Object Persistence in Philosophy and Psychology

BRIAN J. SCHOLL

Abstract: What makes an object the same persisting individual over time? Philosophers and psychologists have both grappled with this question, but from different perspectives—philosophers conceptually analyzing the criteria for object persistence, and psychologists exploring the mental mechanisms that lead us to experience the world in terms of persisting objects. It is striking that the same themes populate explorations of persistence in these two very different fields—e.g. the roles of spatiotemporal continuity, persistence through property change, and cohesion violations. Such similarities may reflect an underlying connection, in that psychological mechanisms of object persistence (especially relevant parts of mid-level visual object processing) may serve to underlie the intuitions about persistence that fuel metaphysical theories. This would be a way for cognitive science to join these two disparate fields, helping to explain the possible origins and reliability of some metaphysical intuitions, and perhaps leading to philosophical progress.

1. Introduction: Perceiving Persisting Objects

Among the most important discoveries of cognitive science has been the realization that most seemingly obvious and effortless mental tasks—e.g. recognizing faces and understanding speech—actually result from complex feats of cognitive processing. Our minds not only meet these challenges but do so in a way that can hide from us that there is a challenge being met in the first place. As a result, many research projects in cognitive science begin by pointing to some aspect of our mental lives that we typically take completely for granted, and then asking how it is possible. This paper is about one such seemingly obvious aspect of our mental lives: our experience of the world in terms of persisting objects.

Visual processing begins with an undivided wash of unbound features, and results in the visual experience of discrete objects and events. Accordingly, a tremendous amount of research has explored the principles by which the visual system segments images into units. However, visual experience consists of more than individual snapshots of the world: we must also bind individual views of objects into dynamic representations which persist across time and motion. Without such a computation of *persisting* objecthood, visual experience would be incoherent. In recent years, psychologists have made great strides in working out the principles that guide the construction and maintenance of representations of portions of the visual field as the *same* objects over time, motion, featural change, and interruptions such as occlusion.

For helpful comments on earlier drafts I thank Tamar Gendler, Frank Keil, Ron Mallon, Shaun Nichols, Nic Noles, Susanna Siegel, Ted Sider, and the members of the Yale Perception and Cognition Laboratory.

Address for correspondence: Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA.
Email: Brian.Scholl@yale.edu

Mind & Language, Vol. 22 No. 5 November 2007, pp. 563–591.

Psychological research on object persistence is not isolated to visual perception. Indeed, object persistence seems like a true case study in cognitive science, since it has been a focus of study in so many different disciplines. In the first place, questions about persistence are not restricted to vision: the same questions arise in the study of audition, for example, wherein we can hear multiple melodic contours persist over time and frequency changes. In this case each melody can be considered as a different auditory 'object' that persists and evolves over time, and researchers have studied how the mind binds sounds into these persisting auditory representations (see Bregman, 1990). In addition, object persistence has also been studied in the context of higher-level cognition, where cognitive psychologists have begun to work out the factors that influence our judgments of object persistence (e.g. Rips, Blok and Newman, 2006). And persistence is also one of the most well-studied phenomena in infant cognition—where it has been explored both explicitly in studies of infants' 'core knowledge' of simple physical principles (e.g. Cheries, Wynn and Scholl, 2006; Spelke, Kestenbaum, Simons and Wein, 1995; Xu and Carey, 1996) and implicitly in studies of numerical cognition (e.g. Wynn, 1992).

Beyond these diverse areas of psychology, object persistence has also been a focus of exploration in several other fields. Some of these projects are applied in nature: for example, computer scientists and engineers have been forced to explore real-world issues of object persistence as they construct computer-vision systems that can track objects over time—cars through traffic, or people through crowds (for a recent review see Yilmaz, Javed and Shah, 2006). Other projects are conceptual in nature: for example, philosophers have long grappled with questions and puzzles about object persistence as central challenges in metaphysics (for a recent review see Kurtz, 2006). These philosophical investigations will be a focus of the present paper.

1.1 Linking Philosophical and Psychological Persistence

Despite their diverse origins, these many studies of object persistence tend to traffic in the same underlying questions, answers, and research themes. The primary goal of this paper is to highlight such connections between studies of object persistence in philosophy and experimental psychology. Research on persistence in these two areas has almost always proceeded completely independently (for rare exceptions, discussed in Section 5.1, see Hall, 1998; Rips *et al.*, 2006; Xu, 1997), but I suggest that the many similarities between these research projects reflect an underlying connection: psychological mechanisms of object persistence (especially relevant parts of mid-level visual object processing) may serve to underlie the intuitions about persistence that fuel metaphysical theories. This would be a way for cognitive science to join these two disparate fields, helping to understand the possible origins and reliability of various metaphysical intuitions, and perhaps leading to philosophical progress.

As a part of analytic philosophy, the study of object persistence in metaphysics takes as its goal to analyze our concept of persistence—and persistence itself, as a

purely metaphysical issue—explaining what it means for an object to persist, and developing criteria that govern when objects do and do not persist (for reviews see Gallois, 1995; Haslanger, 2003; Kurtz, 2006; Sider 2000). The typical method in this area, as in most such philosophical projects, is a priori in nature: analyses are developed and then tested against our intuitions about when objects do and do not persist, both in general and in the context of particular situations and puzzles. In addition, such analyses can be tested for their fit with other analyses of neighboring metaphysical concepts (e.g. of space and time), in pursuit of a coherent theory of the nature of reality. I discuss the methodology of such projects—and in particular their use of intuitions about persistence—in more detail below, in Section 5.

This is a large and fractious literature, with many competing views of persistence, and the ensuing discussion here makes no attempt to be systematically inclusive of the various positions. I instead limit the focus of this paper to three of what I take to be the most foundational and salient themes from this philosophical investigation: (1) spatiotemporal continuity, (2) persistence through property change, and (3) cohesion. Sections 2, 3, and 4 of this paper respectively discuss these three themes in turn.[1]

Each of these sections introduces a theme of object persistence in its conceptual or philosophical context, and then discusses how that same theme has led to empirical discoveries in recent psychological research about the nature of object persistence in perception and cognition. This empirical work has not only explored the factors that determine when we experience objects as persisting (or not), but has also demonstrated that computations of object persistence are critical for understanding many other processes such as visual memory (e.g. Cheries *et al.*, 2006; Flombaum and Scholl, 2006), implicit learning (e.g. Fiser, Scholl and Aslin, 2007), numerical cognition (e.g. Feigenson, Dehaene and Spelke, 2004), motion perception (e.g. Dawson, 1991), search and foraging (e.g. Flombaum, Kundey, Santos and Scholl, 2004; Santos, 2004), and even visual awareness in the first place (e.g. Mitroff and Scholl, 2005; Moore and Lleras, 2005).

This work has been conducted in several subfields of experimental psychology, and the work reviewed below includes studies of adult visual cognition, as well as studies of 'core knowledge' of objects in both human infants and nonhuman primates. This collection of research is not arbitrary, as a growing number of researchers think of these areas as studying the same

[1] It may be worth noting that the 'objects' under investigation as discussed here will be assumed to be typical everyday objects (e.g. cars or basketballs), though these are often studied via their visual representations (as in studies of simple colored shapes on a computer monitor). In contrast, I do not discuss entities that are spatially extended beyond our immediate experience (e.g. lakes) or the sets that philosophers sometimes call 'unnatural' objects (e.g. the 'object' consisting of my left arm plus Wisconsin). What it means to be an 'object' in the first place is of course a deep question in its own right, but it will not be relevant to this discussion—though it has also seen its share of psychological progress in recent years (e.g. see Ben-Shahar, Scholl and Zucker, 2007; Feldman, in press; Scholl, 2001a).

underlying processes of visual cognition (Carey and Xu, 2001; Feigenson, Carey and Hauser, 2002; Scholl and Leslie, 1999). Collectively, this research converges on the idea that there are mental mechanisms for computing object persistence that are *primitive* in three senses: (1) they occur relatively early in perceptual processing, such that they are largely 'hardwired' parts of perception that then constraint later cognition; (2) they arise early in human infancy; and (3) they arose early enough in phylogenetic development to be shared by some of our fellow primates.

2. Theme #1: Spatiotemporal Continuity

What factors determine whether an object at time t1 and an object at time t2 are the same individual? The one intuitive constraint that immediately comes to mind is that of spatiotemporal continuity: for two objects encountered at different locations to be subsequent stages of the same individual, there must be a spatiotemporally continuous path between them. If an object disappears at one location, and an object immediately appears at a different spatially separated location, then those two instances cannot be the same object, since physical laws do not allow for that sort of thing (at least at the spatial and temporal scales that characterize our everyday interaction with objects).

Because of its intuitive appeal, spatiotemporal continuity has long been taken as a key principle of persistence in philosophical theorizing (see Burke, 1980; Coburn, 1971; Hirsch, 1976). '[S]patiotemporal continuity … [is] a logically necessary condition of identity' (Shoemaker, 1963, pp. 4-5), and it is simply 'part of our common-sense concept of a thing that its existence is spatiotemporally continuous. It never leaps gaps in either time or space' (Armstrong, 1963, p. 220). For various statements of spatiotemporal continuity as a key principle of persistence, including formal characterizations, see Coburn, 1971. Some discussions of spatiotemporal continuity seem to take it as a sufficient as well as a necessary criterion. For example: '[I]n a series of events which common sense would regard as belonging to one "thing", the similarity need only be between events not widely separated in space-time' (Russell, 1948, p. 488). Most theorists, however, take it as merely a necessary principle. This is in part because many philosophers are concerned with the persistence of objects considered under sortals (see Hirsch, 1982; Grandy, 2007; Wiggins, 2001; Xu, 1997, in press). For example, an object such as a car that is about to be crushed may soon cease to exist despite its continuous spatiotemporal trace with the resulting cube of metal (Hirsch, 1982); perhaps something continues to exist, but the *car* is gone. (Issues involving sortal concepts will not be considered in the present paper, in part for reasons that will become clear in Section 3.) Even these theories still give spatiotemporal continuity a bedrock role, however. For example, accounts which attempt to avoid large featural transformations of objects typically still consider such constraints in the context of a more basic principle of continuity; 'The basic rule is: Trace an object's career by following a spatiotemporally

and qualitatively continuous path which minimizes changes as far as possible' (Hirsch, 1976, p. 366).

