Copyright

by

Bryce Daniel Dalbey

2020

The Dissertation Committee for Bryce Daniel Dalbey
certifies that this is the approved version of the following dissertation:

How to be Naïve about the Mind

Committee:

Mark Sainsbury, Supervisor

Ray Buchanan

Josh Dever

Michelle Montague

Adam Pautz

Jeff Speaks

Michael Tye

How to be Naïve about the Mind

by

Bryce Daniel Dalbey

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

August 2020

How to be Naïve about the Mind

Publication No. _____

Bryce Daniel Dalbey, Ph.D.
The University of Texas at Austin, 2020

Supervisor: Mark Sainsbury

This dissertation defends **common-sense views of the mind** by combating two widespread tendencies among philosophers. One such tendency is to eliminate: to resolve puzzles about the mind by denying the existence of ordinary mental features. For instance, many think that while we can experience the shape, color, and texture of a baseball, we cannot experience the time it takes for a baseball to fall to the ground or the number of times it bounces — indeed, we cannot experience any temporal features. The other tendency is to inflate: to resolve puzzles about the mind by positing new and unusual mental features. For instance, it is almost universally accepted among philosophers that to allow for the rationality of agents, especially those like Lois Lane and Oedipus, we must posit *guises* (or *senses*, or *modes of presentation*) under which agents think.

Against these tendencies, I argue that we can resolve puzzles about the mind without invoking new features or denying ordinary ones. In chapter one

I confront 'Frege puzzles' concerning Lois Lane and show that there several distinct yet often-conflated issues at play. Moreover, the plausibility of such puzzles depends on an equivocation between them. Once disentangled, it is clear there are simple explanations of Lois' rationality that do not employ *guises*.

In chapter two I confront the Knowledge Argument, which aims to establish that Mary the color scientist learns a non-physical fact upon seeing red for the first time, and by extension that the mind is not physical. The most popular responses to this argument invoke special mental features, including so-called *phenomenal concepts*, *knowledge by acquaintance*, and certain *mental abilities*. I argue for a simple response to the argument which does not invoke any special mental features. On the simple response, Mary is simply misled into thinking she's learned something when she has not.

In chapter three I confront a puzzle about temporal experience that many take to suggest we do not experience temporal features. I argue that experiencing is a process rather than a state (more like *running* than like *being tall*) and that this distinction resolves the puzzle: we experience temporal features over periods of time but not in virtue of experiencing them at instants during that time (just as one runs over periods of time but not in virtue of running at instants during that time).

Table of Contents

Chapter 1

Russellianism and Rationality without Guises

In theorizing about the mind, philosophers often invoke a notion of content. This is perhaps most common in discussions of the so-called propositional attitudes: states like hoping that the weather cools off, fearing that the end is nigh, and supposing that it rains later. Indeed, it is the received view that having a propositional attitude at least partly consists in taking a mental attitude (e.g. hope, wonder, judgment) to some content.[1] Content is in turn understood as something wholly specified by — and, according to many, identical to — an arrangement of certain items.[2] Of course, this received view still leaves much to be settled: are propositional attitudes constituted by anything more than taking an attitude to content? what sort of items figure into the arrangements that specify content?

A particularly natural answer to the former question (to which many subscribe) is negative: propositional attitudes are wholly constituted by taking an attitude to content. Call this Austerity. A particularly natural answer to

---

[1] 'Attitude' is, unfortunately, used both for what one takes to content and for the state of taking some such things to content. It will be clear from context which is meant.

[2] This is often elaborated on by reifying contents as abstract entities which have the relevant items as constituents, then understanding attitudes as relations one bears to such abstract items. Nothing in this paper turns on whether that elaboration is correct.

the latter question (to which many also subscribe) is that content is specified by arrangements of whatever items the propositional attitude is about, and hence generally by arrangements of ordinary worldly items (e.g. ordinary objects, events, properties, relations, etc.).[3] Call this Russellianism.

The conjunction of these two claims — **Austere Russellianism** — thus follows naturally from basic questions about propositional attitudes. It also helps us make substantial progress towards many of the theoretical goals at which theories of mind aim, including a naturalization of the mind, a model of psychological explanation, a simple semantic theory, an ordinary ontology, and an account of our epistemic access to the world.

Despite this, Austere Russellianism is widely rejected.[4] This is often done on the basis of what I'll call **intelligibility objections**, many of which originate with Frege. Such objections take a variety of forms but are generally to the effect that Austere Russellianism cannot allow for the rationality or cognitive explicability of certain individuals, among them the infamous Lois Lane, Oedipus, ancient Babylonians, and Catiline. This is because, it is claimed, propositional attitudes can vary in the rational or explanatory work they do despite not varying in attitude or what they are about.

---

[3]One can also get a grip on what a propositional attitude is about as what it concerns, what the content's truth-makers are, or (in a more semantic mode) what it specifies, denotes, or picks out.

[4]Austere Russellians include Crawford (2004), Fodor (1994), and Thau (2002), though Fodor later abandoned the view. Soames comes close but at (2006, p. 722) appears to reject it. Cappelen and Dever (2018, Ch. 1) expresses sympathy for a nearby view. The semantic analog of Austere Russellianism is, however, fiercely defended. See fn. 6.

To allow propositional attitudes to do such rational or explanatory work, philosophers standardly posit so-called **guises** (sometimes called 'senses,' 'modes of presentation,' or 'ways of thinking') as aspects of propositional attitudes. These philosophers either take the content of propositional attitudes to be specified solely by arrangements of guises (thereby rejecting Russellianism) or take guises to be additional elements of such mental states beyond taking an attitude to content (thereby rejecting Austerity). Indeed, the invocation of guises is so widespread in contemporary philosophy that their existence is often taken as a datum rather than a theoretical posit (though it is worth remembering that even Frege took the existence of senses as a posit).[5]

The reports of Austere Russellianism's death have been greatly exaggerated. Guises are not needed to make rational or explanatory sense of agents. In fact, there is a perfectly plausible Austere Russellian account of rationality and cognitive explanation which evades intelligibility objections, which thereby undercuts much of the theoretical motivation for invoking guises.

In §1.1 I present and taxonomize intelligibility objections and relate them to another well-known form of objection. In §1.2 to §1.4 I present Austere Russellian responses to the most prominent forms of intelligibility objections. In §1.5 I generalize those responses to a systematic Austere Russellian account of agents' rational and explanatory status.

_____

[5]In fact, he originally went without them — see (Frege 1892, pp. 56–7).

1.1 Austere Russellianism and Its Critics

Austere Russellianism is a thesis about the structure of propositional attitudes. It is related to but distinct from a more popular view about the semantic values of attitude ascriptions, namely that they just express attitude relations to Russellian contents.[6] Additionally, Austere Russellianism is neutral on what realizes or gives rise to propositional attitudes, be it relations to internal representations, having informational files, being in functional states, or something else. But on Austere Russellianism, if there are such items, mere differences in them do not make for differences in propositional attitudes. That is, if two agents have propositional attitudes which differ in how they are realized but not in attitude or what they are about then those agents have the same propositional attitude. On Austere Russellianism, mere differences in (say) the internal representations that give rise to propositional attitudes are washed out at the mental level, just as mere differences in the number of neurons that give rise to propositional attitudes are washed out at the mental level.

