

# What Are Centered Worlds?\*

Shen-yi Liao

*The Philosophical Quarterly* (2012)<sup>†</sup>

## Abstract

David Lewis argues that centered worlds give us a way to capture *de se*, or self-locating, contents in philosophy of language and philosophy of mind. In recent years, centered worlds have also gained other uses in areas ranging widely from metaphysics to ethics. In this paper, I raise a problem for centered worlds and discuss the costs and benefits of different solutions. My investigation into the nature of centered worlds brings out potentially problematic implicit commitments of the theories that employ them. In addition, my investigation shows that the conception of centered worlds widely attributed to David Lewis is not only problematic, but in fact not his.

This paper raises a problem for centered worlds and discusses the costs and benefits of different solutions. In recent years, centered worlds have gained uses in a variety of philosophical sub-disciplines, from metaphysics to ethics. Given the diverse applications and theoretical usefulness of centered worlds, the problem and solutions raised in this paper have wide ramifications. Examining the nature of centered worlds brings out implicit commitments of the theories that employ them.

§1 motivates the present investigation into what centered worlds are by reviewing their applications. The original motivation for positing centered worlds is to have a way of adequately capturing the self-locating contents of attitudes such as belief. Since then, philosophers have given centered worlds new uses in the metaphysics of properties, foundations of two-dimensional semantics, and relativist theories in aesthetics, epistemology, metaethics, and philosophy of language.

§2 considers existing answers to the question of what centered worlds are and argues that they are inadequate. Roughly, a centered world is an ordered set of a possible world and a perspective within it, a center. Although some metaphysical issues regarding centered worlds are simply inherited from possible worlds, the issue unique to centered

---

\*I am most grateful to Andy Egan for his patient guidance and constant encouragement. For additional comments and discussions, I thank Derek Ball, Aaron Bronfman, Nate Charlow, Michael Della Rocca, Eric Guindon, Sungil Han, Jason Konek, Neil Mehta, Daniel Nolan, Alex Plakias, Jonathan Schaffer, Wolfgang Schwarz, Dustin Tucker, Jonathan Vertanen, two anonymous referees, and participants of the 2007 UT-Austin Graduate Philosophy Conference and the 2008 Australasian Association of Philosophy Conference.

<sup>†</sup>PENULTIMATE VERSION. For citation and reference, please see the definitive and final version in *The Philosophical Quarterly* 62, pp. 294-316. doi: [10.1111/j.1467-9213.2011.00042.x](https://doi.org/10.1111/j.1467-9213.2011.00042.x).

[end of p. 294] worlds boils down to what centers are. The current literature offers two answers to this question. The *Quinean account* says that a center is an ordered set of space-time coordinates. The *Lewisian account* says that a center is an ordered set of a persisting person and a time coordinate. Both accounts face counterexamples: the former cannot correctly capture co-locating individuals' self-locating beliefs and the latter cannot correctly capture time-traveling individuals' self-locating beliefs. These counterexamples reveal a deeper problem with identifying the centers of centered worlds. Unless the ordered set of features with which a center is identified is exhaustive, there can be two individuals that share the same specified features.

§3 presents two novel solutions to this problem and assesses their costs and benefits. The *exhaustive set account* identifies a center with an exhaustive ordered set of features. The *primitive identification account* treats the identification of centers as primitive, unable to be elucidated through other features. In the end, which solution is more preferable depends on one's other theoretical commitments. Indeed, although there are other ways to evade the counterexamples in §2—for example, denying the possibility of co-location—they inevitably involve controversial and substantive theoretical commitments. Centered worlds are not as theoretically innocent as they initially appear. Given that users of centered worlds are likely to have different theoretical commitments, it is unlikely that they have the same understanding of what centered worlds are.

§4 explores David Lewis's position on centered worlds. Considering Lewis's statements elsewhere and his other theoretical commitments, it seems that he in fact endorses the primitive identification account, and not the Lewisian account. Although the primitive identification account leaves the nature of centers somewhat mysterious, this mysteriousness is to be expected given the main lesson from the problem of essential indexicals. Furthermore, although centers are not to be identified with any other features, they can still be roughly characterized and heuristically indicated by features such as space-time coordinates.

## 1 What Good Are Centered Worlds?

To motivate an investigation into what centered worlds are, I will first explain what centered worlds are good for. This section begins with the problem of essential indexicals. The Lewisian diagnosis of this problem says that positing centered worlds is necessary to represent the self-locating contents of attitudes such as belief. After considering this influential use of centered worlds, I briefly review recent applications of centered worlds in addressing other philosophical issues.

[end of p. 295]

### 1.1 The Problem of Essential Indexicals

The initial motivation for positing centered worlds comes from the problem of essential indexicals, as discussed by John Perry, David Lewis, and Robert Stalnaker.<sup>1</sup> On

---

<sup>1</sup>The seminal articles are, respectively, (Perry 1979), (Lewis 1979), and (Stalnaker 1981). Their discussions, in turn, owe much to Hector-Neri Castañeda's work on the semantics of quasi-indicators; see the articles cited in (Lewis 1979, footnote 6).

one prominent traditional picture, the objects of propositional attitudes such as belief are *propositions*, which are sets of possible worlds.<sup>2</sup> Roughly speaking, **the problem of essential indexicals** is that propositions seem unable to capture finer distinctions that we intuitively make with contents of attitudes. Even supposing that two thoughts have the same content at the propositional level—meaning that they pick out the same set of possible worlds—there might nevertheless be an intuitive difference between the thought that involves an indexical—such as *I*, *here*, or *now*—and the thought that involves a non-indexical designation of the same thing.<sup>3</sup>

To get a concrete grasp of the problem, begin by considering Perry's messy-shopper case, as told from Perry's own perspective:

I once followed a trail of sugar on a supermarket floor, pushing my cart down the aisle on one side of a tall counter and back the aisle on the other, seeking the shopper with the torn sack to tell him he was making a mess. With each trip around the counter, the trail became thicker. But I seemed unable to catch up. Finally it dawned on me. I was the shopper I was trying to catch. (Perry 1979, 3)

Before it dawned on him that he was the messy shopper, Perry already thought to himself *the messy shopper is leaving a sugar trail*. However, after it dawned on him that he was the messy shopper, it seems that Perry came to believe something new and different, *I am leaving a sugar trail*, or more precisely, *I am the messy shopper who is leaving a sugar trail*. One way of seeing the difference between the earlier belief that does not involve the indexical *I* and the later belief that does is to compare their motivational influences. It is only when Perry came to believe *I am the messy shopper who is leaving a sugar trail* that he stopped pushing his cart up and down the aisle in search of the messy shopper.