Perhaps because physical laws simply do not allow for flagrant violations of spatiotemporal continuity, this principle has also become wired into our minds, and helps to control our experience of objects as persisting in the world. The research that demonstrates this indicates that the principle is primitive in three senses. First, a 'principle' of spatiotemporal continuity seems to characterize relatively automatic visual processing. We have demonstrated this, for example, using a multiple object tracking (MOT) task (Pylyshyn and Storm, 1988). This task in some ways forces observers to judge object persistence (but see Pylyshyn, 2004; Scholl, in press). In a typical MOT experiment (see Figure 1), observers initially see a number of qualitatively identical objects. A subset of these are then flashed to indicate their status as targets, after which all of the (again identical) objects begin moving independently and unpredictably about the display. When they stop moving, observers must indicate which of the objects are the original targets.[2] To explore the influence of spatiotemporal continuity on the ability to perform this task, we required observers to track multiple objects as they moved in a display filled with occluders (Scholl and Pylyshyn, 1999). Occlusion is a challenge to the implemention of any principle of continuity, given its ecological ubiquity: objects may not pop into and out of existence, but they do frequently pop into and out of *sight*. To show that a mere disruption of visual stimulation would not impair object persistence, we demonstrated that MOT is just as accurate when the objects frequently disappear for a moment behind the occluders (as in Figure 2b) as when they simply move atop those same contours (as in Figure 2a).

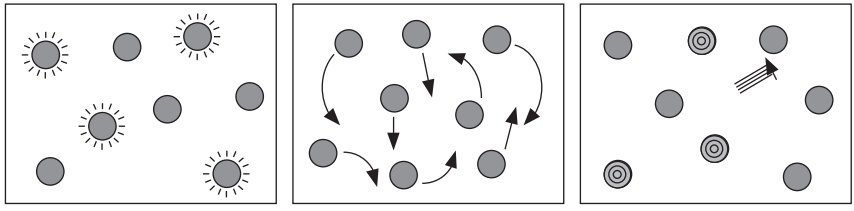

Figure 1 *A schematic depiction of multiple object tracking (MOT). (a) Four items are initially flashed to indicate their status as targets. (b) All items then begin moving independently and unpredictably around the display. (c) At the end of the motion phase, the subject must move the cursor about the screen to highlight the four targets—here the subject has just highlighted three of the targets, and is moving the mouse cursor to the fourth. Animations of this task and many others described in this paper can be viewed or downloaded over the internet via http://www.yale.edu/perception/*

[2] This type of task is inherently dynamic, of course. Readers may experience MOT—and many of the other visual paradigms and phenomena described in this paper—by viewing or downloading animations at http://www.yale.edu/perception/

Unimpaired MOT performance in the context of these occluders, however, required the presence of accretion and deletion cues along fixed contours at the occluding boundaries. In contrast, performance was significantly impaired when objects were present on the visual field at the same times and to the same degrees as in the occlusion conditions, but disappeared and reappeared in ways which did not implicate the presence of occluding surfaces—e.g. by imploding and exploding into and out of existence, instead of accreting and deleting along a fixed contour (see Figure 2c). Thus, when these subtle visual cues indicate that the momentary disappearances reflected the objects going out of sight, the disappearances did not affect tracking—but when these cues indicated that the objects were going out of existence momentarily, observers were no longer able to track them as persisting individuals, despite their brute visual similarity to the 'trackable' displays (Scholl and Feigenson, 2004). Critically, notice that this inability did not reflect observers' beliefs or strategies: they knew perfectly well what was going on in the 'implosion/explosion' conditions, and would liked to have ignored those cues. But they could not—indicating that this particular type of principle of spatiotemporal continuity is encapsulated from our beliefs and preferences, and is part of the basic processes that help to generate visual experience.

Spatiotemporal continuity also appears to constitute a primitive psychological principle in at least two other senses, relating to ontogenetic and phylogenetic development. First, this principle comes online relatively early in development. This can in part be shown using the same manipulations discussed above: 10-month-old infants can maintain a representation of the number of moving objects (2 versus 3) in a display with occluders when those objects occlude and disocclude, but not when

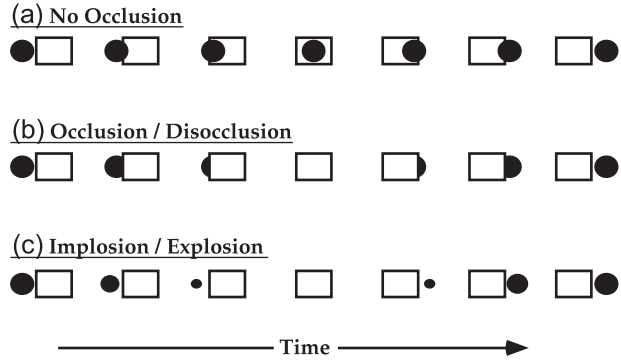

Figure 2 *Depictions of how an object might interact with an occluder in three conditions in studies of multiple object tracking through occlusion (Scholl and Feigenson, 2004; Scholl and Pylyshyn, 1999). (a) The object does not disappear, but simply moves atop the other contours. (b) The object is momentarily occluded—deleting from its trailing edge as it disappears and then accreting from its leading edge upon its reappearance. Tracking in this condition is unimpaired. (c) The object moves through the same trajectory, disappearing at the same time and the same rate as in (b), but now it gradually 'implodes' to a point, and then later 'explodes' from a point. This violates spatiotemporal continuity, greatly impairing tracking*

they implode and explode (Cheries, Feigenson, Scholl and Carey, 2005; see also Bertenthal, Longo and Kenny, 2007; Kaufman, Csibra and Johnson, 2005). Other studies have demonstrated an appreciation for spatiotemporal continuity in infancy using a more direct method. For example, the looking-time behavior of 5-month-old infants who observe a rod that moves sequentially behind two screens (see Figure 3a) indicates that they perceive the event in terms of a single persisting object, but they instead seem to perceive two distinct objects if the rod does not traverse the space in between the two screens (Figure 3b; Spelke *et al.*, 1995; see also Aguiar and Baillargeon, 2002, for similar experiments with even younger infants). Moreover, even nonhuman primates seem to make this inference (Santos, 2004): if rhesus macaques observe a piece of food roll behind a screen and then a piece of food roll between the screens, disappearing behind the second screen (as in Figure 3a, but now with objects rolling down a ramp), their resulting foraging behavior indicates that they think there was only a single piece of food involved. But if the piece of food does not appear between the two screens (as in Figure 3b), then their resulting foraging behavior indicates an appreciation that there were two pieces of food involved. In each of these cases, the critical factor controlling looking times and searching behavior seems to be the principle of spatiotemporal continuity—which in turn seems to be wired into our minds in a deep way, controlling how we experience the world.

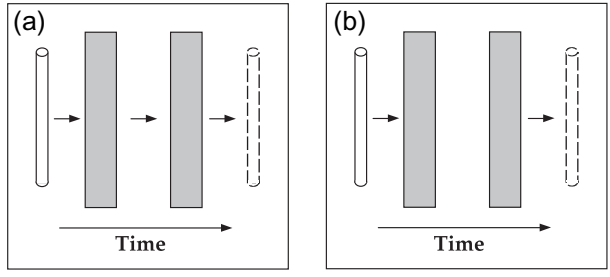

Figure 3 *A schematic depiction of events used to study spatiotemporal continuity in studies of infants (Spelke et al., 1995) and nonhuman primates (Santos, 2004; with objects rolling down ramps behind screens). (a) An object moves behind a first screen, then an object moves between the two screens, then an object emerges from the second screen and continues moving (after which it may change direction and these events occur again in reverse order, in some experiments where the event cycles). It is natural to experience this event (as infants and monkeys seem to do) as a single object that moves behind two screens. (b) When the same event proceeds but without an object ever appearing between the two screens, it is natural to experience it (as infants and monkeys seem to do) as involving two distinct objects, one moving behind each screen*

3. Theme #2: Persistence Through Property Change

A second salient theme in philosophical theories of object persistence involves the need to reconcile persistence over time with changes in the properties of objects. How, for example, can we say that a banana is the same object over time when it

is yellow at time t1, but brown at time t2? At first blush this can seem like a silly question, but it has been one of the primary obsessions for philosophers at least as far back as Plato and Aristotle. Leibniz, for example—as part of the formulation of what we now call Leibniz' Law—noted the difficulty in simultaneously maintaining three beliefs that each individually seem intuitive: (1) Nothing can have conflicting properties (e.g. being both fully brown and fully yellow); (2) Changes may involve conflicting properties (e.g. turning from yellow to brown); but (3) Objects can persist as the same individuals through changes. When these beliefs are combined, there is trouble. This problem is very simple to formulate, but it could not be a more central challenge to metaphysical theorizing about object persistence. (For example, it is essentially the very first issue raised in almost every recent review of this literature; e.g. Haslanger, 2003; Kurtz, 2006; Wiggins, 2001.) It also arises in more particular puzzles, for example about whether a statue is identical to the lump of clay from which it was formed (e.g. Gibbard, 1975).