One might object to Austere Russellianism merely on these grounds, claiming that independent issues of rationality or cognitive explicability it is clear that agents at least sometimes differ in which propositional attitudes they have despite their attitudes being the same with respect to attitude and

---

[6] Defenders and sympathizers of this semantic view include Båve (2008), Braun (1998, 2002), Braun and Saul (2002), Crawford (2004), Fodor (1990, 1994), Frances (1998), Kaplan (1989), McKay (1981), Millikan (1993), Salmon (1986), Saul (2007), Schneider (2005), Sider and Braun (2006), Soames (1988, 2002), Thau (2002), and Tye (1978).

what they are about.[7] Call these **difference objections**. If they succeed then one might reasonably take intelligibility objections to be superfluous. As we will see, however, not only are the two forms of objection closely related but the plausibility of difference objections depends on claims about intelligibility.

1.1.1 The Road from Difference to Intelligibility Objections

Both difference and intelligibility objections are perhaps best presented as concerning cases of a specific sort. The details of such cases vary widely and in theoretically important ways but for now let us focus on the familiar case of Lois Lane. Recall Lois' situation:

> When asked to list superheros, Lois says the name 'Superman' but not the name 'Clark.' When asked to name people who are not superheros she says 'Clark' but not 'Superman.' When asked 'Is Superman the same person as Clark?' she says 'no.' When Superman is in front of her in his cape and tights and someone says 'can you point to Superman?' she points to him. When Superman is in front of her in his cape and tights and someone says 'can you point to Clark?' she does not point to him.

Both difference and intelligibility objections begin with an alleged datum about cases of this sort. They go on to contend, on various grounds, that if Austere

---

[7]Another influential form of objection to Austere Russellianism concerns thought about non-existents. See [REDACTED] for an Austere Russellian response to those objections.

Russellianism is true then that datum is false — that Austere Russellianism does not recognize this difference in attitudes.

Difference objections generally begin with an alleged datum about the presence or absence of propositional attitudes. One such objection, couched in terms of Lois' case, begins with the following: it is at least sometimes true that

(Sup)   Lois believes that Superman is a superhero

and also false that

(Cla)   Lois believes that Clark is a superhero.[8]

Call this a **difference datum**.

Austere Russellianism alone is not inconsistent with the difference datum. One way to put the two in contact (which takes a cue from analogous objections to Austere Russellianism's semantic analog) invokes certain semantic principles, most prominently Millianism, the thesis that the semantic contribution of a name is just its referent.[9] Another way to connect the two bypasses semantic principles and employs, among others, the claim that the

---

[8]'Sometimes' encodes variation in, among other things, the context of evaluation.

[9]Other standardly-invoked semantic principles include i) that-clauses designate contents, ii) names contribute their ordinary semantic values in that-clauses, iii) the semantic value of an expression is determined by the semantic values of its atomic constituents and their place in its syntax, and iv) (Sup) and (Cla) do not differ with respect to their syntax or the semantic value of any atomic constituents beyond certain substituted expressions. v) the names 'Superman' and 'Clark' co-refer.

relevant propositional attitudes are about exactly the same things. Of course, the Austere Russellian can attempt to block difference objections by rejecting some such claims — by, say, arguing that names semantically contribute descriptions rather than referents — but I will grant them. Austere Russellianism together with such supplementary claims deliver that (Sup) is the case just in case (Cla) is the case. This is incompatible with the difference datum.

Of course, (Sup) and (Cla) are at least sometimes both true of Lois as described: suppose that you and I desperately need to locate a superhero to save the city and, unlike Lois, we are privy to the fact that 'Superman' and 'Clark' co-refer; I hear Lois say 'Superman is a superhero' and I immediately report back to you by saying 'we may have found one! Lois thinks Clark is a superhero!' It is not just that we speak this way on rare occasion — this way of speaking is utterly ubiquitous.[10] Given this, one might begin to suspect that (Cla) is in fact true exactly when (Sup) is — it's just that sometimes (Cla) is misleading to say or an otherwise poor choice of words given what one's interlocutor will come to believe as a result.

To more firmly establish the difference datum in light of this suspicion, objectors might appeal to supporting claims about Lois' case. Related appeals in discussions of the semantic analog of Austere Russellianism generally fall into one of two categories: those which concern various (generally intentional

---

[10]Analogues of the difference datum clearly fail for other mental state attributions, including those of the so-called objectual attitudes: Lois hates Superman just in case she hates Clark.

or behavioral) features of Lois and those which concern various (intentional or behavioral) features of speakers using (Sup) and (Cla). For instance, some argue for the difference datum by appeal to the fact that Lois assents to 'Superman is a superhero' but dissents from 'Clark is a superhero.' Others argue for the difference datum by appeal to the fact that informed, competent speakers of English sometimes assent to (Sup) but dissent to (Cla).[11]

In response to such arguments for the difference datum the Austere Russellian can either reject the supporting claim or reconcile it with their view.[12] Supporting claims like those above — which concern what certain people do — are easily reconciled with Austere Russellianism. For instance, it might be that what such agents do is irrational or cognitively inexplicable. Given this, the supporting claims should be couched in a way that does not allow for such reconciliations. A natural way to do this is for the supporting claims to be of the form: so-and-so does such-and-such *rationally* or *explicably*. Of course, the notion of explicability cannot be of just any sort. Austere Russellianism is a comparatively restrictive view of propositional attitudes but it is not similarly restrictive about other sorts of features. So it is only for forms of explanation that employ propositional attitudes for which Austere Russellianism is at a relative disadvantage. One prominent form of explanation

---

[11]See Crawford (2004) for an overview of such appeals.

[12]Claims about competent speakers appear to generate problems analogous to those generated by claims about Lois herself: Lois misidentifies Superman and competent speakers misidentify the content of (Sup). Given this, one would expect the responses to each to be similarly analogous. Some, e.g. Braun (1998), Braun and Saul (2002), and Saul (2007), obey this expectation while others, e.g. Salmon (1986) and Soames (1988, 2002), appear to flaunt it.

that fits the bill is cognitive explanation. Another is a form of rationalization, a kind of explanation that requires mental facts as explanans — mere neural, chemical, or even computational features cannot rationalize any more than features of electromagnetic fields unless they are themselves mental.

At this stage it becomes clear that the work in these objections is mainly being done by claims about intelligibility. Assenting to (Sup) or dissenting to (Cla) are just some among many intentional features evaluable for intelligibility. Indeed, many philosophers object to Austere Russellianism on the basis of claims about intelligibility that are independent of such linguistic issues. They argue that no matter how such semantic facts turn out, Austere Russellians cannot account for facts about the intelligibility of what agents like Lois do or do not do. This transition delivers a second variety of objections.

1.1.2 Intelligibility Objections

Despite their influence, intelligibility objections are often underdeveloped and the relations among them are not always appreciated. To see the full range of intelligibility objections it will be helpful to have a more detailed scenario than the one above. Consider the following continuation of that case:

> One day, Lois arrives at the Daily Planet and says 'does anyone know where there are any superheroes? I need to find one.' Everyone shakes their heads. For unrelated reasons, someone says to her 'by the way, Clark is at the courthouse all day if you need to get ahold of him.' Lois responds 'Thanks. And let me know if

you hear of the whereabouts of any superheroes.' Later, someone says to Lois 'I just heard that Superman is at the courthouse.' She responds 'finally — thanks!' and heads to the courthouse.