[end of p. 296]

Although Perry's earlier belief *the messy shopper is leaving a sugar trail* and Perry's later belief *I am leaving a sugar trail* have contents that correspond to the exact same set of worlds—worlds in which Perry himself or his counterparts are leaving sugar trails—there nevertheless seems to be an important distinction between the two beliefs.<sup>4</sup> When it dawned on him that he was the messy shopper, Perry intuitively learned something new. But he did not learn anything new about what the world is like, or equivalently, which of the possible worlds he is in. Instead, he learned something new about which of the *possible individuals* he is—he is one who is leaving a sugar trail. The indexical *I* in the later belief is essential to expressing what Perry has just learned: a fact about who *he* is.

<sup>2</sup>I follow Lewis in setting aside the structured view of propositions. Even if one adopts the structured view of propositions for mental and linguistic content, one might nevertheless need centered worlds for, say, modeling contexts.

<sup>3</sup>Strictly speaking, indexicals are linguistic items that are part of sentences, not what the sentences express. However, I will use the term loosely. When I talk about thoughts with indexicals, I am really talking about *self-locating* thoughts that are expressible by sentences containing indexicals.

<sup>4</sup>Although speaking of counterparts is presentationally convenient, the problem of essential indexicals does not presuppose the acceptance of counterpart theory.

There are other cases in the literature that illustrate the problem of essential indexicals in other ways. Perry (1979) presents the case of an amnesiac lost in the Stanford library who does not know where *here* is. Stalnaker (1981) presents the case of a person who fell asleep in a trunk who does not know when *now* is. These cases all bring out the same lesson: beliefs involving indexicals importantly differ from beliefs involving non-indexical terms that refer to the same things because only the former have contents about which possible individual one is. Furthermore, although I have only talked about beliefs so far, the same lesson applies to contents of other mental attitudes and linguistic expressions.

## 1.2 The Lewisian Diagnosis

Many philosophers nowadays accept the following **Lewisian diagnosis** of the problem of essential indexicals.<sup>5</sup> Possible worlds are representational devices that illuminate the notion of *de dicto* possibility. Different possible worlds represent the different ways that a *world* might have been.<sup>6</sup> Hence, possible worlds suffice for capturing **world-locating contents**. By locating our world within a set of possible worlds, we rule out other ways the world might have been. However, as it turns out, possible worlds cannot capture all possibility. The ruling-out of possible worlds is insufficient for some finer distinctions we draw in logical space, namely the different ways an *individual* might have been within the same possible world. Sometimes, such as in Perry's messy-shopper case, in addition to ruling out other ways one's *world* might have been, one also wants to rule out other ways *oneself* might have been. Hence, there need to be different representational devices that suffice for capturing **self-locating contents**.

Enter centered worlds. Some illustrations may be helpful to introduce what centered worlds are supposed to be. While a possible world includes all objective information about the way that things in that world could be, no information is given about the different subjective perspectives one might occupy within a world. So if we think of possible worlds as entities that individuate objective information within logical space, then we can similarly think of centered worlds as entities that individuate subjective information within logical space. Borrowing Lewis's metaphor, a possible world is like a map and a centered world is like a map with an arrow pointing to a spot saying "you are here" (Lewis 1979, 520). Given these illustrations, I take a **centered world** to be the combination of a possible world and a perspective within it—the **center**.

Centered worlds are used to capture self-locating content, such as the content of Perry's belief *I am leaving a sugar trail*. One rules out other ways oneself might have been by locating oneself within a set of centered worlds. In the same way that possible worlds

<sup>5</sup>I am being purposely vague about whether this is Lewis's diagnosis. Plausibly, the official account in (Lewis 1979) does not need centered worlds at all, but only properties and self-ascriptions of properties. Notably, Perry (1979) offers a distinct diagnosis of the problem. Lewis (1979) comments on the differences between his and Perry's diagnoses and argues that his is preferable.

<sup>6</sup>On Lewis's modal realism, a possible world is a maximal collection of spatiotemporally-related individuals; see (Lewis 1983a) and (Lewis 1986). Although it is presentationally convenient to speak as if modal realism were true, this paper does not presuppose any particular metaphysical account of possible worlds.

illuminate the notion of *de dicto* possibility, centered worlds illuminate the notion of *de se* possibility, or the notion of how oneself might have been.

There is an important asymmetry: while centered worlds can also illuminate *de dicto* possibility, possible worlds cannot illuminate *de se* possibility. To show that centered worlds can also illuminate *de dicto* possibility, consider how they are able to capture world-locating contents. World-locating contents are really just a special kind of self-locating contents: they locate oneself as a member of a world. To rule out that one is in a possible world where pigs fly is simply to rule out that oneself is a possible individual who is part of a world where pigs fly. Hence, every *de dicto* proposition has a *de se* equivalent.<sup>7</sup> But the converse is not true. As the problem of essential indexicals shows, some *de se* propositions have no *de dicto* equivalents. So possible worlds cannot by themselves illuminate *de se* possibility. On the Lewisian diagnosis, the problem of essential indexicals arises because of the irreducibility of self-locating contents to world-locating contents, or the irreducibility of the *de se* to the *de dicto*.

[end of p. 298]

### 1.3 Recent Applications

Carving the logical space finer than possible worlds has proven useful for philosophical theorizing. The implications of self-locating contents have been further investigated and centered worlds have gained novel applications in various sub-disciplines. Adam Elga has drawn out implications of self-locating belief in decision theory, and his sleeping beauty puzzle has generated an important literature for decision theorists to grapple with (Elga 2000, 2004). In metaphysics, Andy Egan advocates understanding secondary qualities and Sydney Shoemaker's "appearance properties" as self-locating contents (Egan 2006a,b). Centered worlds are important theoretical entities in the relativist theory that Egan and others defend in philosophy of language, metaethics, and aesthetics (Egan et al. 2005; Egan 2009, 2010). They are also important theoretical entities in a competing relativist theory, Berit Brogaard's perspectivalism, which has applications in philosophy of perception and epistemology (Brogaard 2008, 2010). Finally, David Chalmers uses centered worlds as a type of foundational entity in his two-dimensional semantics (Chalmers 2006).