Despite its simplicity, this challenge—'the traditional problem of change' (Sider, 2000)—ends up requiring exceedingly complex, nuanced solutions. Most currently popular solutions come in one of three types, which I will briefly describe because they will re-enter the discussion below, in Section 5, in a discussion of the connection between empirical results and philosophical intuitions. *Perdurance* theories (e.g. Lewis, 1986; Quine, 1950) suggest that objects have *temporal parts* just as they have spatial parts (for extensive discussion, see Sider, 2001a). My right arm is a part of me, but so is the way I was at exactly noon today—and I (as an object) extend into the past just as I extend in height. On this view, an object is simply a sequence of fleeting *stages* ('all the world's a stage'; Sider, 1996)—but because only one of these stages actually exists at any given moment, an object is never fully present at any given time (just as I am never fully present in any single 2D spatial plane). This ducks the problem of change by relativizing changes to time: it's not that the banana is both yellow and brown, but that one of its (temporal) parts is yellow, and another of its parts is brown—and those parts/stages never exist at the same time. But, counterintuitively, change can no longer be thought of as alteration on this view: because the *object* extends backward and forward in time as a 'space-time worm', its properties never actually *change*, per se.

Exdurance theories (e.g. Hawley, 2002; Sider, 1996) are similar: they are also 'stage' theories that appeal to temporal parts, but they claim that an object is identical with a given stage. Thus an object is fully present at any given moment (unlike a space-time worm), but it is never the same object from moment to moment—and so again it cannot truly change. (For our purposes here, the difference between perdurance and exdurance theories is subtle, and will not be important.) Finally, *Endurance* theories (e.g. van Inwagen, 1990) are in some ways the most intuitive accounts of persistence (see Haslanger, 2003), since unlike the stage theories they maintain that the very same object is present from moment to moment: I am here now, and now, and now—and I do not need to be thought of as extending into the past and future the way that I extend in space. Endurantists must thus appeal to other maneuvers in order to relativize properties to individual

times (and thus avoid contravening Leibniz' Law), but in doing so they create other problems, counterintuitively denying that objects can have ('intrinsic') properties irrespective of particular times.

The persistence of objects through property changes has also been a salient theme in empirical work on our experience of object persistence—but less as a challenge, and more as a brute phenomenon. In general, there are two importantly different routes that can lead you to identify an object as the same over time. First, you can note what the object looks like: if you see an animal that is short and brown, and then you see an animal that is tall and yellow, then they are unlikely to be the same animal. Second, you can identify objects as the same based on how and where they move through the local visual environment: if you see an animal appear from behind one tree, and then see an animal appear from behind a different tree, then they cannot be the same animal, for the reasons discussed in the previous section. But critically, this second sort of "sameness" is independent of the first. If spatiotemporal continuity is violated, then the two animals must be different, even if they look the same. In contrast, changes in properties do not necessarily foil sameness, because of property change.

This conceptual asymmetry expresses itself in many empirical demonstrations of *spatiotemporal priority* (Scholl, 2001b): across many different paradigms and phenomena, when deciding whether an object is the same persisting individual from some earlier time, factors relating to how and where that object has moved will always trump factors relating to what the object looks like. Perhaps the most direct demonstration of spatiotemporal priority is in the 'tunnel effect' (Burke, 1952; Michotte *et al.*, 1964/1991): one object moves behind an occluder (the 'tunnel'), and then a very different (-looking) object emerges from the other side of the occluder and continues moving (Figure 4). If the second object emerges at about the time and place that one

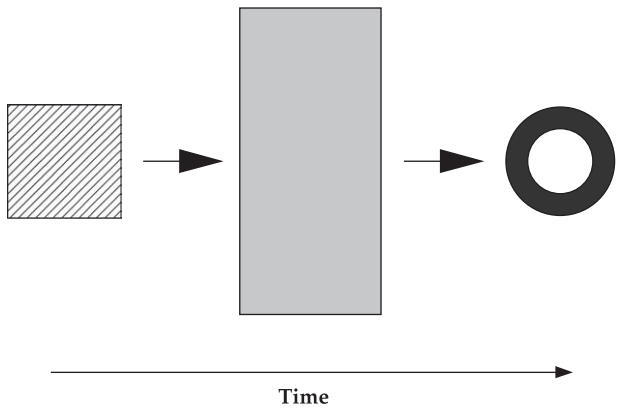

Figure 4 *The 'tunnel effect': an object moves behind an occluder, and then an object with very different surface features emerges on the other side of the occluder. When the emerging object appears just where and when it should had the initial object continued its motion, we irresistibly perceive a single persisting object that changed its surface features (Flombaum and Scholl, 2006; Michotte et al., 1964/1991)*

would expect the first object to emerge, had it continued its motion, we tend to perceive the uninterrupted and uniform motion of a single object. Critically, the tunnel effect obtains even when the object emerging from the tunnel differs from the one that entered in terms of its surface features (e.g. turning from red to green; Burke, 1952) and its kind (e.g. turning from a kiwi to a lemon; Flombaum *et al.*, 2004). In such situations, spatiotemporally continuous motion leads the visual system to discount the featural difference between the pre- and post-tunnel objects, and thus to generate the percept of a single object (which changes its features while occluded) rather than the percept of two separate objects, one of which was initially hidden by the tunnel. In contrast, when an extra temporal delay is introduced between the occlusion of the first object and the emergence of the second, observers perceive the successive motion of two different objects—the first of which must remain hidden by the tunnel (Burke, 1952).

This effect is not simply an isolated visual illusion, but can end up playing a major role in determining the contents of our perceptual experience. This was made especially salient in a recent comparative study exploring free-ranging rhesus macaques' perception of the tunnel effect in a foraging task (Flombaum *et al.*, 2004). As depicted in Figure 5, Monkeys watched as a lemon rolled down a ramp and came to rest behind a tunnel (Occluder 1), and then as a kiwi emerged and became occluded at the end of its path behind a screen (Occluder 2). When the kiwi emerged at about the time that the lemon should have (had it continued its motion), subjects searched for food only behind Occluder 2—apparently perceiving

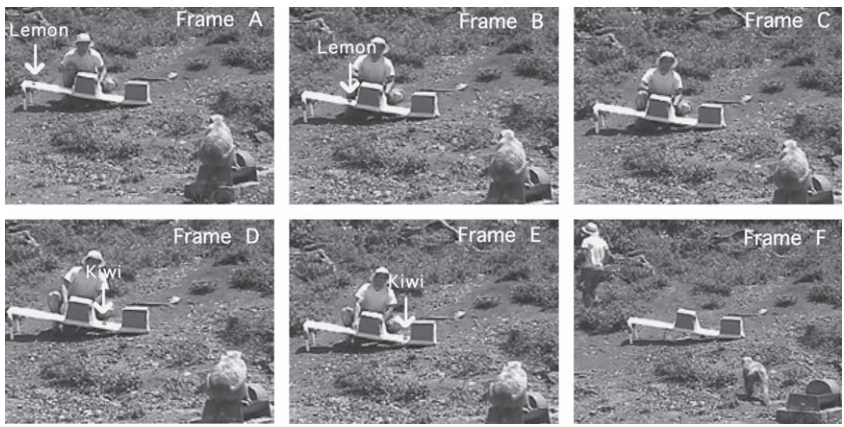

Figure 5 *Sample stills from a demonstration of the tunnel effect in rhesus monkeys (Flombaum, Kundey, Santos and Scholl, 2004). The subject watches as a lemon rolls down the ramp (Frames A-B) and becomes concealed by Occluder 1 (Frame C), and then as kiwi rolls down the remainder of the ramp (Frames D-E) and becomes concealed by Occluder 2 (Frame F). Finally, the experimenter walks away, and the subject is given the opportunity to search (Frame F). When the kiwi appears to continue the spatiotemporal trajectory of the lemon (with no temporal gap), the subjects search for food only behind the second occluder—apparently perceiving a single persisting object which changed from a lemon to a kiwi*

the lemon transform into a kiwi on the basis of spatiotemporally continuous motion. (We included other control conditions to verify that the monkeys did in fact recognize the featural difference.) In contrast, when a brief pause interrupted the occlusion of the lemon and the emergence of the kiwi, monkeys searched for food behind both occluders—apparently perceiving two distinct objects, and assuming (on the basis of that featural difference) that the lemon must have remained in the tunnel. In this way the tunnel effect directly influenced not only the monkeys' percepts, but also their subsequent spontaneous behaviors.

This type of spatiotemporal priority also seems to be a primitive part of adult visual cognition: it is not a decision that we make, but a reflection of the underlying machinery of perception. In human adults, for example, the tunnel effect mediates automatic visual working memory (Flombaum and Scholl, 2006). In these experiments, observers had to detect color changes in displays where several objects oscillated behind occluders and occasionally changed color. Across comparisons with several types of spatiotemporal gaps, as well as manipulations of occlusion versus implosion (as described in the previous section), performance was better when objects' kinematics gave the impression of a persisting individual. In other words, the results revealed a *temporal same-object advantage*: better change detection across temporal scene fragments bound into the same persisting object representations—suggesting that persisting objects are the underlying units of visual memory. Critically, this was true despite the intentions of observers in this experiment: they knew that the presence of spatial or temporal gaps was irrelevant to their task, yet these manipulations (and others such as implosion/explosion) irresistibly impaired their performance.

