Cognitive Psychology* 30:111–53.

Correia, Fabrice. 2017. "Real Definitions." *Philosophical Issues* 27(1):52–73.

Dasgupta, Shamik. 2015. "Inexpressible Ignorance." *The Philosophical Review* 124(4):441–80.

Dasgupta, Shamik. 2018. "Essentialism and the Nonidentity Problem." *Philosophy and Phenomenological Research* 96(3):540–70.

Della Rocca, Michael. 2005. "Two Spheres, Twenty Spheres and the Identity of Indiscernibles." *Pacific Philosophical Quartertly* 86(4):480–92.

DeRose, Keith. 1991. "Epistemic Possibilities." *The Philosophical Review* 100(4):581–605.

Dorr, Cian. 2016. "To Be F Is To Be G." *Philosophical Perspectives* 30(1):39–134.

Elgin, Samuel. Forthcoming. "The Semantic Foundations of Philosophical Analysis.".

Fine, Kit. 1999. "Things and Their Parts." *Midwest Studies in Philosophy* 23:61–74.

Fine, Kit. 2003. "The Non-Identity of a Thing and its Matter." *Mind* 112:195–234.

Forbes, Graeme. 1987. "Is there a Problem of Persistence?" *Proceedings of the Aristotelian Society* 61:137–55.

Frege, Gottlob. 1879. *Begriffsschrift: Eine der Arithmetischen Nachgebildete Formelsprache des Reinen Denkens*. Halle.

Frege, Gottlob. 1892. "Sense and Reference." *Zeitschrift für Philosophie und Philosophische Kritik* 100:25–50.

Geach, Peter. 1967. "Identity." *The Review of Metaphysics* 21(1):3–12.

Goldman, Alvin. 1979. What is Justified Belief? In *Justification and Knowledge*, ed. George Pappas. Reidel Publishing Company pp. 1–23.

Goldman, Alvin. 1986. *Epistemology and Cognition*. Harvard University Press.

Hacking, Ian. 1967. "Possibility." *Philosophical Review* 76(2):143–68.

Hacking, Ian. 1975. "The Identity of Indiscernibles." *The Journal of Philosophy* 72(9):249–56.

Hawthorne, John. 2003. Identity. In *Oxford Handbook of Metaphysics*, ed. John Hawthorne. Oxford University Press.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.

Hawthorne, John. 2012. "Knowledge and Epistemic Necessity." *Philosophical Studies* 158:493–501.

Hofweber, Thomas. 2005. "Number Determiners, Numbers and Arithmetic." *The Philosophical Review* 114.

Hume, David. 1738. *Treatise of Human Nature*. Oxford Philosophical Texts.

Johnston, Mark. 1992. "Constitution is not Identity." *Mind* 101:89–105.

Johnston, Mark. 2006. "Hylomorphism." *The Journal of Philosophy* 103(12):652–98.

Joyce, Jim. 2004. Williamson on Evidence and Knowledge. In *Philosophical Books*. Blackwell pp. 296–305.

Kocurek, Alex. 2018. "Counteridenticals." *The Philosophical Review* 127(3):321–38.

Koslicki, Kathrin. 2004. "Constitution is not Identity." *Mind* 101:89–105.

Koslicki, Kathrin. 2008. *The Structure of Objects*. Oxford University Press.

Kripke, Saul. 1971. Identity and Necessity. In *Identity and Individuation*, ed. M. Munitz. New York University Press pp. 83–94.

Kripke, Saul A. 1980. *Naming and Necessity*. Harvard University Press.

Lewis, David. 1976. Survival and Identity. In *The Identity of Persons*, ed. A. Rorty. University of California Press.

Lewis, David. 1986. *On the Plurality of Worlds*. OUP Oxford.

Lewis, David. 2009. Ramseyan Humility. In *Conceptual Analysis and Philosophical Naturalism*, ed. Robert Nola and David Braddon-Mitchell. MIT Press pp. 203–22.

Linnebo, Øystein. 2014. "'Just is'-Statements as Generalized Identities." *Inquiry* 57(4):466–482.

Locke, John. 1690. *An Essay Concerning Human Understanding*. Penguin Books.

Lowe, Edward. 1995. "Coinciding Objections: In Defense of the 'Standard Account'." *Analysis* 55:171–8.

MacFarlane, John. 2011. Epistemic Modals are Assessment-Sensitive. In *Epistemic Modality*, ed. Andy Egan and B. Weatherson. Oxford University Press.

Marcus, Ruth Barcan. 1947. "Identity of Individuals in a Strict Functional Calculus of Second Order." *Journal of Symbolic Logic* pp. 12–15.

McGinn, Colin. 2000. *Logical Properties: Identity, Existence, Predication, Necessity, Truth*. Oxford University Press.

Noonan, Harold. 1993. "Constitution is Identity." *Mind* 102(405):133–46.

Oderberg, David. 1996. "Coincidence under a Sortal." *The Philosophical Review* 105:145–71.

Parfit, Derick. 1971. "Personal Identity." *The Philosophical Review* 80:3–27.

Quine, Willard V. 1964. "Review of P. T. Geach, *Reference and Generality*." *The Philosophical Review* 73:100–4.

Rayo, Agustín. 2013. *The Construction of Logical Space*. OUP Oxford.

Reid, Thomas. 1785. *Essays on the Intellectual Powers of Man*. Pennsylvania State University Press.

Schaffer, Jonathan. 2005. "Quiddistic Knowledge." *Philosophical Studies* 123(1):1–32.

Shoemaker, Sydney. 1984. Personal Identity: A Materialist's Account. In *Personal Identity*, ed. Sydney Shoemaker and Richard Swinburne. Blackwell.

Shoemaker, Sydney. 2003. "Realization, Micro-Realization and Coincidence." *Philosophy and Phenomenological Research* 67:1–23.

Stalnaker, Robert. 1976. "Propositions." *Issues in the Philosophy of Language* pp. 79–91.

Stanley, Jason. 2005. "Fallibilism about Concessive Knowledge Attributions." *Analysis* 65(2):126–31.

Swinburne, Richard. 1997. *The Evolution of the Soul*. Oxford University Press.

Teller, Paul. 1972. "Epistemic Possibility." *Philosophia* 2(4):303–20.

Thomson, Judith. 1983. "Parthood and Identity Across Time." *Journal of Philosophy* 80(201-20).

Thomson, Judith. 1998. "The Statue and the Clay." *Noûs* 32:148–73.

Unger, Peter. 1979. "There Are no Ordinary Things." *Synthese* 41:117–54.

Van Inwagen, Peter. 1990. *Material Beings*. Cornell University Press.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.

Yablo, Steve. 1987. "Identity, Essence and Indiscernibility." *Journal of Philosophy* 84:293–314.

Yalcin, Seth. 2007. "Epistemic Modals." *Mind* 116(464):983–1026.