These various uses of centered worlds fall into two broad categories: epistemic uses, in which centered worlds are used to illuminate *de se* epistemic possibilities, and metaphysical uses, in which centered worlds are used to illuminate *de se* metaphysical possibilities. While I cannot give a complete taxonomy, let me roughly sketch the divide between these two theoretical roles that centered worlds play. Epistemic uses of centered worlds model contents of speech and thought. Most existing uses that involve *de se* content fall into this category. Metaphysical uses of centered worlds draw out metaphysical distinctions. Egan's use of centered worlds to capture the notion of secondary qualities and appearance properties exemplifies this category. Given this divide, it is conceivable that centered worlds are suitable representational devices for

[end of p. 299]

<sup>7</sup>Most philosophers have assented to this claim, originally made in (Lewis 1979). However, Nolan (2006) objects that centered worlds cannot capture contents of selfless desires. Turner (2010) offers one response to Nolan's objection.

*de se* metaphysical possibilities but not for *de se* epistemic possibilities, and vice versa. For simplicity, however, for the rest of this paper I assume that both epistemic and metaphysical uses involve the same set of representational devices, and note where delineating the two uses may have ramifications.

Let me further motivate the project of investigating what centered worlds are through a comparison to investigations of what possible worlds are. It is true that philosophers can make good use of possible worlds without settling on what they are. Similarly, I concede that philosophers can make good use of centered worlds while setting aside the question of what they are. However, examining what possible worlds are brings out commitments of theoretical frameworks which make use of possible worlds that might have been previously obscured, including potential internal inconsistencies. Similarly, examining what centered worlds are holds the promise of bringing out previously-obscured commitments of theoretical frameworks that make use of them. For example, it is worth considering whether the various applications of centered worlds mentioned earlier indeed involve the same theoretical entities. It would be an interesting discovery if it turned out that *de se* epistemic possibilities and *de se* metaphysical possibilities need to be captured by two different sets of entities. As is the case with possible worlds, although substitutes may be found, discovering what substitutes are available is nevertheless a part of the investigation into what centered worlds are. Considering the diverse applications of centered worlds, answering the question of what they are has wide ramifications.

## 2 What Centers Are Not

After motivating an investigation into what centered worlds are, I now consider two existing accounts and why they are unsatisfactory. Although some metaphysical issues regarding centered worlds are simply inherited from possible worlds, the issue unique to centered worlds boils down to what centers are, or how possible individuals are individuated. Both accounts offered in the existing literature, the Quinean account and the Lewisian account, are inadequate because they face counterexamples. To put the problem simply: the existing set-theoretic apparatuses do not allow centered worlds to perform their intuitive role of picking out possible positions, situations, or predicaments. Although the counterexamples initially appear to be driven by exotic cases, they can be sufficiently generalized to reveal a deeper problem that affects other potential accounts of centered worlds.

### 2.1 What Are Centers?

At the most basic level, a centered world is the combination of a possible world and a center. This basic characterization is insufficient as an account of what centered worlds are; an account needs to specify what possible worlds are and what centers are. For example, the answers that Lewis apparently gives are that possible worlds are concrete entities on par with the actual and centers are time-individual pairs (Lewis 1979, 532–533). Given the notorious controversy surrounding what possible worlds

are, most philosophers who employ centered worlds leave open what possible worlds are. Their assumption, one made by Lewis himself, is that whatever possible worlds turn out to be, the notion of centered worlds can be adapted accordingly (Lewis 1979, 533). Following this assumption, the rest of this paper will ignore the metaphysical issues that centered worlds inherit from possible worlds. Instead, there is an issue unique to centered worlds: what are centers?

There is an easy answer to the question of what centers are, but it merely raises similar issues in different terms. The easy answer, suggested by the foregoing discussion, is that each center corresponds to a possible individual, or a privileged perspective within a possible world. While this is indeed an answer, and one that is surely right, it appears uninformative because it tells us little about what exactly centers are picking out. Are possible individuals instantaneous time-slices, or are they temporally extended? The question of what centers are, I claim, is intimately related to the question of how possible individuals are individuated.<sup>8</sup>

[end of p. 301]

These two questions are intimately related because centers select possible individuals, and selecting is a way of stipulating: what we use to select a possible individual implicitly stipulates how possible individuals are to be individuated.<sup>9</sup> If possible individuals were individuated differently than we thought, then our actual selections of them might depart from our intentions, and that is theoretically undesirable. But as theorists, it is up to us to stipulate how technical terms such as ‘possible individual’ are to be used. We have a way of doing so: by selecting them in a manner that we think gets our intentions right. Thus, how we philosophers in fact select possible individuals gives an implicit stipulation of how possible individuals are to be individuated. It is for this reason that the question of how possible individuals are individuated is intimately related to the question of what centers are. On this understanding, the easy answer is in fact hardly an answer. If no informative answers can be given to one question, then no informative answers can be given to the other, intimately-related question. While it is possible to take the individuation of possible individuals to be primitive—after all, everyone needs primitives in theorizing—it is the last resort after other plausible alternatives have been rejected.

Philosophers who employ centered worlds appear to think that there are other alternatives because they have proposed or adopted accounts of what centers are. Lewis takes Quine to give the answer that possible individuals are distinct regions of spacetime

---

<sup>8</sup>Lewis implicitly accepts this intimate relationship. For Lewis, sets of centered worlds and self-ascriptions of properties play the same theoretical role and are little more than notational variants. He makes this especially clear in (Lewis 1983a). On the Lewisian theory of properties, properties are simply sets of possible individuals. Hence, sets of centered worlds are, for all intents and purposes, self-ascriptions of sets of possible individuals. For they to play the same theoretical role, there must be a close connection between centers and possible individuals.

<sup>9</sup>If one is a universalist about composition, then any mereological sum of any two possible individual is also a possible individual. In which case, there will be no matter of fact about how possible individuals might be individuated without specifying some level of individuation that is of interest. The relevant level of individuation here is one that suffices for philosophers’ various uses of centered worlds. To restate this point more carefully: what we use to select a possible individual implicitly stipulates how the possible individuals *that we care about in our uses of centered worlds* are to be individuated. This interest-sensitive individuation of possible individuals is the sense that I will employ for the rest of this paper.

[end of p. 302] (Quine 1969; Lewis 1979, 531–532). On the **Quinean account**, each unique possible individual can be picked out by its spacetime coordinates. As I note earlier, Lewis himself appears to give the answer that centers are pairs of an individual and their temporal coordinate (Lewis 1979, 532). As most philosophers who follow Lewis understand it, the individuals involved in this specification are persisting persons (or, more generally, entities with intentional states), as opposed to (primitively-individuated) possible individuals.<sup>10</sup> On the **Lewisian account**, each unique possible individual can be picked out by the persisting person it is a (proper or improper) part of and an external temporal coordinate. I will examine these accounts in greater detail next. The important point for now is that philosophers who employ centered worlds appear to think that possible individuals can be picked out by some special features that they possess, and consequently, possible individuals are individuated by their differences with respect to a set of these special features.