Plausibly, throughout the case, Lois believes that Superman is a superhero and believes that Clark is not a superhero. Moverover, before she hears 'Superman is at the courthouse' Lois believes that Clark is at the courthouse and does not believe that a superhero is at the courthouse. Then after hearing 'Superman is at the courthouse' Lois comes to believe that a superhero is at the courthouse. As above, it is available to the Austere Russellian to block various intelligibility objections by rejecting some of these claims but I will grant them here.[13] (For ease of discussion I will assume that co-referring names are intersubstitutable in attitude ascriptions and that descriptions in attitude contexts are non-referential. So, I will assume that Lois believing that Superman is a superhero suffices for Lois believing that Clark is a superhero but does not suffice for Lois believing that the savior of Metropolis is a superhero.)

Among the most common intelligibility objections are those that employ alleged data of the following sort (couched in terms of Lois' case):

(ExpNot)   there is a cognitive explanation of the fact that Lois does not believe that a superhero is at the courthouse before she hears

---

[13]A promising strategy of this sort is to reject that Lois does not believe that a superhero is at the courthouse. In fact, beliefs of that sort are frequently attributed to agents in situations like Lois' (cf. Hawthorne and Manley (2012)). That being said, I will set aside this strategy.

'Superman is at the courthouse'.[14]

(ExpCome)   there is a cognitive explanation of the fact that Lois comes to believe that a superhero is at the courthouse after she hears 'Superman is at the courthouse'.[15]

(NoPosBel)   Lois is not in a position to rationally believe that a superhero is at the courthouse before she hears 'Superman is at the courthouse'.[16]

(NoPosKno)   Lois is not in a position to know that her belief that Superman is a superhero and her belief that Clark is a reporter are about the same person.[17]

(RatCont)   Lois rationally believes that Superman is a superhero and Clark is not a superhero.[18]

I'll call claims of this sort **intelligibility data**. Each corresponding objection then proceeds, on various grounds, that if Austere Russellianism is true the

---

[14]Advocates of nearby objections include Arjo (1996), Aydede (1997, 1998), Aydede and Robbins (2001), Braun (2004), Devitt (1996), Heck (1995, 2002), and Richard (1990). Classic discussions of surrounding issues include Block (1986) and Loar (1988).

[15]Heck (2012) presents an objection in this spirit.

[16]Claims of this sort and the two following are found widely in the literature on Russellianism but seldom developed at length (though see Heck (2012)).

[17]Cf. related discussions on reconciling forms of content externalism with varieties of self-knowledge.

[18]A related objection concerns multiple beliefs with contradictory contents. Related constraints are often imposed on theories of guises, such as the Intuitive Criterion of Difference (Evans 1982, p. 18) and Frege's Constraint (Schiffer 1990, p. 252). See also discussions surrounding Kripke (1979).

relevant claim is false.[19]

These intelligibility data form a heterogeneous lot. Fortunately, they fall into a general taxonomy of available intelligibility data. They are all alike in being assessments of Lois having some sort of feature. They differ in the kind of assessment and the kind of feature, both of which vary along several dimensions. Understanding these dimensions will help make clear distinctions among seemingly equivalent objections as well as similarities among seemingly unrelated objections.

The features that figure into intelligibility data are generally intentional since intentional features are rationally evaluable and admit of cognitive explanation. Such features vary in at least four relevant ways. One dimension of variation governs what kind of intentional state is being evaluated. Candidates include belief, action, intention, and judgment.

A second governs whether the feature is, as we might say, positive (e.g. believing, acting) rather than negative (e.g. not believing, not acting). (B) and (RatCont) above are of the former sort while the other three are of the latter.

A third dimension of variation governs whether the feature is a state

---

[19]It is worth noting that a variety of extant theories of guises allow for cases that generate the same problems for them that Lois' case generates for Austere Russellianism. That is, these theories allow for cases that are structurally like Lois' and such that the agent's propositional attitudes involve the same guises. For many such theories, Paderewski cases are of this sort (Kripke 1979). In allowing for this possibility, these views are in largely the same position as Austere Russellianism with respect to intelligibility objections.

(e.g. believing, not believing) or a change (e.g. coming to believe, ceasing to believe). At least (A) and (RatCont) are of the former sort while at least (B) is of the latter.

A fourth dimension of variation governs the content of the intentional features. Some, like (NoPosKno), concern metacognitive content. Others might concern metalinguistic content, such as that 'Superman' and 'Clark' refer to the same thing (though none on the list above are of this sort). All but (NoPosKno) concern contents not about language or mind.[20]

The kinds of assessment that enter into intelligibility objections vary in at least four relevant ways as well. One dimension of variation governs whether the assessment originates from a normative domain (e.g. rationality and rationalization) rather than from a scientific domain (e.g. cognitive explanation and cognitive explicability - cognitive explicability stands to cognitive explanation as rationality stands to rationalization). (NoPosKno) and (RatCont) are of the former while the rest are of the latter.[21]

A second dimension of variation governs whether the assessment is couched in terms of explanation (e.g. cognitive explanation, rationalization) rather than in non-explanatory terms (e.g. rationality, cognitive explicability). Only the former quantify over explanations. (A) and (B) are of the former while the rest are of the latter.

---

[20]Even if the objects of other intentional states are not contents of this sort, a similar distinction will apply, e.g. actions directed at one's mind or directed at pieces of language.

[21]Being in a position to infer, at least in the relevant discussions, is akin to being in a position to rationally infer.

A third concerns whether the assessment is positive (e.g. being rational) rather than negative (e.g. not being rational). (A), (B), and (RatCont) are of the former sort while (NoPosBel) and (NoPosKno) are of the latter.

A fourth dimension of variation in kinds of assessment requires more exposition. The intelligibility data that concern rationality come in two forms. Some, like (RatCont), are assessments of features that require the agent to have those features: rationally believing that p requires that one believe that p. Others, like (NoPosBel), are assessments of features that do not require the agent to have those features: being in a position to rationally believe that p does not require that one believe that p.[22] Plausibly being in a position to rationally believe that p is a condition on rationally believing that p. A natural way to understand the former is as rational permission to believe or rational possibility of believing.[23] A way to understand the latter is as one believing that p in an appropriate way (e.g. in a way that appropriately takes advantage of one's permission to believe that p). As a result, rationality assessments vary with respect to whether they are assessments concerning one's rational permission to believe rather than assessments of one believing appropriately.

There is an analogue of this distinction for cognitive explicability. One can intelligibly believe that p, which requires believing that p. And one can be cognitively able to believe that p, which does not require believing that

---

[22]These differ in something like the way so-called propositional justification differs from so-called doxastic justification.

[23]This also reveals a kind of rational necessity or obligation. As these can be understood in terms of rational possibility or permission, I will set them aside.

p. The latter is a kind of cognitive possibility. The former is believing that *p* appropriately. As a result, the fourth dimension of variation in assessment governs whether the assessment concerns a feature being (rationally or cognitively) possible for one rather than one having a feature in a (rationally or cognitively) appropriate way.

These distinctions generate a multitude of intelligibility objections, not all of which are worth considering separately here. For reasons of length, I will restrict my focus to intelligibility objections that concern Lois and belief that is not about language or the mind. In discussing each, I am mainly concerned with presenting an account of rationality and cognitive explicability that either accepts the corresponding intelligibility datum or which provides an explanation of it being false.

One difference in intelligibility data that results from some of these distinctions is that between those like (A), (B), (NoPosBel), and (NoPosKno) which concern (cognitive or rational) unavailablility and those like (RatCont) which concern (cognitive or rational) availability. These generate importantly different problems for Austere Russellianism. I begin with the former, confronting objections that employ claims like (A) and (B) in §1.2 and confronting those that employ claims like (NoPosBel) and (NoPosKno) in §1.3. Then, in §1.4, I turn to objections that employ claims like (RatCont).