Finally, this type of spatiotemporal priority also manifests itself in terms of our visual percepts themselves, and not only in the tunnel effect. For example, in apparent motion, two featurally-dissimilar flashes will still be perceived as subsequent stages of a single object so long as they occur in quick enough succession and in nearby locations (Kolers and Pomerantz, 1971). As in Figure 6, for example, you might see frame 1 followed by frame 2—quickly enough that the visual system is led to perceive motion. But this presents a correspondence problem: which went where? Based on your knowledge about the possibility and likelihood of property

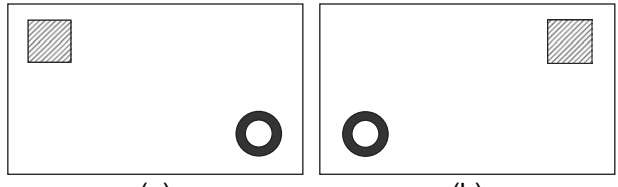

Figure 6 *An example of figural change in apparent motion. Observers see the two frames flashed in quick succession, so that the visual system infers motion. Despite the salient figural identities, however, observers will not see horizontal motion: instead, the vertical proximity trumps surface-feature processing, and observers see two vertically translating objects that dramatically change their visual features as they move*

changes, you might judge that the square in frame 1 moved to the right, while the disc moved to the left. But that is not what you'll *see*: since the objects are separated by a greater distance in the horizontal direction, spatiotemporal priority will dominate, and you'll irresistibly see the objects translate vertically, swapping all of their features in the process (e.g. Burt and Sperling, 1981; Navon, 1976).

The studies described here collectively illustrate that persistence through featural change is a part of how we experience the world in a primitive sense: our visual systems appear to be constructed so that featural changes are often irrelevant for determining object persistence. As described in more detail in the following section, the visual system appears to represent objects via representations called 'object files', and the automatic processes that control the construction, destruction, and maintenance of object files seem wired to discount featural changes. Such changes may still be salient, but they are limited to updating the *contents* of object files, and cannot directly control their creation or maintenance (Mitroff and Alvarez, in press; but see Moore, Mordkoff and Enns, 2007).

4. Theme #3: Cohesion

If spatiotemporal continuity is the most obvious principle involved mediating object persistence (as described in Section 2), then perhaps the most obvious problems for this principle are violations of *cohesion*—wherein an object fails to maintain a single cohesive boundary over time. In such cases, there is no single continuous space-time 'worm' to track; instead, this worm splits into multiple strands. One of the simplest types of cohesion violations that has been much discussed in the philosophical literature is that of *fission*, wherein one object splits into two (e.g. Parfit, 1971; Williams, 1956-1957; cf. Locke, 1689/1975). The 'objects' in many such examples are people: via various science-fiction scenarios (e.g. involving Star-Trekesque transporter malfunctions), one person ends up resulting in two (often exactly similar) people. The conceptual possibility of such examples is taken to pose a deep problem for theories of persistence: diachronic numerical identity cannot hold between an single object and two distinct objects, but by hypothesis in such examples there is nothing to prefer one possible continuant over the other. (In contrast, an example of cloning doesn't present a serious problem, since there would then be reasons related to spatiotemporal continuity for preferring the 'original' person as the same persisting individual, rather than the clone.) Such 'fission cases' have been central to philosophical theorizing about object persistence ever since such projects began, at least back to the 18th century (see Martin, Barresi and Giovannelli, 1998).

The metaphysical problem posed by such examples echoes in other areas of philosophy—e.g. in ethics, where we might wonder which of the resulting two people can or should be held responsible for past crimes committed by the initial single person. As a result, fission examples are especially popular in the enormous literature devoted specifically to issues of *personal identity* (e.g. see Baillie, 1993;

Gendler, in press a; Noonan, 1993), and these debates often end up turning on psychological issues unrelated to physical objects—e.g. involving memory. Nevertheless, the essential problem posed by fission examples also applies to physical objects—and so while I will sometimes mention examples involving people, this will not be important: this type of puzzle 'is not so much about *personal* identity through fission, as it is about the identity through fission of any dividing object … . After all, the fission of any object … might seem 'doubly perplexing' for those who believe that there must be a true answer to questions about its identity in any describable case' (Merricks, 1997, p. 181).

In any case, there are also several hallowed philosophical problems involving cohesion violations in more mundane physical objects. For example, suppose a watchmaker, in repairing a watch, takes it apart and temporarily stores the various pieces in different locations before putting it back together (e.g. Wiggins, 2001, chapter 3). In this case we are intuitively drawn to say that the resulting watch is the same individual that was initially taken apart—but how can this be reconciled with the fact that it traced no single spatiotemporally continuous path during its repair? Perhaps the most famous puzzle involving a cohesion violation, of course, is that of the 'Ship of Theseus' (e.g. Hobbes, 1672/1913): a ship is repaired gradually, one plank at a time, until none of its initial materials remain. Meanwhile, a scavenger has collected each discarded plank, and eventually uses them to rebuild the original ship. The puzzle: which of the two resulting ships is the same individual as the initial ship? In other words, which is the Ship of Theseus?

Puzzles involving cohesion violations end up requiring one of several types of revisions to theories of object persistence: for example, they may appeal to the idea that there were really somehow two objects that constituted the initial entity or that there is a single resulting object with two spatially separated parts (e.g. Lewis, 1976)—or they may (radically) give up on the one-to-one logic of the identity principle (e.g. Gallois, 1998; Geach, 1997). For our purposes, the important point about fission is simply that such cases are problematic precisely because they challenge our intuitions about what it means for an object to persist over time. It is not just that we do not know what to say in such cases, but that we are led directly to apparent conflicts. For example, if half of an object erodes away over time, we are inclined to intuit that it remains that same object. But if instead the object is simply broken in two in an instant, we are inclined to intuit that neither of the resulting halves is identical to the initial object—even if one of them were somehow the exact duplicate of the eroded object.

The literature on object persistence in experimental psychology has also grappled with cohesion violations, and for many of these same reasons. Here cohesion is often considered as the one of the principles of "core knowledge": an object must maintain a single bounded contour in order to be represented as persisting over time (e.g. Spelke, 1990, 2000). Indeed, this principle may be uniquely important, in that it helps define what counts as an object in the first place. If you want to know what an object is, just "grab some and pull"; the stuff that comes with your hand is the object, and the stuff that doesn't is not. This has led some theorists to

claim that cohesion is perhaps the single most important principle of what it means to be an object (e.g. Bloom, 2000; Pinker, 1997).

Some of the resulting empirical literature has explored cohesion violations via the contrast between objects and nonsolid substances. For example, infants' ability to keep track of small sets of objects is disrupted if the objects (e.g. sand piles) pour from one location to another rather than moving as unified wholes (Huntley-Fenner, Carey and Solimando, 2002). This is true for adult visual cognition as well. Using the multiple-object tracking task described in Section 2, for example, observers were asked to track spatially extended objects that repeatedly moved in a particular type of noncohesive motion: each object began as a small square, but then split into many smaller units and moved in a non-rigid manner—essentially 'pouring' from one location to another, as would a nonsolid substance. This manipulation greatly impaired tracking, despite the fact that the 'objects' still followed the same trajectories as typical MOT control conditions (vanMarle and Scholl, 2003). (This is the sort of inherently dynamic display that isn't really worth trying to depict in a figure, but the animations can be viewed online; see footnote 2.) These examples involve cohesion violations—similar but more extreme than watch disassembly—but not in the sense that characterizes most fission examples in philosophy. After all, subjects (or their visual systems) may conclude from such displays that the relevant entities are not (and never were) objects at all, but were rather nonsolid substances. In contrast, the objects in most philosophical fission cases—e.g. the watch, the person, or the Ship of Theseus—were bona fide objects before the fission occurred.

Additional research, however, has also explored the ways in which visual processing copes with fission cases that are almost a direct analogue of the relevant philosophical examples (albeit not using Star-Trek-like transporters). This research has exploited the 'object reviewing' paradigm as illustrated in Figure 7. In this task

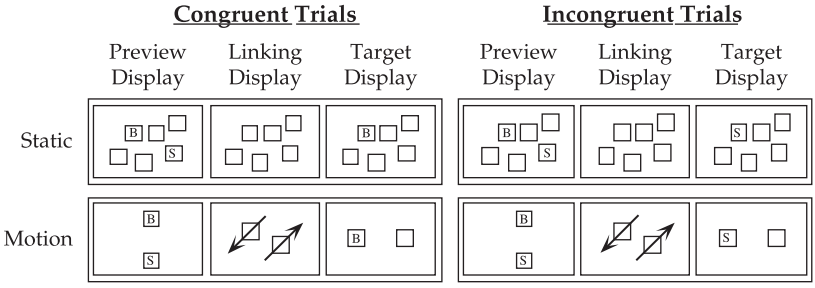

Figure 7 *Sample displays used in the object-reviewing paradigm (Kahneman et al., 1992). Subjects see two preview letters in an initial display and a single probe letter in a final display, and must simply indicate whether the final probe was present in the initial preview display. Responses are faster when the probe letter had initially appeared in that same object during the preview display, compared to when it had appeared in the other object. In the static displays, the target is seen as the same object as one of the previews, because it appears on the same object, in the same location. Objecthood and location are unconfounded in the moving displays*

(Kahneman, Treisman and Gibbs, 1992), a number of objects are initially presented, and distinct letters (or other distinguishable features) appear briefly on some of them. The objects then move about the visual display for a brief period, after which a single probe letter appears on just one of the objects. The subject's task, in one variant of this paradigm (see Noles, Scholl and Mitroff, 2005) is to use a keypress to indicate whether the final probe letter was the same as *any* of the initially presented letters. This response will be slightly faster when the probe is the same letter that initially appeared on *that same object*, compared to when it was the letter that had initially appeared on a different object—an *object-specific preview benefit* (OSPB). This effect can thus be used as an index of persisting objecthood: manipulations which degrade enduring object representations will result in attenuated OSPBs (e.g. Gao and Scholl, under review; Mitroff, Scholl and Wynn, 2005; Noles *et al.*, 2005).

One recent study reasoned that if cohesion was really such a powerful 'core' principle of object persistence, then even the simplest possible cohesion violation—a single object splitting into two—should impair processing (Mitroff, Scholl and Wynn, 2004). We tested this using the object reviewing paradigm with displays in which one object split into two resulting objects, as depicted in Figure 8. (This figure, not accidentally, is similar to those figural illustrations of fission cases in many philosophical treatments—e.g. Johnston, 1989, p. 383; Sider, 1996, p. 435.)