(A note on terminology: I am reserving **features** as a folk, or non-technical, term for qualities that one has, such as *being tall* or *being red*, and reserving **properties** as the corresponding technical, philosophical notion. Although I will mostly use ‘feature’ and its cognates for the subsequent discussion, they may be mentally replaced with ‘property’ and its cognates for ease of understanding. The importance of keeping the two terms apart will become apparent in §3.1.)

## 2.2 The Quinean Account

[end of p. 303] Let us first examine the Quinean account, which answers that possible individuals are distinct spacetime regions of different possible worlds, and accordingly, are picked out by their spacetime coordinates  $(x, y, z, t)$ . In other words, on this account, one’s location in logical space is picked out by one’s physical location in a possible world. To see how to translate talk of mental and linguistic self-locating contents into talk of this technical notion, consider Perry’s messy-shopper example again. Let us stipulate that when Perry came to believe that he was the messy shopper, Perry was located at spacetime coordinate  $(a, b, c, d)$  in the actual world @. When he has the self-locating belief *I am the messy shopper*, the content of his belief thus corresponds to a set of centered worlds including  $(@, (a, b, c, d))$ . More concisely, the self-locating belief **selects** the centered world  $(@, (a, b, c, d))$ .

The Quinean account faces problems when location in logical space comes apart from physical location in a possible world. Lewis raises this problem in a brief parenthetical remark, but it is worthwhile to spell the problem out in more detail.<sup>11</sup> Consider a co-location variation on Perry’s messy-shopper case. The case is as described before. Let us again stipulate that when Perry came to believe that he was the messy shopper, Perry was located at spacetime coordinate  $(a, b, c, d)$  in the actual world @.

<sup>10</sup>For example, (Egan 2006a, footnote 34) takes the choice between the Quinean account and the Lewisian account to be a mere matter of convention. On Egan’s understanding, since the Quinean account does not involve primitive possible individuals, neither does the Lewisian account.

<sup>11</sup>Lewis remarks, “(Here I assume that one centered world cannot be centered on two different cats, cats who occupy the same place at the same time. To avoid that assumption, as perhaps we should, we might redefine centered worlds as pairs of a world and a designated inhabitant thereof.)” (Lewis 1979, 532).

Unbeknownst to Perry, and thus missing from his description of the original case, was the fact that there was a conscious ghost, Ghost, located at the same exact spacetime coordinate (a, b, c, d) in the actual world @. Ghost was not the messy shopper. While this scenario is exotic, it is nevertheless a possibility.<sup>12</sup>

The co-location messy-shopper case shows that the Quinean account cannot adequately capture the talk of mental and linguistic self-locating contents. Intuitively, Perry's self-locating belief *I am the messy shopper* should select messy-shopper centered worlds. Since Perry is a messy shopper, the centered world associated with Perry should be amongst those that are selected by Perry's self-locating belief. In contrast, since Ghost is *not* a messy shopper, the centered world associated with Ghost should *not* be amongst those that are selected by Perry's self-locating belief. The problem is that, on the Quinean account, Ghost and Perry are associated with the same centered world ((@, (a, b, c, d)). The set-theoretic apparatus provided by the Quinean account thus leaves no room for Perry's self-locating belief to be true of him but false of Ghost, as it is intuitively the case.

Simply put, Perry's self-locating belief is about *him*, and not about someone at his physical location. The space of centered possibilities allows for two distinct predicaments at the same place, same time, and same possible world. However, the representation afforded by the Quinean account does not adequately capture such a rich space of centered possibilities, and the co-location messy-shopper case forcefully brings out this problem. Therefore, this account is an unsatisfactory answer to what centers are, or how possible individuals are individuated.

### 2.3 The Lewisian Account (Standard Interpretation)

Next, let us examine the Lewisian account, which answers that possible individuals are, roughly, instantaneous slices of persisting persons. Accordingly, each center is a persisting person plus a temporal location (*i*, *t*). This account is "Lewisian" because it is frequently attributed to Lewis and adopted by philosophers who employ centered worlds.<sup>13</sup> §4 will consider whether Lewis actually endorses this account.

[end of p. 304]

On the **standard interpretation** of the Lewisian account, the *i* parameter corresponds to persisting persons and the *t* parameter corresponds to external, real, objective time. To see how to translate talk of mental and linguistic contents into talk of this technical notion, consider Perry's messy-shopper example again. Let us stipulate that when Perry came to believe that he was the messy shopper, the persisting person, John

<sup>12</sup>I am taking such a scenario to be both metaphysically and epistemically possible. If the reader does not think that the following case is a genuine metaphysical possibility, then consider my counterexample to be restricted only to the epistemic uses of centered worlds. Importantly, denying that co-location is metaphysically possible is itself a controversial and substantive theoretical commitment that users of centered worlds are unlikely to all agree to take on. The same disclaimer applies, *mutatis mutandis*, to the other counterexamples in this paper.

<sup>13</sup>For example, Stalnaker writes, "One influential answer to the question of content, defended by David Lewis, is this: If the contents of ordinary beliefs about objective facts can be represented by sets of possible worlds, then the contents of self-locating beliefs can be represented by sets of *centered* possible worlds, where a centered possible world is a pair consisting of a world plus a center, which is a designated time and person" (Stalnaker 2008, 49).

Perry, was located at temporal coordinate  $d$  in the actual world  $@$ . When he has the self-locating belief *I am the messy shopper*, the content of his belief thus corresponds to a set of centered worlds including  $(@, (\text{John Perry}, d))$ . More concisely, the self-locating belief selects the centered world  $(@, (\text{John Perry}, d))$ .