1.2 "Why didn't she...?"

I begin with objections that concern the rational assessment of Lois coming to believe that a superhero is at the courthouse and not believing that a superhero is at the courthouse.

1.2.1 Coming to Believe and Not Believing

The first intelligibility objections I'll confront are those which employ the claim (concerning Lois once she hears 'Superman is at the courthouse'):

(RatCome)   There is a rationalization of Lois coming to believe that a super-
             hero is at the courthouse.[24]

To see the force of these objections, consider that, plausibly, had Lois not heard 'Clark is at the courthouse,' rationalization of her coming to believe that a superhero is at the courthouse would be her coming to believe that Superman is at the courthouse. Yet on Austere Russellianism, Lois did not come to believe that Superman is at the courthouse when she heard 'Superman is at the courthouse,' as she already came to believe this when she heard 'Clark is at the courthouse.' This suggests that Austere Russellianism precludes the apparently ordinary rationalization — not merely because of the time between Lois coming to believe that Superman is at the courthouse and her coming to believe that a superhero is at the courthouse but because it appears that

---

[24]Our discussion will make clear a response to objections which turn on the claim that Lois rationally comes to believe that a superhero is at the courthouse.

whatever mental state Lois forms when she hears 'Clark is at the courthouse' plays no rational role in her later cognitive change.

Rationalization, on this objection, is a form of explanation that employs mental explanans. Rationalization of one coming to believe that p does not, on this objection, concern the basis upon which one believes that p — that upon which one bases one's belief that p. Even on Austere Russellianism, Lois might base her belief upon something she learned quite a bit earlier. Rather, rationalization of one coming to believe that p concerns the changes that prompt one to believe that p. Call this **prompting rationalization**.

On Austere Russellianism there is a prompting rationalization of Lois coming to believe that a superhero is at the courthouse. Perhaps the easiest way to see it is by first considering related cases. The focus of much of the literature on intelligibility objections has been on cases involving the use of two names, as with Lois'. But crucially, intelligibility objections do not require the use of names. In fact, they do not even require that the agent acquires information linguistically. (Nor do they require that the relevant item be a concrete object (e.g. Superman) rather than, say, a feature.) To get a wider view of the phenomenon, consider the following case:

> After a long vacation, Henry goes to the parking garage to find his car. Having forgotten exactly where he parked it, he wanders around searching. At one point he sees it in the back corner but because of the poor lighting and odd angle he does not recognize it and continues searching. Eventually he circles around and sees

it again, this time from the other side and in better lighting. He then heads towards his car, relieved to have found it.

Call Henry's car Carl. Plausibly, Henry comes to have various beliefs about Carl in a way that matches the structure of Lois' case: he believes that Carl is his car (and has since he bought it, though he would not put it this way), that Carl is not his car (since seeing it on his first trip around the garage), and that Carl is in the corner (again, since seeing it the first time around); he initially (the first time around) does not believe that his car is in the corner; later (the second time around) he comes to believe this.

The analog of the present objections for Henry employ the claim that there is a rationalization of Henry coming to believe that his car is the the back corner. Like for Lois, the relevant rationalization is not that Henry came to believe that Carl is in the back corner nor is it that he came be believe that Carl is his car — neither of those are what prompts him to believe. Any proposed rationalizer which itself requires rationalization regenerates the problem. So there must be some ultimate rationalizer — something which rationalizes but does not itself require rationalization — available.

Indeed, there is a perfectly natural rationalizer of Henry coming to believe that his car is in the corner. We might quibble about the details but it is clearly related to the visual state he is in the second time around the garage. A particularly natural ultimate rationalizer is that he saw Carl well-

18

illuminated.[25] This is a mental features that can rationalize Henry coming to believe, does not itself require rationalization, and which (even on Austere Russellianism) Henry only comes to have the second time around the garage.

Of course there are other, non-rationalizing explanations of Henry coming to believe that his car is in the corner. If certain computational theories of cognition are true then some of these explanations concern the way his cognitive system encodes various pieces of information. This is a realizational explanation — an explanation of his coming to believe in terms of him coming to have something that realizes this. The Austere Russellian need not reject such explanations. But they are not rationalizations and do not compete with the rationalization just given.[26] Likewise, there are explanations of why Henry has the proposed rationalizing mental feature. Since this rationalizing mental feature does not itself require rationalization, such explanations are not rationalizations and may concern (say) neural realizers without threat of undermining Austere Russellianism.

Turning back to Lois, the rationalization of her coming to believe that a superhero is at the courthouse is not as obvious as it was for Henry. Yet

---

[25]Other candidates include his perceptually experiencing Carl to be well-illuminated and it seeming to him that his car is in the corner as the result of his visual state.

[26]One gets the feeling that an equivocation between rationalizing and realizational explanations is partly responsible for the widespread invocation of internal representations and the like as guises. One asks: why did Henry come to believe that the car was in the corner? A natural interpretation of this question is a request for a rationalization (e.g. he experienced it to have certain features). Another natural interpretation of the question is a request for a realizational explanation (e.g. his cognitive system came to encode the information in the right way). Not clearly distinguishing the two, one offers the answer that concerns computation while taking the question to concern rationalization.

there is an analog. The rationalization of Henry coming to believe is related to his visual state. The rationalization of Lois coming to believe is not since vision does not play an important role in her case. Instead, it is related to her auditory state. The rationalization of Henry coming to believe is related to him seeing Carl. The rationalization of Lois coming to believe is not related to her hearing Superman. Rather, is related to her hearing something about Superman.[27] In other respects, the rationalization of Lois coming to believe is like the rationalization of Henry coming to believe. A natural analog is that she was told that Superman is at the courthouse by the use of the name 'Superman.'[28] This is a mental feature that can rationalize, does not require rationalization, and which (even on Austere Russellianism) Lois only comes to have when she hears 'Superman is at the courthouse'. It does not require that Lois think about which names are used, just as the rationalization we gave of Henry coming to believe did not require him to think about the lighting.

One might respond to this proposal that the invocation of names heard demands the invocation of guises — demands that the propositional attitudes thereby formed are not exhausted by their attitude and what they are about. But there is no such demand. The proposal merely specifies the details of Lois' perceptual state. Differences in the names one hears (or the lighting conditions in which one sees) often leads to differences in what the resultant

---

[27]Henry's case might have gone similarly: he sees a poorly-lit picture of Carl and then later a well-lit picture of Carl.

[28]Other candidates include that she auditorily experiences the use of the name 'Superman' in being told that Superman is at the courthouse and that it seems to her that a superhero is at the courthouse as the result of her auditory state.

propositional attitudes are about. They also often lead to differences in neural or computational features. But names playing this role does not demand that there are additional items which constitute the propositional attitudes Lois formed, just as differences in lighting do not demand that there are additional items which constitute the propositional attitudes Henry forms. Seeing in one sort of lighting can all by itself rationally explain why one forms certain beliefs and not others. Likewise for hearing something by the use of one name rather than another.

With the above rationalization in hand we can also respond on behalf of Austere Russellianism to the intelligibility objections which employ the claim (concerning Lois before she hears 'Superman is at the courthouse'):

(RatNot)   There is a rationalization of Lois not believing that a superhero is at the courthouse.[29]

To see the Austere Russellian response to such objections, consider the rationalization of Henry at first not believing that his car is in the corner. The form of rationalization at play here is not prompting rationalization as it is not rationalization of a change. Rather, it is rationalization of a failure to change. A natural rationalization is that Henry only saw Carl poorly-illuminated. This is a mental feature that can rationalize him not coming to believe that his car is in the corner. For Lois, an analogous rationalization is that she was only

---

[29]Our discussion will make clear a response to objections which turn on the claim that Lois rationally does not believe that a superhero is at the courthouse.

told that Superman is at the courthouse by the use of the name 'Clark.' This is not a name she is unused to (compare Henry: he is not used to finding Carl in poor lighting) but it is a name she is unused to using when trying to find out about superheroes.