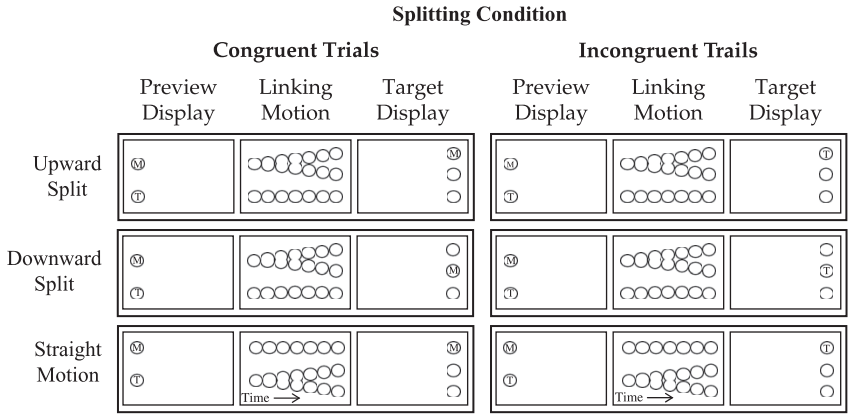

Figure 8 *A subset of the trial types from a study of cohesion violations using the object reviewing paradigm (Mitroff, Scholl and Wynn, 2004). In each case observers simply responded as quickly as possible whether the final letter appeared anywhere in the initial preview display on that trial. Objects either traveled a straight trajectory, a curved trajectory, or smoothly split into two separate objects. (The actual experiment also included No-Match trials for each condition, in which the final probe letter was neither of the initially presented letters, and factors such as the relative orientations were always counterbalanced and randomized, such that the final probe letter could appear on any final object.) As described in the text, these cohesion violations produced considerable costs in visual processing, relative to controls that are not depicted here*

Critically, this splitting was smooth and symmetrical, so that there was nothing to bias one of the resulting objects over the other. (A second non-splitting object is also present in each display as a control.) The theoretical question concerns what happens to the underlying object representation of the initial object—its 'object file'—as a result of the fission. There are several possibilities that correspond to various philosophical positions on fission cases: (1) the cohesion violation might obliterate the object file's contents entirely; (2) the information about the preview letter might survive the 'splitting' intact, but stay bound to only one of the two resulting objects, indicating that object files cannot themselves split into two; or (3) the object file's contents might essentially be 'copied' to both of the two resulting objects. In other words, this experiment assesses how the visual system copes with the same fission cases that have proven so problematic for philosophy.

The results of this study indicated that although we can of course readily perceive an object that splits into two, this produces great underlying costs in visual processing. In the first case, the OSPBs that resulted from such manipulations were often only half the size of those that occurred in control conditions with similar movement but without cohesion violations. In addition, such processing disrupted the maintenance of object-specific information on other objects in the display—e.g. destroying the OSPB for the additional non-splitting object in Figure 8. These costs suggest that the visual system ends up effectively splitting the actual object representations that underlie the object(s) in this event, such that one object file is eventually replaced with two object files, into which the initial contents of the first file are copied. (There is no indication in these studies of bimodal response patterns, which would suggest that the object file 'went with' only one of the post-split objects.) The fact that this process resulted in substantial costs to visual processing, however, indicates that this event is seen as exceptional, and requires more resources to handle than is the norm in the perception of simple events.

Emboldened by this result, we later showed that a single object splitting into two also severely disrupts infants' persisting object representations (Cheries, Mitroff, Wynn and Scholl, in press). In a forced-choice crawling paradigm, 10- and 12-month-old infants witnessed crackers being sequentially placed into containers, and they subsequently crawled toward the container with the greater cracker quantity. When one of the crackers was visibly split in half, however, infants failed to represent the relative quantities, despite careful controls for the overall quantities and the motions involved: now they crawled randomly toward one of the two containers, rather than preferring the one with the greater cracker quantity. This impairment was especially striking given that infants of this age must frequently see real objects (e.g. pieces of food) that are split into two. The collective punchline of all these studies is then similar to those in the previous sections: these experiments indicate that cohesion—and the processing of cohesion violations—is represented in a primitive way in visual processing, such that it operates automatically in both adults visual cognition and in infancy.

5. Can Psychological Discoveries about Object Persistence Drive Philosophical Progress?

The previous sections described three salient themes from philosophical treatments of object persistence—spatiotemporal continuity, persistence through featural change, and cohesion—that have also each been foci of empirical research from the study of visual cognition in experimental psychology. Why might such connections be important?

5.1 A Heuristic for Generating Ideas (and an Aside on 'Experimental Philosophy')

In the first place, pointing out such connections could have some heuristic value in generating ideas and hypotheses in each field. This has been the case, for example, in a few previous papers that have attempted to link philosophical and psychological accounts of object persistence. Sometimes this has involved importing philosophical concepts into empirical work—e.g. in 'sortal'-based theories of infant individuation (Xu and Carey, 1996). Such identifications may then lead back to implications for philosophical theories, if the empirical evidence is taken to lead them in new directions (Xu, 1997). In other cases, psychologists have simply taken various puzzles from the philosophical literature, and then tested subjects' intuitions on various cases that systematically manipulate relevant variables, leading to discoveries about the importance of new factors, or the developmental progression of the relevant intuitions. For instance, the 'Ship of Theseus' problem has been used in this way (Hall, 1998), as have other fission examples (Rips *et al.*, 2006). This work might be considered as an example of 'experimental philosophy' (cf. Nichols, 2004), since its methods essentially replicate what philosophers do. Whereas philosophers may consult their own intuitions about various cases, these psychologists are consulting others' intuitions about such cases, albeit using various types of more sensitive measures (e.g. rating scales) and then quantitatively comparing the results of various manipulations.

This sort of experimental philosophy can be dangerous, though, and I suspect that such studies, rather than telling us anything about underlying mental mechanisms, may instead often tell us more about how subjects respond to bizarre questions and scenarios. In particular, the value obtained by the systematic manipulation of variables in such contexts—especially when this is done within-subjects—is offset by a cost: such designs may end up introducing task demands that can easily lead to misinterpretations of the data. To take a concrete example, one experimental study of fission cases (Rips *et al.*, 2006) had subjects judge whether an object was the same individual after it was described as undergoing a bizarre procedure in which it was destroyed, but a sophisticated computer then reproduced a copy. In these examples, the researchers varied—within subjects—the percentage of molecules of the object that were described as having been perfectly copied from the original. When later asked to judge whether the copy

(or in other cases which of two copies) was the same individual as the original object, they found that the subjects' answers were systematically affected by the percentage of identical molecules: more of the same molecules, more likely to be the same individual. They concluded that such experiments 'provide evidence about which of the experimental factors affect decisions about individual identity' (p. 11)—but I doubt this is so. Naïve subjects in such experiments are forced to make *some* response, but what are they to do if they are simply baffled by the questions, or if they have no strong intuitions? In this case, they grasp for the obvious variable that is being manipulated across trials, and use that to structure their answers in a systematic way—even though this may not reflect any true judgments about identity. Indeed, there are great task demands in such experiments to respond in this manner—i.e. to make use of a salient manipulation; to do otherwise in such cases would often require giving an identical indeterminate answer on each and every trial of the experiment, and there is always implicit pressure against doing this. This is one way in which such examples of 'experimental philosophy' can tell us more about heuristic getting-through-the-experiment strategies than about actual metaphysical intuitions.

In contrast, note that the experimental work described in the previous sections is not 'experimental philosophy' in the methodological sense, and that each of these experiments uses more implicit response measures that help to ensure that the results reflect underlying mental mechanisms of object persistence. In any case, the connections I wish to draw here between philosophical and psychological work on object persistence are very different from those noted above. While it is true that these connections may sometimes help to heuristically motivate new experiments (e.g. Hall, 1998; Rips *et al.*, 2006)[3] or interpretations of experiments (Xu, 1997), I will focus instead on how psychological mechanisms of object persistence (especially relevant parts of mid-level visual object processing) may interact with philosophical intuitions in the first place.

5.2 Origins of Metaphysical Intuitions

On its face, the philosophical project of analyzing object persistence seems to be a completely different type of investigation than the empirical project of determining how persisting object representations are constructed in the visual system. One is about considered judgment and about the world itself, while the other is about our experience of the world. I think there is a deeper link between these projects, though, and this link turns on the use of intuitions as 'data' in philosophical theorizing. Philosophy can easily strike empirical researchers as a curious discipline,

[3] Of course, study of the philosophical literature could also be of heuristic use in helping to extend the types of studies discussed in this paper. For example, it would be relatively straightforward to adapt our object reviewing experiments to cases of fusion as well as fission, and to study additional factors such as the constraint that two objects cannot be located in the same place at the same time (see Mitroff, Scholl and Wynn, 2005).

because of its reliance on intuition. As noted in the opening paragraph of this paper, perhaps the central lesson of cognitive science has been that certain intuitions (about the nature and fidelity of cognitive processing) are radically mistaken, and that intuitions in general are a notoriously unreliable guide to how the mind actually works. Yet, intuitions play an extremely powerful role in constraining philosophical theorizing. This is obvious in projects of conceptual analysis, but in practice intuitions also seem to play a dominant role in 'pure' metaphysical theorizing about the world itself.