The Lewisian account faces problems when location in logical space comes apart from an individual plus a temporal location. Consider a time-travel variation on Perry's messy-shopper case. The case is as described before. Let us again stipulate that when Perry came to believe that he was the messy shopper, Perry was located at spacetime coordinate  $(a, b, c, d)$  in the actual world  $@$ . Unbeknownst to Perry, and thus missing from his description of the original case, was the fact that Perry is a time-traveler. When he was older, he traveled back to this moment,  $d$ , and was spatially located elsewhere; let us stipulate that he was located at spacetime coordinate  $(a', b', c', d)$ . For the sake of clarity, call the former Young-Perry and the latter Old-Perry. Old-Perry was not the messy shopper. While this scenario is exotic, it is nevertheless a possibility.<sup>14</sup>

This time-travel messy-shopper case shows that the Lewisian account cannot adequately capture the talk of mental and linguistic self-locating contents. Intuitively, Young-Perry's self-locating belief *I am the messy shopper* should select messy-shopper centered worlds. Since Young-Perry is a messy shopper, the centered world associated with Young-Perry should be amongst those that are selected by Young-Perry's self-locating belief. In contrast, since Old-Perry is *not* a messy shopper, the centered world associated with Old-Perry should *not* be amongst those that are selected by Young-Perry's self-locating belief. The problem is that, on the Lewisian account, Young-Perry and Old-Perry are associated with the same centered world  $(@, (\text{John Perry}, d))$ . The set-theoretic apparatus provided by the Lewisian account, on the interpretation that individuals are persisting persons, leaves no room for Young-Perry's self-locating belief to be true of him but false of Old-Perry, as it is intuitively the case.

[end of p. 305]

Simply put, Young-Perry's self-locating belief is about *him-at-that-time*, and not about a future part of him that happens to be at the same external time. The space of centered possibilities allows for two distinct predicaments of the same persisting person located at the same external time. However, the representation afforded by the Lewisian account does not adequately capture such a rich space of centered possibilities, and the time-travel messy-shopper case forcefully brings out this problem. Therefore, this account is an unsatisfactory answer to what centers are, or how possible individuals are individuated.

## 2.4 The Lewisian Account (Other Interpretations)

Perhaps this rejection of the Lewisian account is too hasty. The standard interpretation assumes that the  $i$  parameter corresponds to persisting persons and that the  $t$  parameter corresponds to external, real, objective time. This assumption has frequently been adopted by philosophers who have employed centered worlds in their theories, and so even if other interpretations emerge from dropping this assumption, the problem raised in the previous section remains significant. Nevertheless, given that this paper

<sup>14</sup>See the disclaimer in footnote 12.

is about what centered worlds are, we should consider whether other interpretations of the Lewisian account afford better answers.

The **temporal parts interpretation** proposes to understand possible individuals as temporal parts of persisting persons.<sup>15</sup> Consider the following widely-accepted definition of temporal parts, as offered by Theodore Sider in his exposition of the ontology of temporal parts.

*x* is an **instantaneous temporal part** of *y* at instant *t* =<sub>df</sub> (1) *x* exists at *t*, but only at, *t*; (2) *x* is part of *y* at *t*; and (3) *x* overlaps at *t* everything that is part of *y* at *t*. (Sider 2001, 59)

Let us return to the time-travel messy-shopper case and apply this definition. What is the persisting person John Perry's temporal part at the instant *d*? It is not Young-Perry because it does not overlap at *d* everything that is part of John Perry at *d*. Clause (iii) of Sider's definition is not satisfied because, in this case, Young-Perry does not overlap Old-Perry, which is also a part of John Perry at *d*. For analogous reasons, Old-Perry is not the temporal part of John Perry at *d* either. Instead, the temporal part of John Perry, applying Sider's definition of temporal part, is the mereological sum of Young-Perry and Old-Perry. This mereological sum is clearly not the relevant possible individual in question. This interpretation has the same problem as the standard interpretation: it leaves no room for Young-Perry's self-locating belief to be true of him but false of Old-Perry, as it is intuitively the case. Young-Perry's self-locating belief *I am the messy shopper* is about *him-at-that-time*, and not about some strange mereological sum of his younger and older selves. A straightforward application of the definition of temporal parts in the time-travel messy-shopper case shows that possible individuals cannot be temporal parts of persisting persons.

[end of p. 306]

The **personal time interpretation** keeps the *i* parameter as corresponding to persisting persons, but proposes that the *t* parameter corresponds to *personal time* rather than external time. Personal time is not a distinct dimension of time. In fact, strictly speaking it is not time, but "that which occupies a certain role in the pattern of events" (Lewis 1976a, 147). Considering the time-travel case described earlier helps to make the distinction between personal time and external time apparent. People's lives tend to follow a pattern of events that involve change to many properties, such as the accumulation of memories and the loss of hair. Old-Perry exhibits features that are typical of a later stage in the patterns of events that characterize people's lives. It has more memories and less hair than Young-Perry. The changes that typify the patterns of events allow us to truly say that although Old-Perry is external-time simultaneous with Young-Perry, it is personal-time later than Young-Perry. We can assign time-like coordinates, or personal-time coordinates, to the different stages of John Perry's life because the stages exhibit different stages of regularities that are typical of persons. On this interpretation, since Young-Perry and Old-Perry are assigned different personal-time coordinates, and

<sup>15</sup>At face value, Lewis might be read as favoring this interpretation in (Lewis 1983a, 29): "The subject's alternatives will typically be possible people, or subjects rather like people; or better, they will be temporal stages thereof." As we will see, however, person stages are not always identical to instantaneous temporal parts of persons.

are thus associated with distinct centered worlds, the time-travel messy shopper case no longer presents a problem for the Lewisian account.

[end of p. 307] Yet, the invocation of personal time has its own problems. The assignment of personal-time coordinates, by the definition of personal time, depends on there being some regularities or patterns of events on which changes occur. When a possible individual is in a world with no changes, there can be no informative assignment of personal time. Consider the following scenario. The external time of the world extends infinitely in both directions, and there is only one person in this world. This person persists through time with the same sensation of pain at all instants. While there are intuitively many different stages of this person at different instances, no assignment of personal-time coordinates can pick out any stage out uniquely. Thus, the personal time interpretation does not work for all possible scenarios either.

Finally, consider the **primitive stage interpretation**. Rather than bringing in instantaneous temporal parts or personal time, the proposal is to take the  $i$  parameter to correspond to person stages.<sup>16</sup> As the problems with the previous two interpretations show, there is an intuitive sense of person stages on which they correspond to neither instantaneous temporal parts of persisting persons or personal times of persisting persons. So let us take person stage as a primitive.

While this interpretation does not face the problems that the previous two face, it faces other worries. First, how possible individuals are individuated remains mysterious without an account of how person stages are individuated. It is unclear what theoretical advantage this proposal offers over a proposal that simply takes centers, and therefore the individuation of possible individuals, as primitive. In fact, this proposal may incur additional ontological costs. On the surface, at least, invoking person stages involves a commitment to the ontology of temporal parts. This additional theoretical commitment is unlikely to be one that all users of centered worlds would be willing to take on. Second, and more importantly, this interpretation does not seem to preserve the spirit of the Lewisian account because it makes parameter  $t$  otiose. Given that each person stage, by their nature, can be only located at one instant in external time (and one instant in personal time), there is no need to specify the external (or personal) temporal location once we have specified which person stage we are selecting. This second worry makes it especially unlikely that the primitive stage interpretation is what the philosophers who have endorsed the Lewisian account has in mind.