We can isolate at least three upshots from the preceding discussion. First, the resources available for rationalization are much wider than one might have thought — one need not only look at ordinary beliefs. Second, the rationalizations of agents like Lois and Henry not believing need not be of the same form. Third, in asking "Why didn't she do such-and-such?" one should be careful not to equivocate over the notion of explanation at play. There is certainly an explanatory role to be played by internal realizers of propositional attitudes, but such explanations needn't come in when giving rationalizations.

1.2.2 Rationalization and Having Sufficient Reason

There is a worry which I suspect underlies many philosophers' motivation for pressing the above arguments. It is that the putative rationalizations just proposed are in fact no rationalization at all because they do not concern Lois' reasons to believe that a superhero is at the courthouse. One antecedently might have thought that if there is a rationalization of one not believing that p then one lacks **sufficient reason** to believe that p. Or, put in terms of rationality rather than rationalization, one antecedently might have thought

(SuffReas)   If one has sufficient reason to believe that p and one does not

believe that p then one irrationally does not believe that p.[30]

Plausibly, given Austere Russellianism and the above proposed rationaliza-
tions, when Lois hears 'Superman is at the courthouse' she does not gain a
reason to believe that a superhero is at the courthouse — the proposed ratio-
nalizations do not specify one nor make salient what one would be. Since Lois
has sufficient reason to believe that a superhero is at the courthouse after she
hears 'Superman is at the courthouse' she must have sufficient reason before
hearing this as well. So by (SuffReas) Lois irrationally does not believe that
a superhero is at the courthouse and hence there is no rationalization of her
not believing.[31]

Fortunately for the Austere Russellian, there are clear counter-examples
to (SuffReas). Consider the following case:

> Hannah craves a salty treat — a pretzel, some salted peanuts, even
> popcorn would do — but believing that there is nowhere nearby
> for her to get one she turns her attention to other matters. Later,
> she has a conversation about the gentrification of the surrounding
> neighborhood and it is mentioned that there are now three bars, a
> pizzeria, and a pretzel shop within walking distance. But Hannah

---

[30]Cf. Broome (1999, 2013). Related issues arise in debates surrounding the uniqueness
thesis.

[31]The objection from Graeme Forbes discussed at Braun (2000, fn. 36) has a somewhat
similar flavor to this objection. I think it is susceptible to a response similar to the one
below.

does not then go get a pretzel. Later, she suddenly exclaims 'wait a second, I could have gotten a salty treat this whole time!' and leaves to get one.

After the conversation Hannah believes that pretzels are salty treats and that there are pretzels around the corner. It might be very important to Hannah that she get a salty treat and she might have nothing better to do. She is not a particularly absent-minded person — this sort of thing happens to us all. Yet she does not (say) believe that there are salty treats around the corner, despite apparently having sufficient reason to do so. Hannah appears no less rational than Lois. Yet by (SuffReas) she irrationally does not believe that there are salty treats around the corner.

For another counterexample to (SuffReas), consider formal logician Jon who foolheartedly attempts to do a proof after a long and draining day. Jon knows that what he's trying to prove is of form Z, he knows that the claims he's already proved are of X and Y, and (being a trained logician) he knows that sentences of form X and Y jointly entail sentences of the form Z. Yet he does not come to believe that he can finish the proof, despite having sufficient reason to. He is not particularly irrational — he's simply too tired to think clearly. He just does not see it. Yet by (SuffReas) he irrationally does not believe that he can finish the proof.

An additional counter-example to (SuffReas) is Henry's case. Henry does not gain a reason to believe that his car is in the corner when he sees Carl the second time around. Since he has sufficient reason to believe that

24

his car is in the corner after he sees Carl the second time, he must have had sufficient reason before as well. So, by (SuffReas), Henry irrationally does not believe that his car is in the corner and hence there is no rationalization of him not believing so. But clearly there is such a rationalization.

Henry and Lois are in a position similar to that of Jon and Hannah. In fact, if anything Lois and Henry are better off with respect to rationality as they have better excuses for not believing: it is much more difficult for Lois and Henry to discover their errors than it is for Hannah and Jon to discover their own.

(Perhaps despite all this (SuffReas) is true and Hannah and Jon irrationally do not believe. One would not outright accuse Hannah or Jon of being irrational, of course, but one must recognize a difference between being irrational and it being appropriate and worth the effort to point out that one is irrational. And perhaps even the same is true of Henry. If so, it is little surprise that Lois is the same and no cost for Austere Russellianism to allow so.[32])

One might attempt to improve the present response to the proposed rationalizations by modifying (SuffReas) in a way that applies to Lois but not Hannah and Jon. There are of course many ways of distinguishing them — the task is to find a difference that generates a true principle. One might add to the antecedent of (SuffReas) that the agents consider the relevant claims

---

[32]Millikan (1993, esp. 289-90) makes a distinction between theories of rationality that is related to the distinction between those that accept (SuffReas) and those that deny it.

all at once. But such a principle will still apply to Jon as well as to a version of Hannah. One might add to the antecedent that one be in a position to rationally believe that p. But this also applies to Hannah and Jon, even more clearly than it does to Lois.

One might take a different tack and add to the antecedent that the agents be ideal in a certain respect: Hannah and Jon cannot be ideal given the details of the case while Lois could be ideal consistent with the details of her case. For instance, one could, so to speak, fix Hannah and Jon so that they do not get into situations of their sorts whereas one could not fix Lois so that she does not get into at least some situations of her sort. Given this, so goes the response, their rational status is different.

Set aside whether adding this to the antecedent generates a true principle. It is true that you could not get Lois to avoid all scenarios like the one in which she finds herself (at least without preventing her from doing anything at all). This is because you could not, so to speak, build a system that is perfectly sensitive to the sameness of features in the world. But likewise, you could not get Jon and Hannah to avoid all situations of their sort either. Jon and Hannah fail, as we might say, to put two and two together. That is, they fail to be appropriately sensitive to the relevant sameness of information they possess. And for the same reason you cannot build something that is perfectly sensitive to the sameness of external items, you cannot build something that is perfectly sensitive to the sameness of internal items.

One might instead object that there is no non-trivializing view of ra-

tionality on which (SuffReas) is false, so it must be true. In fact there is a quite defensible view which denies it, to which we now turn.

1.2.3 Rationalization and Cognitive Explanation

A natural view of rationalization on which (SuffReas) is false is one which assimilates it to notions of explanation in the special sciences. It will be easiest to understand this view by considering special science generalizations.[33] This has the added benefit of allowing us to respond to intelligibility objections that employ claims about cognitive explicability.

Consider the generalization

(Salty) If one wants a salty treat and believes that there are salty treats around the corner then *ceteris paribus* one will go around the corner.

This is a generalization of cognitive psychology. It is *ceteris paribus*, i.e. it allows for tolerable exceptions: cases which satisfy the antecedent but not the consequent and yet are consistent with the generalization.