Sometimes the role of intuition is made explicit in analytic philosophy: the entire project is to devise a conceptual scheme which accounts for the intuitions in various cases. Our intuitions are not always the sole constraints in such projects—since for example we may also want to test philosophical accounts for their coherence with other neighboring accounts of related concepts (in pursuit of 'reflective equilibrium')—but they do play a decisive role in practice. This is certainly true in much (though perhaps not all) of the metaphysical theorizing discussed in the present paper. The most salient methodological practice in the philosophical literature on object persistence, at least in my reading of it, is to (1) create clever new puzzles and scenarios, (2) note what various theories are committed to saying about them, (3) note that these commitments contrast with our intuitions about what we *should* say in such cases, and then (4) attempt to revise the theories to match our intuitions. Sometimes this use of intuitions is made explicit. In theorizing about object persistence, for example, '[m]any philosophers share a deep commitment to … metaphysical theses that express some of our surest and most firmly held intuitions … . Our goal … should be to strike the best balance between achieving philosophical beauty in terms of elegant, coherent metaphysical theories of persistence and *matching our intuitions*' (Kurtz, 2006, pp. 1, 15).[4] In other cases, this use of intuitions is simply apparent in practice, for example when a sufficient reason to reject a theory is just that 'This strikes me as obviously false' (Thomson, 1983, p. 213, commenting on the claim of some 'temporal parts' theories that objects exist only momentarily, being constantly replaced by non-identical objects). In any case, many philosophers seem unwilling to simply abandon our intuitions about object persistence entirely.[5]

In meta-philosophical discussions, there seem to be a wide range of attitudes toward the use of intuitions (for a representative collection of views see DePaul and Ramsey, 1998). Some philosophers offer broad defenses of their power, in support of classical conceptual analysis (e.g. Bealer, 1996, 2000). Others claim that intuition-mongering (and the conceptual analysis that it fuels) are of little use

---

[4] Note, though, that it is sometimes difficult to tell just what philosophers mean by 'intuitions'—since in some cases this term may refer to any beliefs about abstract things that are not a direct result of perception—even if those beliefs are the result of lengthy consideration by experts.

[5] One exception, from an article entitled 'Can amoebae divide without multiplying?': '[I]f we wish to recognize, as we should, the possibility of persistence through fission, … we must swallow our intuitions, based as they are on habits of thought and speech adapted to situations in which such possibilities rarely, if ever, eventuate, and revert to the explicitly identity-based idioms of pure metaphysics' (Robinson, 1985, p. 319).

except as ways of aiming empirical inquiry in useful directions (e.g. Cummins, 1998; Kornblith, 1998; see also Stich, 1990; Weinberg, Nichols and Stich, 2001). Others take an intermediate view (e.g. Levin, 2004). Regardless of how ultimately useful intuitions are in philosophical theorizing, though, it seems well worth asking about where they come from, and how they are produced.

I suggest that that psychological mechanisms of object persistence—especially the kinds of mid-level visual object processing reviewed above—may serve to underlie the intuitions about persistence that fuel metaphysical theories. Many intuitions, after all, are not appreciated in an abstract, Platonic sense, but are rather part of how we experience and imagine the world. Intuitions may generate 'pull on the intellect' (Sider, 2001b, p. 201) in some cases, but this pull may not have an intellectual origin (e.g. arising as a direct result of other beliefs); rather, it may arise because of the way our minds work. With regard to blatant violations of spatiotemporal continuity, for example, it may be true that 'Our physical laws do not allow us to believe in the possibility of such spatiotemporal translations' (Scaltsas, 1981, p. 402). But this connection between physical laws and beliefs is not direct; rather, it is indirectly mediated by mental reflexes that were presumably fixed during natural selection because they matched so well with the laws that constrain the behavior of matter in the real world. And because those laws—and the character of our environments—were unchanged (or even unchangeable) for so long during the course of our evolution, these mental habits may now be firmly entrenched in perception. In short, our intuitions about the nature of persisting individuals may derive from the way we *experience* the world in terms of persisting individuals. And those experiences may in turn arise from specific representations and mechanisms of mind—such as 'object files' and mid-level object-based visual routines—whose nature we can empirically determine.

The kind of connection I am suggesting may hold not only in 'online' perceptual experience, but also in the imagery, simulation, and tacit reasoning that accompanies the evaluation of philosophical puzzles. These bring online the relevant representations and processes—and they bring online memories of relevant experiences when those representations and processes were also operating—and these processes and memories may in turn influence our attitude toward the cases we are evaluating. Indeed, these mechanisms of mind may in some sense be inescapable. Because they are primitive processes—operating, as do many or most visual processes, in reflexive ways that are unswayed by our current goals and beliefs (see Pylyshyn, 1999)—they effectively serve as lenses through which we cannot help experiencing and imagining the world. Even when we formulate higher-level beliefs about what may be occurring in a certain scenario, for example, we may still not be able to help perceiving it in a different way. In the 'tunnel effect' described in Section 3, for example, we simply *see* strange events such as lemons turning into kiwis (Flombaum and Scholl, 2006; Flombaum *et al.*, 2004; cf. Figure 5) even when we know that cannot be really happening. As such, these experiences ('perceptual intuitions?') may still 'pull on the intellect' despite our higher-level judgments, leading us to feel that certain interpretations are more or less natural.

This connection between psychological mechanisms of object persistence and philosophical intuitions about object persistence may hold in a general sense, but it may also help drive our intuitions in particular cases:

Why do we judge that spatiotemporal continuity is so intrinsic to object persistence (Section 2)? Perhaps because that principle is wired into how we experience the world in the first place. In some cases we irresistibly perceive objects as persisting individuals because of their spatiotemporally continuous trajectories, or we find ourselves unable to track objects as persisting individuals because they violate spatiotemporal continuity (Cheries *et al.*, 2005; Scholl and Feigenson, 2004; Scholl and Pylyshyn, 1999; Spelke *et al.*, 1995). Because this principle is part of how we process object tracking in the first place, our minds may simply not have the ingredients or experiences that would allow us to imagine persistence in its absence. There might be philosophical reasons for wanting to do so, but everything about our experiences with persisting objects in the world will pull on the intellect in the opposite direction.

Why do we judge that objects continue to exist despite changes to their properties (Section 3)? Perhaps because our visual *representations* of objects persist through property changes. These representations—'object files'—may store and update visual features, but they are controlled primarily (or even solely) by spatiotemporal factors (Flombaum and Scholl, 2006; Flombaum *et al.*, 2004; Kahneman *et al.*, 1992; Mitroff and Alvarez, in press). From this perspective, the relevant observation is not that we *can* perceive objects as persisting despite changes to their features, but that we *cannot help* but see objects as persisting in some cases, despite changes to those properties, as in Figure 6—and this is because in turn our visual systems are wired to represent the relevant objects via a single updated 'object file' representation. This 'representational persistence' in mid-level visual cognition may then in turn fuel our intuitions about actual metaphysical persistence in the world itself. Endurance theories of persistence, for example, may simply strike many readers as much more natural and compelling than perdurance or exdurance theories because endurance theories are more closely matched to the actual way in which we experience the world. (It may be a metaphysical question whether an object persists over time in a given scenario, but it is an increasingly tractable empirical question whether the object in that scenario would be represented by the same online visual representation over time—though our laboratory, at least, has not yet managed to procure a Star Trek transporter.) Philosophers may think, for example, that they are preferring an endurance theory because 'a number of our practices and forms of self-understanding depend on the idea that there are enduring things' (Haslanger, 2003, p. 351)—but I suggest that what many of these 'practices' and 'forms of self-understanding' may come down to in practice are the results of the mental machinery we have that is devoted to computing enduring representations of objects over time.

Finally, why do we judge that cohesion violations are so likely to disrupt object persistence, as in fission examples? Perhaps because these very cohesion violations, when they are actually experienced, *do* disrupt our persisting object representations

(Cheries *et al.*, in press)—or at least give rise to considerable costs in perceptual processing (Mitroff *et al.*, 2004). In contrast, more complicated fission examples such as the Ship of Theseus (as discussed in Section 4) seem to leave us intuitively adrift: when many people contemplate such examples, there just seems to be a breakdown in judgment. The 'renovated' ship seems to be the same object in some ways, but the 'scavenged' ship seems like the same object in other ways. What's going on here? This might be a case where conflicting intuitions arise due to judgments in different parts of the mind (see Nichols, 2004). Perceptual simulation might continue to suggest that the renovated ship is the Ship of Theseus, because if you were to watch it sail around while any particular renovation was occurring, you would represent it via the same object file. More abstract judgments based on brute similarity or origin might still—and at the same time—suggest to you that the scavenged ship is the Ship of Theseus, though, since the parts of your mind that generate these judgments are distinct from those parts that are cleaving to spatiotemporal continuity.

5.3 Judging Intuitions, and the Heuristic Nature of Perception

If understanding psychological mechanisms of object persistence could help us to understand where some metaphysical intuitions come from, then following through on this project could also help to fuel philosophical progress. In particular, if such links could be made more concrete we might gain some insight into which intuitions we should trust, and in what circumstances. This is true because while some psychological mechanisms may be reflections of underlying physical laws, most are not. Psychological principles—including those principles that are wired into the visual system and that collectively constitute the 'logic of perception' (Rock, 1983)—are typically fallible *heuristic* principles that get things right most of the time, but fail spectacularly in some special circumstances. This, for example, is what drives almost all visual illusions. Perception is fraught with ambiguity, and cannot be implemented as a type of deductive inference. Instead, the visual system must make assumptions about the world in order to arrive at useful percepts. Such assumptions—for example that objects are rigid and three-dimensional (e.g. Ullman, 1979), or that there is a single overhead light source (Hershberger, 1970)—can be wrong in certain circumstances, however, which leads to the familiar visual illusions of depth and lightness. In short, the visual system involves (and must involve) probabilitistic assumptions about the nature of the environment (see Scholl, 2005) and thus will fail in those unusual circumstances when the assumptions are not satisfied.