[end of p. 308] Instead, the primitive stage interpretation comes close to the earlier proposal of treating centers, and therefore the individuation of possible individuals, as primitive. What considering this interpretation shows is that perhaps such a proposal has more advantages than one might initially think, insofar as it avoids the problems that plague the Quinean account and the other interpretations of the Lewisian account. Nevertheless, as I note earlier, such an account is a last resort that I am going to save for, well, last. For now, it suffices to reiterate that the Lewisian account, on interpretations that preserves its spirit, remains an unsatisfactory answer to what centers are, or how

---

<sup>16</sup>This interpretation is inspired by the stage view of persistence that Sider (1996, 2001) develops and argues for. Sider's stage view of persistence, in turn, takes inspiration from Lewis's account of personal identity in (Lewis 1976b).

possible individuals are individuated, just like the Quinean account.

## 2.5 The Problem Behind the Counterexamples

The counterexamples above reveal a deeper problem with identifying the centers of centered worlds. Each messy-shopper variation case presents a possible scenario where the account considered delivers a counterintuitive result. As the co-location messy-shopper case shows, when taking a center to be specified by an ordered set of space-time coordinates, there is a possible scenario where two possible individuals share those features. As the time-travel messy-shopper case shows, when taking a center to be specified by an ordered set of a persisting person and a temporal location, there is a possible scenario where two possible individuals share those features. Both cases have the same underlying structure: they both present a possible scenario where two possible individuals share all the features that an account uses to select possible individuals.

The recognition of this common structure, in turn, suggests a recipe for generating counterexamples. The accounts considered so far only mention in their set-theoretic apparatuses a subset of all features that a possible possible individual could have. So, their set-theoretic apparatuses cannot distinguish two possible individuals who differ only with respect to one of the unmentioned features. As long as an account of centered worlds employs a set-theoretic apparatus that leaves out a feature that a possible individual could have, we can construct a scenario where two possible individuals share all features mentioned by its set-theoretic apparatus but differ with respect to an unmentioned feature.<sup>17</sup> The appropriate abundance of possible worlds guarantees such a scenario is possible.<sup>18</sup> There is thus a reliable way to get such an account to deliver an unintuitive result: present a possible scenario where the account mistakenly associates two possible individuals with the same centered world.

[end of p. 309]

To illustrate how to apply this recipe, let us briefly consider another account of centered worlds. Looking at the counterexamples earlier, one might notice something interesting. On the one hand, in the co-location messy-shopper case where the Quinean account delivers a counterintuitive result, the Lewisian account delivers the intuitively correct result. On the other hand, in the time-travel messy-shopper case where the Lewisian account delivers a counterintuitive result, the Quinean account delivers the intuitively correct result. A natural suggestion, then is to put the two accounts together.

<sup>17</sup>The claim needs to be qualified to only scope over a maximal domain of *modally independent* features. For example, *being red* and *being colored* are not modally independent features because the former necessitates the latter. A possible individual could not be red but be not colored. Two features can also fail to be modally independent of one another in more complex ways. For example, *being red* and *being blue* are not independent because they are exclusive determinates of the same determinable, *being colored*. A possible individual could not both be red and be blue. The notion of modal independence that I am employing is akin to what Sider (2005) calls “non-overlapping” and what Saucedo (2009) calls “determinably-distinct”. A domain of *basic* features, whatever these turn out to be, is an example of a maximal domain of modally independent features. This qualification is implicit for the rest of this paper.

<sup>18</sup>There is an appropriate abundance of possible worlds when logical space includes all possible ways of recombining modally independent features. Although formulating a precise recombination principle is notoriously difficult, the basic idea of recombination suffices for the purpose here.

On the **combined account**, each center is a combination of a persisting person and a physical location in a possible world, written as  $(i, x, y, z, t)$ .

Unfortunately, there is a counterexample to the combined account too; the recipe above shows us how to generate it. The goal here is to find a possible scenario where two possible individuals are parts of the same persisting person and in the same physical location. The previous variations on Perry's messy-shopper case suggest one: a scenario that includes a time-traveling, co-locating conscious ghost. In this possible scenario, there are two possible individuals, call them Young-Ghost and Old-Ghost, who are parts of the same persisting conscious ghost and in the same physical location. The combined account mistakenly associates Young-Ghost and Old-Ghost with the same centered world. The set-theoretic apparatus provided by the combined account leaves no room for Young-Ghost's self-locating belief to be true of him but false of Old-Ghost. Therefore, the combined account is also an unsatisfactory answer to what centers are. The foregoing illustration shows that the recipe above is powerful because it can be used to rule out many potential accounts of centered worlds, namely accounts that simply add more (but not all possible) parameters to their set-theoretical apparatuses.

[end of p. 310]

### 3 What Centers Might Be

A series of counterexamples have shown existing answers to the question of what centered worlds are to be unsatisfactory. Moreover, the recipe developed in §2.5 shows that the problem behind the counterexamples is a general one. I now suggest two novel accounts of what centered worlds are—neither of which are explicitly considered in the existing literature—that evade this problem. One account identifies a center with an exhaustive ordered set of features. The other account treats the identification of centers as primitive. Considering the costs and benefits of each account brings out their respective theoretical commitments.

#### 3.1 Exhaustive Set

On the **exhaustive set account**, the way out is to let each center to be picked out by an exhaustive ordered set of features.<sup>19</sup> The ordered set of features being exhaustive guarantees that there cannot be a possible scenario where two individuals share all of the mentioned features but differ on an unmentioned feature. Since the ordered set is exhaustive, there cannot be any other feature to differ on. So no counterexample against this account can be generated using the recipe developed in §2.5. On the exhaustive set account, possible individuals are individuated by all the features they could have. If two possible individual differ on any feature, then they are distinct.

Despite its initial attractions, this account faces a considerable problem when we begin to consider what features are. My earlier discussion is purposefully conducted in terms of features, the non-technical notion, in order to set aside this question. But this question can no longer be avoided. The most natural answer is the folk notion of

<sup>19</sup>An exhaustive ordered set of features needs to only include all *modally independent* features that a possible individual could have. See the qualification in footnote 17.

features correspond to philosophers' notion of properties. However, on one popular view of what properties are—a view advanced by Lewis—features cannot be properties if they are what individuates possible individuals.