Tolerable exceptions to generalizations like (Salty) are of two sorts. Those of one sort are still intelligible by the relevant science. For instance, someone who satisfies the antecedent of (Salty) but has more important things to do or believes that they cannot easily get around the corner will generally be a tolerable exception to (Salty). However, other generalizations from cognitive

---

[33]For proposals related to the one here, though which still avert to guises, see Braun (2000) and Schneider (2005).

27

psychology will capture such agents,including: if one wants a salty treat and believes that there are salty treats around the corner but believes that one cannot easily get around the corner then *ceteris paribus* one will not go around the corner.[34]

Tolerable exceptions of the other sort are those not capturable by the relevant science. Only generalizations of some other science — usually concerning realizing features — can capture such cases. Someone who satisfies the antecedent of (Salty) but has an aneurysm and does not go around the corner is this kind of tolerable exception. Only a principle invoking neurological features can capture such an agent. Tolerable exceptions of this kind are said to be unintelligible from the perspective of the original science and to exhibit a kind of pathology.[35]

In light of this distinction in tolerable exceptions, we can more informatively articulate (Salty) as

(Salty*) If one wants a salty treat and believes that there are salty treats
around the corner and one does not go around the corner then *ceteris*

---

[34]This generalization captures that particular tolerable exception to (Salty) but does so at the expense of decreased generality. It cannot explain the action of someone who satisfies its antecedent but decides to take the chance anyway, nor can it explain the actions of agents who do not believe that they cannot easily get around the corner. Such is the nature of special science explanation: increased accuracy brings decreased breadth.

[35]It is sometimes claimed that this kind of tolerable exception must be relatively rare. I disagree. Many true special science generalizations are in fact rarely followed. The ideal gas law is almost never obeyed exactly. The generalization 'if one flips a coin n times, the coin comes up heads $n/2$ times' is never followed when $n = 1$, only followed half the time when $n = 2$, never followed when $n = 3$, only followed three eights of the time when $n = 4$, etc.

paribus one unintelligibly does not go around the corner.

Tolerable exceptions to (Salty) that are capturable within cognitive psychology are also tolerable exceptions to (Salty*). Tolerable exceptions to (Salty) that are not capturable within cognitive psychology obey (Salty*) and exhibit unintelligibility.

(Salty*) is almost the same in form as (SuffReas). In light of this, a natural thought is that principles concerning rationality are akin to principles of the special sciences and that if there is any true principle in the vicinity of (SuffReas) it is *ceteris paribus*, like

(SuffReas*) If one has every reason to believe that p and one does not believe that p then *ceteris paribus* one irrationally does not believe that p.

A straightforward development of this thought is that rationalization and explanation in the special sciences are analogously structured, with the former involving principles like (SuffReas*) and the latter generalizations like (Salty*). On a view of this sort, the rationalization previously articulated of Lois not believing that a superhero is at the courthouse is a specification of what makes Lois a tolerable exception to (SuffReas*).

On such views, rationalization is closely related to cognitive and folk psychological explanation. This puts us in a position to confront additional

intelligibility objections, namely those that employ (A) and (B) from §1.1.2.[36] Folk cognitive psychology certainly recognizes the rationalizations that we've given for Henry as well as the analogues that we've given for Lois. So just as the rationalizations we've given for Lois exempt her from (SuffReas*), so too do they exempt her from an analog of (SuffReas*) couched in terms of cognitive explicability.

1.3   Being in a Position to Rationally Believe

We've so far seen responses to several intelligibility objections, all concerning (rational or cognitive) assessments of features Lois has. Those responses reconcile the relevant intelligibility datum with Austere Russellianism. This might suggest that Austere Russellian responses to intelligibility objections always take the form of reconciliations of the relevant intelligibility datum. Not so. At least some forms of intelligibility objections that turn on an agent being in a position to rationally or explicably believe employ alleged data that are false.[37] Consider, for instance, intelligibility objections that

---

[36]For responses to such objections consistent with Russellianism see Braun (2001) and Schneider (2005).

[37]There are a variety of surrounding notions in which related objections are sometimes couched: being rationally entitled to believe, being rationally permitted to believe, having reason to believe, being able to rationally believe, and being such that one could rationally believe. Though these notions differ, they fall into two relatively clear groups. Those of one group are kinds of rational allowance or entitlement (e.g. 'is rationally permitted to,' 'is rationally entitlement to,' 'has reason to'). Those of the other group are kinds of rational capacity (e.g. 'could rationally believe,' 'is able to rationally believe,' one reading of 'is in a position to rationally believe'). Plausibly, the former is necessary but not sufficient for the latter. The points made below at least apply to the former.

employ the claim

(NoPosBel) Lois is not in a position to rationally believe that a superhero is at the courthouse before she hears 'Superman is at the courthouse'.

It is not always clear what the evidence for (NoPosBel) is meant to be. One consideration given in its favor is that were Lois to come to believe that a superhero is at the courthouse at that time in her present state, she would not do so rationally. Of course, the truth of a counter-factual like this is not sufficient for claims like (NoPosBel). It might be that the basis upon which one would believe are different than those that put one in a position to rationally believe. To distinguish those cases from Lois', the counter-factual should specify the basis upon which one believes as well, as in

(WereBel−) If one is in a position to rationally believe that p on the basis of $q, r, s, ...$ then were one to believe that p on the basis of $q, r, s, ...$ in one's present state one would rationally believe that p.

Yet this is still not enough. Some agents are rationally entitled to believe that p on some basis and yet were they to believe it on that basis they would do so by employing a non-rationally entitling method.[38] To distinguish agents like these from Lois the counter-factual must also specify the method employed:

---

[38]Views of what counts as a rationally entitling method vary. Candidates include the employment of elementary valid forms of inference, reliable belief-forming faculties, and meaning-constituting inferential capacities.

(WereBel+)   If one is in a position to rationally believe that p on the basis of $q, r, s, ...$ then were one to believe that p on the basis of $q, r, s, ...$ in one's present state *and by employing a rationally entitling method* one would rationally believe that p.

Yet (WereBel+) also does not distinguish Lois from other entitled agents because Lois obeys the relevant counter-factual.

To see this, consider again the formal logician Jon. He also violates the counter-factual in (WereBel−) and obeys the counter-factual in (WereBel+). Given his beliefs about the form of the claims he's proved, the form of his hoped-for conclusion, and there being an entailment relation between claims of those forms, Jon is in a position to believe that he can complete the proof. Yet were he to believe that he can complete the proof in his present state (tiredness and all) even on the basis of those beliefs he would not rationally believe that p. Given how tired he is, if he were to form that belief on that basis (or any other) it would not be by using a rationally entitling method. Yet unlikely as it is, were he to believe that he can complete the proof in his present state on the relevant basis and by using a rationally entitling method, he would rationally believe that he can complete the proof.

Lois is just the same. Were she to believe that a superhero is at the courthouse on the basis of the relevant claims in her present state she would not use a rationally entitling method. Yet she has the information she needs, and unlikely as it is, were she to believe that a superhero is at the courthouse on the

basis of the relevant claims in her present state and do so via rationally entitling method, she would rationally believe that a superhero is at the courthouse. She will not do that, of course, and we should not expect her to. But she nevertheless is in a position to. We (incorrectly) thought that Lois is not in a position to rationally believe on the basis of (correctly) thinking that she wouldn't rationally believe it even if she did so on the right basis.

Like Jon, Lois falsifies the counter-factual in (WereBel−) but not the one in (WereBel+). Agents like Lois, just as those like Jon, have a difficult time rationally believing. And being epistemically responsible agents, they will not believe via a non-rationally entitling method. But they are still rationally entitled to believe. This is an unfortunate situation for them but not an unusual consequence of Austere Russellianism as situations of this sort are commonplace.