But 'unusual circumstances', to say the least, is what philosophical puzzles in metaphysics (and beyond) seem to be all about! Such puzzles—with their fake barns, stopped watches, swampmen, and Star Trek transporter malfunctions—are deemed useful precisely because they are so unusual, and thus hopefully diagnostic of the relevant intuitions. In practice, this frequently causes our intuitions to lean in unexpected directions, which in turn puts pressure on the philosophical theories

that do not accord with those unexpected intuitions. It is possible, though, that these intuitions are like visual illusions: they are simply *mistakes* brought on by the unusual circumstances, giving us the wrong answer. This will be hard to diagnose conceptually, though: such intuitions may seem just as compelling as those based on more pedestrian cases, just as a visual illusion of depth is no less compelling than a percept of actual depth.

This analysis might seem to lead us to question the use of such outlandish examples in cases involving object persistence (e.g. Gendler, 1998, 2002, in press b). The outlandishness per se may not be problematic, but these scenarios may violate assumptions about the world that are made in a reflexive way by our perceptual and cognitive processes. This is consistent with previous discussions of the danger of such thought experiments. For example: '[C]ertain patterns of features that coincide only fortuitously may none the less play a central role in the organization of our concepts. To the extent that imaginary scenarios involve disruptions of these patterns, our first-order judgments about them may be distorted or even inverted' (Gendler, 2002, p. 34). And on fission cases in particular: 'Fission cases simply confirm … what all of us … already know: the facts that inform our everyday judgments of identity over time are such that it is broadly logically possible that they lead us astray, or simply fail to lead us' (Merricks, 1997, p. 180). In summary, a philosophical intuition that is fueled in part by mechanisms of visual cognition may be suspect, at least to the extent that it involves contexts in which those perceptual mechanisms generate incorrect interpretations. Such perceptual illusions may, in the service of metaphysical theorizing, become cognitive illusions that can lead us astray.

5.4 Conclusions: Philosophical Progress?

As my philosophical colleagues are constantly reminding me, experimental psychology grew out of philosophy in the first place. Having diverged rather dramatically in the intervening century, however, they are now growing closer again in some subtle ways. Some of this new flirtation is due to empirical progress on topics that have previously resisted scientific investigation—e.g. involving consciousness, or reference. And some of these new connections are also due to the recent injection of experimental techniques into philosophy itself (see Knobe and Nichols, in press). These interactions are not always comfortable, though, and often such connections can feel threatening or deflationary. Certainly the preceding discussion could be read in this way, as it questions the value of intuitions—especially those derived from bizarre metaphysical scenarios.

I suggest, though, that the empirical study of psychological mechanisms of object persistence could also help drive progress in philosophy in a positive sense, and for just these reasons. If we can learn more about where our metaphysical intuitions come from, then we may be able to make more informed choices about which intuitions to give up when they conflict. This could be useful, to the extent that philosophical theorizing requires having to favor one intuition over another.

But, at least in metaphysical theorizing about object persistence, this is a very great extent. Indeed, the entire project of object persistence in philosophy has been characterized in terms of deciding which intuitions to prioritize: '[T]he real problem of persistence remains one of balancing trade-offs. To explain how objects persist …, we must revise and/or forfeit some of our basic intuitions' (Kurtz, 2006, p. 24). In this context, understanding the origins of our metaphysical intuitions in various psychological mechanisms could help us understand when they are worth revising or forfeiting in our philosophical theories, especially if there is reason to think that those psychological mechanisms may yield unreliable results in the particular contexts in which they are being asked to operate. (If, for example, a theory of persistence demands that we give up the intuition that property changes involve true alteration of persisting objects, then we may be comforted to know that the reason we have that intuition in the first place is not because of any intellectual obligations due to other beliefs, but only because in visual experience we represent property changes by altering persisting *representations* of objects.) In general, understanding where metaphysical intuitions come from is one way that psychology might usefully contribute to philosophy, rather than trying to replace it.

Department of Psychology
Yale University

References

Aguiar, A. and Baillargeon, R. 2002: Developments in young infants' reasoning about occluded objects. *Cognitive Psychology*, 45, 267–336.

Armstrong, D. 1963: Absolute and relative motion. *Mind*, 72, 209–223.

Baillie, J. 1993: Recent work on personal identity. *Philosophical Books*, 34, 193–206.

Bealer, G. 1996: A priori knowledge and the scope of philosophy. *Philosophical Studies*, 81, 121–142.

Bealer, G. 2000: A theory of the a priori: Intuition, evidence, concept-possession. *Pacific Philosophical Quarterly*, 81, 1–30.

Ben-Shahar, O., Scholl, B.J. and Zucker, S.W. 2007: Attention, segregation, and textons: Bridging the gap between object-based attention and texton-based segregation. *Vision Research*, 47, 845–860.

Bertenthal, B.I., Longo, M.R. and Kenny, S. 2007: Phenomenal permanence and the development of predictive tracking in infancy. *Child Development*, 78, 350–363.

Bloom, P. 2000: *How Children Learn the Meanings of Words*. Cambridge, MA: MIT Press.

Bregman, A.S. 1990: *Auditory Scene Analysis: The Perceptual Organization of Sound*. Cambridge, MA: MIT Press.

Burke, L. 1952: On the tunnel effect. *Quarterly Journal of Experimental Psychology*, 4, 121–138.

Burke, M.B. 1980: Cohabitation, stuff, and intermittent existence. *Mind*, 89, 391–405.

Burt, P. and Sperling, G. 1981: Time, distance, and feature trade-offs in visual apparent motion. *Psychological Review*, 88, 171–195.

Carey, S. and Xu, F. 2001: Infant knowledge of objects: Beyond object files and object tracking. *Cognition*, 80, 179–213.

Cheries, E.W., Mitroff, S.R., Wynn, K. and Scholl, B.J. in press: Cohesion as a principle of object persistence in infancy. *Developmental Science*.

Cheries, E.W., Wynn, K. and Scholl, B.J. 2006: Interrupting infants' persisting object representations: An object-based limit? *Developmental Science*, 9, F50–F58.

Cheries, E., Feigenson, L., Scholl, B.J. and Carey, S. 2005: Cues to object persistence in infancy: Tracking objects through occlusion vs. implosion [Abstract]. *Journal of Vision*, 5, 352a.

Coburn, R.C. 1971: Identity and spatio-temporal continuity. In M.K. Munitz (ed.), *Identity and Individuation*. New York: NYU Press, pp. 51–101.

Cummins, R. 1998: Reflections on reflective equilibrium. In M. DePaul and W. Ramsey (eds), *Rethinking Intuition*. Lanham, MD: Rowman and Littlefield, pp. 113–127.

Dawson, M.R. 1991: The how and why of what went where in apparent motion: Modelling solutions to the motion correspondence problem. *Psychological Review*, 98, 569–603.

DePaul, M. and Ramsey, W. 1998: *Rethinking Intuition*. Lanham, MD: Rowman and Littlefield.

Feigenson, L., Carey, S. and Hauser, M. 2002: The representations underlying infants' choice of more: Object files versus analog magnitudes. *Psychological Science*, 13, 150–156.

Feigenson, L., Dehaene, S. and Spelke, E.S. 2004: Core systems of number. *Trends in Cognitive Sciences*, 7, 307–314.

Feldman, J. in press: Formation of visual 'objects' in the early computation of spatial relations. *Perception and Psychophysics*.

Fiser, J., Scholl, B.J. and Aslin, R.N. 2007: Perceived object trajectories during occlusion constrain visual statistical learning. *Psychological Bulletin and Review*, 14, 173–178.

Flombaum, J.I. and Scholl, B.J. 2006: A temporal same-object advantage in the tunnel effect: Facilitated change detection for persisting objects. *Journal of Experimental Psychology: Human Perception and Performance*, 32, 840–853.

Flombaum, J.I., Kundey, S.M., Santos, L.R. and Scholl, B.J. 2004: Dynamic object individuation in rhesus macaques: A study of the tunnel effect. *Psychological Science*, 15, 795–800.

Gallois, A. 1995: Identity over time. In E. Zalta (ed.), *Stanford Encyclopedia of Philosophy* Spring 2005 Edition.

Gallois, A. 1998: *Occasions of Identity*. Oxford: Clarendon Press.

Gao, T. and Scholl, B.J. under review: When are objects required for object files?: The roles of segmentation and postdiction in computing object persistence. Manuscript submitted for publication.

Geach, P. 1997: Selections from 'Reference and Generality'. In M. Rea (ed.), *Material Constitution*. Lanham, Maryland: Rowman and Littlefield. pp. 305–312.

Gendler, T.S. 1998: Exceptional persons: On the limits of imaginary cases. *Journal of Consciousness Studies*, 5, 592–610.

Gendler, T.S. 2002: Personal identity and thought experiments. *Philosophical Quarterly*, 52, 34–54.

Gendler, T.S. in press a: Personal identity and metaphysics. In B. McLaughlin and A. Beckermann (eds), *Oxford Handbook of the Philosophy of Mind*. Oxford: Oxford University Press.

Gendler, T.S. in press b: Philosophical thought experiments, intuitions, and cognitive equilibrium. *Midwest Studies in Philosophy: Philosophy and the Empirical*.

Gibbard, A. 1975: Contingent identity. *Journal of Philosophical Logic*, 4, 187–221.

Grandy, R. 2007: Sortals. In E. Zalta (ed.), *Stanford Encyclopedia of Philosophy* Summer 2007 Edition.

Hall, D.G. 1998: Continuity and the persistence of objects: When the whole is greater than the sum of the parts. *Cognitive Psychology*, 37, 28–59.

Haslanger, S. 2003: Persistence through time. In M. Loux and D. Zimmerman (eds), *Oxford Handbook of Metaphysics*. Oxford: Oxford University Press, pp. 314–354.

Hawley, K. 2002: *How Things Persist*. New York: Oxford University Press.