Lewis advocates building properties out of sets of possible individuals. At the most basic level, he advocates an unstructured conception of properties. A property is simply the set of all its instances. Not the set of all its *actual* instances, of course, but the set of all its *possible* instances. Hence, properties are sets of possible individuals (Lewis 1986, 55). As with other loaded theoretical terms in philosophy, there are multiple conceptions associated with the term 'property'. Lewis recognizes this and notes another available conception of properties, more closely tied to the meaning of their names. He grants that this alternative structured conception of properties may have stronger expressive power in certain cases. Nevertheless, he argues that structured properties can be built from unstructured properties and higher-order relations between those properties (Lewis 1986, 56). Even on this structured conception, then, properties are still given in terms of sets and possible individuals. Therefore, on the view of properties that Lewis advances, one that is commonly accepted by contemporary philosophers, possible individuals are metaphysically *prior* to properties: what properties are depends on what possible individuals are.

[end of p. 311]

The exhaustive set account is incompatible with this *Lewisian view of properties*. Recall that on the exhaustive set account, possible individuals are individuated by all the features they could have. If features correspond to properties, then the doctrine entails that possible individuals are individuated by all the properties they could have. Since possible individuals are given in terms of sets and properties, properties are metaphysically *prior* to possible individuals: what possible individuals are depends on what properties are. The order of metaphysical priority that the exhaustive set account demands is contrary to the order that the Lewisian view of properties demands. Even if one finds the notion of metaphysical priority somewhat mysterious, the circularity between what properties are and what possible individuals are remains apparent. The incompatibility gives those who are sympathetic to the Lewisian view of properties a reason to reject the exhaustive set account.

Although there exist possible patches, they all involve controversial and substantive theoretical commitments that users of centered worlds are unlikely to all agree to take on. Someone who is sympathetic to both the exhaustive set account and the Lewisian view of properties could argue that even if the circularity involved in the notions of properties and possible individuals makes it impossible to reduce one to the other, articulating the relationship between the two is nevertheless philosophically illuminating. Alternatively, she might reject the claim that features correspond to properties and offer an alternative technical notion that better corresponds to the folk notion of features. Of course, a defender of only the exhaustive set account could also reject the Lewisian view of properties and endorse instead a view of properties that is not metaphysically dependent on possible individuals. In the end, which of these patches is most preferable depends on one's other theoretical commitments. Indeed, whether the exhaustive set account itself is ultimately preferable depends on one's other theoretical commitments.

[end of p. 312]

### 3.2 Primitive Identification

On the **primitive identification account**, the way out is to let each center to be picked out by a cheap “identity” property. On the Lewisian view of properties, since every set of possible individuals is a property, trivially there is a property that corresponds to every possible individual (or more precisely, every unit set of an individual).<sup>20</sup> Call such properties “*identity*” properties. Trivially, since every possible individual has its own unique “identity” property, no two possible individuals can have the same “identity” property. So no counterexample against this account can be generated using the recipe developed in §2.5.

As the scare quotes around ‘identity’ suggest, the primitive identification account tells us nothing about the identity conditions of individuals or how they are individuated. “Identity” properties are cheap because every individual gets one, regardless of how they are individuated. On this account, the identification of centers and the individuation of possible individuals are primitive, unable to be elucidated through other non-trivial features. We have now come to the last resort. Let me offer two defenses of why the last resort may not be as bad as it seems.

The first defense is that, despite the theoretical awkwardness, the primitive identification account can preserve philosophers’ existing uses of centered worlds in their theories. On this account’s formal terminology, each center is simply a possible individual (*i*). All we need to select a possible individual is the *i* parameter, corresponding to the “identity” properties. Nevertheless, it is harmless, and perhaps even convenient, to add additional parameters to our specification of centers, such as the time parameter *t*. While these parameters are, strictly speaking, otiose, they offer a rough characterization and a heuristical indication of the possible individuals we intend to refer to.<sup>21</sup> For example, even though earlier examples show that possible individuals cannot always

[end of p. 313]

be picked out using spacetime coordinates or the combination of a persisting person and an external temporal coordinate, these non-trivial features that possible individuals possess can still serve as tools of convenience in talking about the intended possible individuals.

The second defense is that although there remains something mysterious about what centers are or how possible individuals are individuated—despite the rough characterizations available—this mystery is to be expected given the main lesson from the problem of essential indexicals: the *de se* cannot be reduced to the *de dicto*. There is something special about learning who *oneself* is that cannot be captured in learning about what features one possesses, even if that list of features is exhaustive. There seems to be a fundamental conceptual distinction between ascribing properties to *oneself* and ascribing properties to an individual possessing a unique and exhaustive list of non-trivial properties. Hence, the mysteriousness involved in the primitive identification

<sup>20</sup>These are what Lewis calls strict haecceities: “A unit set of an individual is one especially strict sort of haecceity” (Lewis 1986, 225).

<sup>21</sup>Egan (2009) proposes a similar idea in response to an early ancestor of this paper. A difference between Egan’s proposal and mine is that Egan takes *predicaments* that persisting individuals might be in, rather than possible individuals themselves, to be primitive.

account is in fact necessary to respect the main lesson from the problem of essential indexicals.

In the end, while I find the primitive identification account more attractive, given my sympathies with the Lewisian view of properties, I acknowledge that someone with different theoretical commitments could easily find the exhaustive set account more attractive. The foregoing discussion articulates the costs and benefits of each account, and reveal the interactions that exist between the debate about centered worlds and metaphysical debates elsewhere, such as the debate about properties. Acknowledging these interactions helps to bring out the implicit theoretical commitments of frameworks that make use of centered worlds.

## 4 What Centers Are, for Lewis

A historical curiosity remains: what is Lewis's real position on centers?

It seems that the primitive identification account is most compatible with Lewis's overall metaphysical picture. Consider his other theoretical commitments: he advocates building properties from sets of possible individuals, he believes in the existence of "identity" properties, and he takes the problem of essential indexicals seriously. While he may not fully endorse what I call the Lewisian diagnosis—his official account in (Lewis 1979) involves only properties and self-ascriptions of properties, not centered worlds—he sees no additional problems with positing centered worlds and formulating the solution to the problem of essential indexicals in those terms.<sup>22</sup> The primitive identification account is compatible with all these theoretical commitments, and evades the problems that plague other accounts.