Analogous claims hold for intelligibility objections that turn on claims of being in a position to explicably believe that a superhero is at the courthouse. Lois has the information she needs, just as Jon does. But she will not form the belief, and like Jon (and Henry and Hannah), she has a perfectly good excuse for not doing so.

1.4 Contradictory Beliefs

To conclude our whirlwind tour of intelligibility objections, let us consider those that turn on claims such as that Lois is in a position to rationally believe that Superman is a superhero and Clark is not a superhero, or, more

strongly,

(RatCont)   Lois rationally believes that Superman is a superhero and Clark is not a superhero.

(Analogues of these concerning cognitive explicability are not normally given.) Intelligibility objections that employ (RatCont) are not often spelled out in detail but the objectors often allege a rational prohibition on having contradictory beliefs, or a weaker prohibition on contradictory beliefs in select circumstances that include Lois'. (One should tread carefully here. It may be a norm of rationality that one not have contradictory beliefs. But the move from this to the prohibition on rational belief is substantial step. Compare going from it being a norm of morality that one not break one's promises to it never being moral to break one's promises.)

Objections that turn on claims like (RatCont) generate importantly distinct issues from those generated by the objections we've already discussed. The issue raised by the intelligibility data discussed in §1.2 and §1.3 was that Austere Russellianism makes (cognitively or rationally) *available* certain intentional features that should be unavailable. The issue raised by (RatCont), on the other hand, is that Austere Russellianism makes (rationally) *unavailable* certain intentional features that should be available.

Austere Russellians can allow that (RatCont) is true. An initial way to see this is by noting that on the view of rationality proposed here, rationality and rationalization are closely related to cognitive explanation. Since contra-

dictory beliefs are cognitively explicable, we should expect that they admit of rationalization as well.

Another way to see why these objections fail is by noting that there are cases which both bear a striking resemblance to Lois' and plausibly concern agents that rationally hold logically incompatible beliefs. The preface paradox provides one. Agents in preface paradox cases apparently rationally believe, say, of each claim made in a book that it is true, of the claims made in the book that they are all the claims made in the book, and that not all claims made in the book are true. Indeed, such agents might even know that their beliefs are inconsistent and yet rationally maintain them.

One might worry that it is plausible that preface paradox agents believe rationally because they have attitudes concerning a sufficiently large number of claims. Since Lois has attitudes concerning only two claims, one might take preface paradox lessons to not apply to Lois. On the contrary, there are cases in which agents have a small number of logically inconsistent and yet plausibly rational beliefs. Gilbert Harman discusses cases in which an agent becomes aware of a conflict in their beliefs but because each is practically important and they do not have the time to resolve the inconsistency when they notice it, resolves to keep the inconsistency temporarily until it can be resolved.[39] (In fact, if anything Lois is in a better position with respect to rationality than those in many Harman-style and preface paradox cases since unlike those

---

[39]Harman (1986, pp. 15–6).

agents, Lois does not know that her beliefs are inconsistent and, indeed, rationally believes that they are not.) Other cases of agents with rational beliefs concerning a small number of logically inconsistent claims include those of certain logicians who deny true logical principles. For instance, it is rational for Vann McGee to accept the premises of his famous alleged counter-example to *modus ponens* while denying the conclusion, even if *modus ponens* is valid and his alleged counter-example is an instance of *modus ponens*.[40] Versions of Hannah and Jon also provide clear cases of this. Hannah might, on general inductive grounds, rationally come to believe that there are no pretzels around the corner and Jon might likewise rationally come to believe that his hoped-for conclusion does not follow from the premises. Henry provides another case: he might well rationally form the conjunctive belief that Carl is his car and Carl is not his car.

One feature of these cases is that the agent's reasons for holding the relevant beliefs originate from distinct sources. For instance, in preface paradox cases: reasons for believing of each claim made in the book that it is true originate from evidence concerning the subject-matter of the book; reasons for believing of the claims made in the book that they are all the claims made in the book originate from having read the book; and reasons for believing that not all claims made in the book are true originate from evidence of one's fallibility. Similarly with McGee and Henry. It is this detail that, at least in part, seems to motivate the rationality judgment. Interestingly, just this

---

[40]McGee (1985).

feature is present in Lois' case. Her reasons for believing that Superman is a superhero originate from quite a different source (e.g. from her seeing him fly around) than her reasons for believing that Clark is not a superhero (e.g. from her knowing that reporters are generally not superheros). Lois' case patterns with those that constitute a data point for theories of rationality. (Or perhaps despite all this McGee, Harman's agent, Hannah, Jon, Henry, and those in preface paradox cases do not rationally believe what they do. If so, it is no surprise that Lois is included among them and no cost for Austere Russellianism to allow so.)

1.5   Conclusion

The general view of rationality and cognitive explicability defended here is one on which Lois is very much like Hannah, Jon, and McGee. These agents are all in a position to rationally and explicably believe their respective claims: for Hannah, that there are salty treats around the corner; for Jon, that he can finish the proof; for McGee, the conclusion of the alleged counter-example to *modus ponens*. Yet these agents rationally and explicably do not form such beliefs as they have rationalizing excuses for not doing so. And these agents (or versions of them) rationally have contradictory beliefs, as they have independent reason to believe each of the jointly contradictory claims.

The account defended here can be generalized to intelligibility objections that concern metacognitive and metalinguistic beliefs. For instance, a similar account can be given regarding the fact that Lois is in a position to ra-

tionally believe that her beliefs are about the same thing, though she rationally does not believe so and indeed rationally believes that her beliefs are about different things. We can also generalize the present account to intelligibility objections that concern intentional features besides belief, such as action or knowledge. In fact, some of the intelligibility objections discussed above do not have plausible analogues for other mental features. For instance, analogues of objections that concern having contradictory beliefs do not obviously arise for action (as one cannot perform incompatible or impossible actions) nor for knowledge (as one cannot know something contradictory).

An upshot of our discussion is that the motivation for imposing guises is much like the motivation for invoking an epistemological given. It is tempting to think that when one is in a position to infer something on the basis of other claims, this relation is laid plain to one, and so some aspect of entertaining those claims guarantees that this is laid plain. But as with other subject-matters, such relations are never simply given — agents have to figure it out like they have to figure everything else out. Sometimes this is easy (e.g. when one is well-rested, when the lighting does not change, when the same names are used) and sometimes it is hard.

Chapter 2

A Simple View of Mary

The mere knowledge of a fact is pale; but when you come to realize your fact, it takes on color. It is all the difference between hearing of a man being stabbed to the heart, and seeing it done.

- Mark Twain

Many physicalists respond to the knowledge argument that there is no new fact or information that Mary acquires when she leaves her room — she only receives old information. A central task for such physicalists is to explain or explain away various facts about Mary that look quite unusual if she does not acquire any new information. I argue that there is a simple physicalist account of Mary that does not invoke any special mental features.

2.1   The Big Reveal

Since the case of Mary the color scientist is familiar to most, I will give it briefly.[1] Mary is the foremost expert on color and color experience. She eventually comes to possess all the physical information in her chosen areas of

---

[1]The case comes from Jackson (1982).

study. However, Mary has spent her entire life in a black and white room and she has only ever visually experienced things to be black, white, and shades of gray. One day she exits the room and sees a red apple. *How incredible!* she exclaims.