Hershberger, W. 1970: Attached-shadow orientation perceived as depth by chickens reared in an environment illuminated from below. *Journal of Comparative and Physiological Psychology*, 73, 407–411.

Hirsch, E. 1976: Physical identity. *Philosophical Review*, 85, 357–389.

Hirsch, E. 1982: *The Concept of Identity*. New York: Oxford University Press.

Hobbes, T. 1672/1913: *Metaphysical System of Hobbes in Twelve Chapters*. Chicago, IL: Open Court Publishing.

Huntley-Fenner, G., Carey, S. and Solimando, A. 2002: Objects are individuals but stuff doesn't count: Perceived rigidity and cohesiveness influence infants' representations of small groups of distinct entities. *Cognition*, 85, 203–221.

Johnston, M. 1989: Fission and the facts. *Philosophical Perspectives*, 3, 369–397.

Kahneman, D., Treisman, A. and Gibbs, B.J. 1992: The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24, 174–219.

Kaufman, J., Csibra, G. and Johnson, M.H. 2005: Oscillatory activity in the infant brain reflects object maintenance. *Proceedings of the National Academy of Sciences*, 102, 15271–15274.

Knobe, J. and Nichols, S. (eds.) In press: *Experimental Philosophy*. New york: Oxford University Press.

Kolers, P. and Pomerantz, J. 1971: Figural change in apparent motion. *Journal of Experimental Psychology*, 87, 99–108.

Kornblith, H. 1998: The role of intuition in philosophical inquiry: An account with no unnatural ingredients. In M. DePaul and W. Ramsey (eds), *Rethinking Intuition*. Lanham, MD: Roman and Littlefield, pp. 129–141.

Kurtz, R.M. 2006: Introduction to persistence: What's the problem? In S. Haslanger and R.M. Kurtz (eds), *Persistence: Contemporary Readings*. Cambridge, MA: MIT Press, pp. 1–26.

Levin, J. 2004: The evidential status of philosophical intuition. *Philosophical Studies*, 121, 193–224.

Lewis, D. 1976: Survival and identity. In A. Rorty (ed.), *The identities of persons* pp. 17–40: University of California Press.

Lewis, D. 1986: *On the Plurality of Worlds*. Oxford: Blackwell.

Locke, J. 1689/1975: On identity and diversity. Chapter 27 of *Essay Concerning Human Understanding*. Reprinted in J. Perry (ed.), *Personal Identity*. Berkeley, CA: University of California Press, pp. 33–52.

Martin, R., Barresi, J. and Giovannelli, A. 1998: Fission examples in the Eighteenth and early Nineteenth century personal identity debate. *History of Philosophy Quarterly*, 15, 323–348.

Merricks, T. 1997: Fission and personal identity over time. *Philosophical Studies*, 88, 163–186.

Michotte, A., Thinès, G. and Crabbé, G. 1964/1991: Les complements amodaux des structures perceptives. In *Studia Psychologica*. Louvain: Publications Universitaires. Reprinted and translated as: Michotte, A., Thinès, G. and Crabbé, G. 1991: Amodal completion of perceptual structures. In G. Thines, A. Costall and G. Butterworth (eds) *Michotte's Experimental Phenomenology of Perception*. Hillsdale, NJ: Erlbaum.

Mitroff, S.R. and Alvarez, G.A. in press: Space and time, not surface features, underlie object persistence. *Psychonomic Bulletin and Review*.

Mitroff, S.R. and Scholl, B.J. 2005: Forming and updating object representations without awareness: Evidence from motion-induced blindness. *Vision Research*, 45, 961–967.

Mitroff, S.R., Scholl, B.J. and Wynn, K. 2004: Divide and conquer: How object files adapt when a persisting object splits into two. *Psychological Science*, 15, 420–425.

Mitroff, S.R., Scholl, B.J. and Wynn, K. 2005: The relationship between object files and conscious perception. *Cognition*, 96, 67–92.

Moore, C.M. and Lleras, A. 2005: On the role of object representations in substitution masking. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 1171–1180.

Moore, C.M., Mordkoff, J.T. and Enns, J.T. 2007: The path of least persistence: Object status mediates visual updating. *Vision Research*, 47, 1624–1630.

Navon, D. 1976: Irrelevance of figural identity for resolving ambiguities in apparent motion. *Journal of Experimental Psychology: Human Perception and Performance*, 2, 130–138.

Nichols, S. 2004: Folk concepts and intuitions: From philosophy to cognitive science. *Trends in Cognitive Sciences*, 8, 514–518.

Noles, N., Scholl, B.J. and Mitroff, S.R. 2005: The persistence of object-file representations. *Perception and Psychophysics*, 67, 324–334.

Noonan, H. (ed.) 1993: *Personal Identity*. Dartmouth Publishing Company.

Parfit, D. 1971: Personal identity. *Philosophical Review*, 80, 3–27.

Pinker, S. 1997: *How the Mind Works*. New York: Norton.

Pylyshyn, Z.W. 1999: Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22, 341–423.

Pylyshyn, Z.W. 2004: Some puzzling findings in multiple object tracking MOT: I. Tracking without keeping track of object identities. *Visual Cognition*, 11, 801–822.

Pylyshyn, Z.W. and Storm, R.W. 1988: Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, 3, 179–197.

Quine, W.V. 1950: Identity, ostension, and hypostasis. *Journal of Philosophy*, 47, 621–633.

Rips, L., Blok, S. and Newman, G. 2006: Tracing the identity of objects. *Psychological Review*, 113, 1–30.

Robinson, D. 1985: Can amoebae divide without multiplying? *Australasian Journal of Philosophy*, 63, 299–319.

Rock, I. 1983: *The Logic of Perception*. Cambridge, MA: MIT Press.

Russell, B. 1948: *Human Knowledge, its Scope and Limits*. New York: Simon and Schuster.

Santos, L.R. 2004: 'Core Knowledges': a dissociation between spatiotemporal knowledge and contact-mechanics in a non-human primate? *Developmental Science*, 72, 164–174.

Scaltsas, T. 1981: Identity, origin, and spatiotemporal continuity. *Journal of Philosophy*, 56, 395–402.

Scholl, B.J. 2001a: Objects and attention: The state of the art. *Cognition*, 80, 1–46.

Scholl, B.J. 2001b: Spatiotemporal priority and object identity. *Cahiers de Psychologie Cognitive*, 20, 359–371.

Scholl, B.J. 2005: Innateness and Bayesian visual perception: Reconciling nativism and development. In P. Carruthers, S. Laurence and S. Stich (eds), *The innate mind: Structure and contents*, pp. 34–52: Oxford University Press.

Scholl, B.J. in press: What have we learned about attention from multiple object tracking (and vice versa)? In D. Dedrick and L. Trick (eds), *Computation, Cognition and Pylyshyn*. Cambridge, MA: MIT Press.

Scholl, B.J. and Feigenson, L. 2004: When out of sight is out of mind: Perceiving object persistence through occlusion vs. implosion [Abstract]. *Journal of Vision*, 4, 26a.

Scholl, B.J. and Leslie, A.M. 1999: Explaining the infant's object concept: Beyond the perception/cognition dichotomy. In E. Lepore and Z. Pylyshyn (eds), *What is Cognitive Science? Oxford*: Blackwell, pp. 26–73.

Scholl, B.J. and Pylyshyn, Z.W. 1999: Tracking multiple items through occlusion: Clues to visual objecthood. *Cognitive Psychology*, 38, 259–290.

Shoemaker, S. 1963: *Self-knowledge and Self-identity*. Ithaca, NY: Cornell University Press.

Sider, T. 1996: All the world's a stage. *Australasian Journal of Philosophy*, 74, 433–453.

Sider, T. 2000: Recent work on identity over time. *Philosophical Books*, 41, 81–89.

Sider, T. 2001a: Criteria of personal identity and the limits of conceptual analysis. *Philosophical Perspectives*, 15, 189–209.

Sider, T. 2001b: *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford: Oxford University Press.

Spelke, E.S.1990: Principles of object perception. *Cognitive Science*, 14, 29–56.

Spelke, E.S. 2000: Core knowledge. *American Psychologist*, 55, 1233–1243.

Spelke, E., Kestenbaum, R., Simons, D.J. and Wein, D. 1995: Spatiotemporal continuity, smoothness of motion and object identity in infancy. *British Journal of Developmental Psychology*, 13, 113–142.

Stich, S. 1990: *The Fragmentation of Reason*. Cambridge, MA: MIT Press.

Thomson, J.J. 1983: Parthood and identity across time. *Journal of Philosophy*, 80, 201–220.

Ullman, S. 1979: The interpretation of structure from motion. *Proceedings of the Royal Society of London*, B203, 405–426.

Van Inwagen, P. 1990: Four-dimensional objects. *Noûs*, 24, 245–255.

vanMarle, K. and Scholl, B.J. 2003: Attentive tracking of objects vs. substances. *Psychological Science*, 14, 498–504.

Weinberg, J., Nichols, S. and Stich, S. 2001: Normativity and epistemic intuitions. *Philosophical Topics*, 29, 429–460.

Wiggins, D. 2001: *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.

Williams, B. 1956-1957: Personal identity and individuation. *Proceedings of the Aristotelian Society*, 57, 229–252.

Wynn, K. 1992: Addition and subtraction by human infants. *Nature*, 358, 749–750.

Xu, F. 1997: From Lot's wife to a pillar of salt: Evidence that physical object is a sortal concept. *Mind & Language*, 12, 365–392.

Xu, F. in press: Sortal concepts, object individuation, and language. *Trends in Cognitive Sciences*.

Xu, F. and Carey, S. 1996: Infants' metaphysics: the case of numerical identity. *Cognitive Psychology*, 30, 111–153.

Yilmaz, A., Javed, O. and Shah, M. 2006: Object tracking: A survey. *ACM Computing Surveys*, 384, #13.