[end of p. 314]

Lewis's statements elsewhere further suggest that his real position on centers is the primitive identification account, and not the Lewisian account as it is standardly understood. In *On the Plurality of Worlds*, he remarks that it is person stages, and not persons, who are bearers of attitudes:<sup>23</sup>

The same person can have different systems of belief at different times. Suppose it is true, as I think it is, that a person persists through time by consisting of many different momentary stages located at different times. [...] Then we can say first that the various stages have various systems of belief; and then that the continuing person has a system of belief at a time

<sup>22</sup>In fact, he explicitly says that "centered worlds amount to presentations of possible individuals, so a proposal to use centered worlds differs little from my proposal to use the individuals themselves" (Lewis 1983a, footnote 18). If one accepts Lewis's assumption that possible individuals are world-bound, then even the world parameter  $w$  is otiose; centered worlds themselves can be identified with possible individuals.

<sup>23</sup>However, Lewis makes another modification later: "I have been speaking as if the assignment of content were an assignment directly to a given subject. But I would rather say that the content belongs to some state—a brain state, perhaps—that recurs in many subjects" (Lewis 1986, 39). Lewis makes this modification in response to a concern with subjects whose behaviors fail to fit their brain states. Since this concern falls outside of the scope of this paper, we can reasonably set this modification aside.

by having a stage at that time which has that system of belief. (Lewis 1986, 29)

Since contents of attitudes are properties, and properties are just sets of possible individuals, contents of attitudes correspond to sets of person stages. Hence, possible person stages play the role of possible individuals for Lewis.<sup>24</sup>

[end of p. 315] Importantly, as §2.4 shows, person stages are not extensionally equivalent to instantaneous temporal parts of persisting persons. In his writings on personal identity, Lewis makes it clear that person stages are metaphysically *prior* to persons: the individuation of persons depends on the individuation of person stages. In “The Paradoxes of Time Travel”, he remarks that “A time traveler, like anyone else, is a streak through the manifold of space-time, a whole composed of stages located at various times and places” (Lewis 1976a, 146). He repeats similar remarks in “Survival and Identity” (1976b).<sup>25</sup> Nowhere does he attempt to define how person stages are to be individuated. Indeed, what Lewis does not say is perhaps even more indicative of his position on centered worlds than what he does say. Considering that Lewis is such a systematic philosopher and that he never explicitly gives an account of the individuation of possible individuals, [end of p. 316] it seems likely that he had the primitive identification account in mind all along.

---

<sup>24</sup>I thank an anonymous referee for emphasizing this point and referring me to the textual evidence available in Lewis’s other writings, especially those on personal identity.

<sup>25</sup>In the postscripts to “Survival and Identity”, Lewis hesitates to commit to the claim that person stages are more basic than persons: “When I say that persons are maximal R-interrelated aggregates of person-stages, I do *not* claim to be reducing ‘constructs’ to ‘more basic entities’. [...] Whatever ‘more basic’ is supposed to mean, I don’t think it means ‘smaller’” (Lewis 1983b, 77). However, it seems that his hesitation is primarily due to the notion of basicness, which he finds puzzling, and not due to the nature of persons and person stages.

## References

- Brogaard, B. (2008). In defence of a perspectival semantics for ‘know’. *Australasian Journal of Philosophy*, 86(3):439–459.
- Brogaard, B. (2010). Perspectival truth and color primitivism. In Wright, C. D. and Pedersen, N. J. L. L., editors, *New Waves in Truth*, chapter 16, pages 249–266. Palgrave Macmillan, New York.
- Chalmers, D. J. (2006). The foundations of two-dimensional semantics. In Garcia-Carpintero, M. and Macia, J., editors, *Two-Dimensional Semantics*, chapter 4, pages 55–140. Oxford University Press, New York.
- Egan, A. (2006a). Appearance properties? *Nous*, 40(3):495–521.
- Egan, A. (2006b). Secondary qualities and self-location. *Philosophy and Phenomenological Research*, 72(1):97–119.
- Egan, A. (2009). Billboards, bombs and shotgun weddings. *Synthese*, 166:251–279.
- Egan, A. (2010). Disputing about taste. In Feldman, R. and Warfield, T. A., editors, *Disagreement*, chapter 10, pages 247–292. Oxford University Press, New York.
- Egan, A., Weatherson, B., and Hawthorne, J. (2005). Epistemic modals in context. In Preyer, G. and Peter, G., editors, *Contextualism in Philosophy*, chapter 6, pages 131–168. Oxford University Press, New York.
- Elga, A. (2000). Self-locating belief and the sleeping beauty problem. *Analysis*, 60(2):143–147.
- Elga, A. (2004). Defeating Dr. Evil with self-locating belief. *Philosophy and Phenomenological Research*, 69(2):383–396.
- Lewis, D. K. (1976a). The paradoxes of time travel. *American Philosophical Quarterly*, 13:145–152.
- Lewis, D. K. (1976b). Survival and identity. In Rorty, A. O., editor, *The Identities of Persons*, chapter 1, pages 17–40. University of California Press, Berkeley.
- Lewis, D. K. (1979). Attitudes *de dicto* and *de se*. *The Philosophical Review*, 88(4):513–543.
- Lewis, D. K. (1983a). Individuation by acquaintance and by stipulation. *The Philosophical Review*, 92(1):3–32.
- Lewis, D. K. (1983b). Postscripts to “Survival and identity”. In *Philosophical Papers*, volume 1, pages 73–77. Oxford University Press, New York.
- Lewis, D. K. (1986). *On the Plurality of Worlds*. Blackwell, Malden, MA.

- Nolan, D. (2006). Selfless desires. *Philosophy and Phenomenological Research*, 73(3):665–679.
- Perry, J. (1979). The problem of the essential indexical. *Nous*, 13(1):3–21.
- Quine, W. V. O. (1969). Propositional objects. In *Ontological Relativity, and Other Essays*, chapter 6, pages 139–160. Columbia University Press, New York.
- Saucedo, R. (2009). *Things, Places, and Possibilities: An Inquiry Into Structural Deviations*. PhD thesis, Cornell University, Ithaca, NY.
- Sider, T. (1996). All the world's a stage. *Australasian Journal of Philosophy*, 74(3):433–453.
- Sider, T. (2001). *Four-Dimensionalism*. Clarendon Press, New York.
- Sider, T. (2005). Another look at Armstrong's combinatorialism. *Nous*, 39(4):679–695.
- Stalnaker, R. C. (1981). Indexical belief. *Synthese*, 49:129–151.
- Stalnaker, R. C. (2008). *Our Knowledge of the Internal World*. Oxford University Press, New York.
- Turner, J. (2010). Fitting attitudes *de dicto* and *de se*. *Nous*, 44(1):1–9.