Many have been tempted by the thought that when she sees the apple, Mary receives some new information — information she did not have in the room — about color or color experience. This observation, or one like it, is employed in arguments against physicalism.[2] These arguments vary in a variety of ways but a particularly clear version goes as follows:

(P1)   All physical information about color is such that Mary has it in her room.

(P2)   Some information about color is such that Mary does not have it in her room.

(C1)   So, some information is not physical.

(P3)   If some information is not physical then physicalism is false.

(C2)   So, Physicalism is false.

This is what has come to be known as **the knowledge argument**. The **standard physicalist reply** to the knowledge argument is that while Mary

---

[2]A related but importantly distinct argument concerning Mary turns on questions of *a priori* deducability.

does receive some information upon seeing the apple, all the information she receives is information she already had in the room (so in this formulation, (P2) is false).

Those who endorse the standard reply generally don't rest there, however. They articulate complex theories of new mental features — e.g. knowledge by acquaintance, practical knowledge, exercising phenomenal concepts — that Mary comes to have when she sees the apple. It is worth pausing to see why this is done. The reason, I take it, is that there are certain apparent facts about Mary which suggest that (P2) is true and which, absent further explanation, look quite mysterious on the standard reply. Among the most widely discussed of such claims are that Mary comes to know something, that she comes to know a new fact, that she gains some knowledge, that she gains some propositional knowledge, that she learns something, that she learns something propositional, that she learns a new fact, that she makes a discovery, that she discovers something, that she discovers some information, and that she makes epistemic progress. Much of the literature on the knowledge argument focuses on which of these claims are true and which mental features best allow advocates of the standard reply to reconcile them with the falsity of (P2).

Though many of those claims are plausible they are not beyond question. There are, however, certain claims about Mary which are even more obvious and which, absent further explanation, still look mysterious on the standard reply. One of the most basic of such claims is that upon seeing the

apple, at least some of the information Mary receives *strikes* her as new (or, more or less, that some such information *seems* new to her). Indeed, in some respect it is unavoidable that upon seeing the apple, the information Mary receives strikes her as new. Even if Mary is cognitively ideal — e.g. has perfect working and long-term memory and flawless reasoning abilities — this will still be the case. Granted, there are things that would change Mary in this regard: one could, for instance, reconfigure her cognitive system to be like that of someone who has had many experiences of red. But holding fixed that Mary is ordinary in the ways allowed by the case, upon seeing the apple the information she receives will strike her as new.

In this respect Mary is quite unlike someone in ordinary circumstances who knows that p and then later is, say, told that p. Suppose one learns from the news that it will rain this evening and then is told by a coworker that it will rain this evening. In all likelihood, upon being told so by one's coworker this information will not strike one as new. Absent further explanation, the difference between such agents and Mary looks quite mysterious. Thus, it seems a basic task for the advocate of the standard reply is to explain:

(Strike)   upon seeing the apple, Mary receives old information that strikes her as new.

Absent such an explanation, the standard reply appears largely unmotivated.[3]

––––––––––––––––––––––––––

[3]To be sure, not even (Strike) is completely obvious. Certainly something strikes Mary as new but that *information* strikes her as new is not completely obvious. Still, I will assume

In fact, there is a simple explanation of (Strike) which does not require appeal to special mental features. Moreover, that explanation undermines our reasons for thinking that the more controversial claims taken to motivate (P2) are true, thereby undermining the motivation for invoking special mental features to defend the standard reply. In the next section we see the simple explanation of (Strike) and then in §2.3 apply it to other seemingly mysterious facts about Mary.

2.2 Explanations of (Strike)

The extant version of the standard reply best equipped to explain (Strike) is also perhaps the most popular version of the standard reply: the so-called **phenomenal concept view**. Advocates of this view claim that one can receive information — and entertain, believe, and know that information — via different concepts. Moreover, differences in the concept one employs can make for important differences in one's mental life. For instance, if one receives information via one concept and then again via another that information will likely strike one as new, whereas employing the same concept both times will likely not result in the information striking one as new. In addition to this, advocates of the phenomenal concept view hold that there is a concept Mary did not (and given the setup, could not) acquire in her room. She only came to employ that concept upon seeing the apple, because to possess the concept one must have a certain sort of experience. On this proposal, the ex-

_____

this.

43

planation of why the information Mary receives upon seeing the apple strikes her as new is that upon seeing the apple Mary comes to employ a new concept in receiving that information.

Cases in which old information strikes one as new are widespread. The most tempting form of explanation of why old information strikes one as new in such cases is of a piece with the phenomenal concept explanation: the agent employs a new concept in receiving the old information. Yet there is another kind of explanation of why old information strikes one as new which does not require one to use a new concept, and explanations of this kind are needed in at least some cases.

One might doubt that it is possible for old information to strike one as new when one does not employ a new concept and hence doubt that there can be an explanation of old information striking one as new that does not require one to use a new concept. But in fact this is possible. Consider the following cases:

> Sosseh learned years ago that Jon's birthday is December 4th. Though she hasn't thought about it in a year, she can still recall it. A mutual friend says to her 'by the way, Jon's birthday is December 4th.' She responds 'oh dang, that's coming up. Are we doing something?'

> Alex has an beginner-to-intermediate level of competency in Spanish. He sees on his calendar that his friend Jon's birthday is December 4th. Later, a friend rushes up to him with a worried look

and says with urgency "el cumpleaños de Jon es quatro de diciem-
bre!" He responds with surprise and then, translating what was
said to his native language, realizes that he already knew this.

These are both plausibly cases in which an agent receives old information that
strikes them as new despite them not employing a new concept.[4]

Indeed, we should expect there to be such cases based on the nature
of psychology. One theoretical role that concepts play is allowing us to for-
mulate true psychological laws. Without relativization to the use of concepts,
apparently true psychological laws like

(A)   If one wants that p and knows that ϕing is a necessary means
           to p then *ceteris paribus* one will ϕ

appear clearly false. Granted, the principle has a *ceteris paribus* clause, which
allows for some tolerable exceptions to the principle (i.e. cases which satisfy the
antecedent but not the consequent and yet owing to some special features are
not counter-examples to the principle). But even so, without being relativized
to what concepts one employs, such a principle appears hopelessly strong. So
see why, consider Lois Lane, who wants to find Superman and knows that
going to Clark Kent's office is a necessary means of finding Clark yet does

---

[4]Cases of this sort can be generated in other ways as well, such as from someone mistak-
enly taking one person to be two distinct people and noting the same fact about them on
two occasionsand from an agent having good (say, probabilistic) reason to think that they
were told different information when in fact it was the same.

not go to Clark's office. She satisfies the antecedent but not the consequent of (A) and yet doesn't appear to be special in a way that would make her a tolerable exception to it. By building into the law that one must employ the same concepts in order to satisfy each clause in the antecedent, we can avoid the problem posed by Lois: she fails to satisfy the antecedent of (A) because she employs different concepts in thinking about Clark.

Of course, even relativized to using the same concepts throughout, (A) is not exceptionless. Sometimes one has better things to do than ϕ or thinks one cannot easily ϕ. Such is the case in psychology: there are few if any theoretically interesting exceptionless psychological laws.[5] This is true for essentially all non-fundamental sciences.

The same goes for laws which relate one receiving old information to the information striking one as new. For instance, the following is a law of psychology:

(B)   If one knows that p and one receives the information that p then *ceteris paribus* the information one receives will not strike one as new.

As a law of psychology, this law admits of tolerable exceptions even when relativized to employing the same concepts throughout. So, there are cases in which one receives old information that strikes one as new despite one

---

[5]The exception to this is for laws that relate mental states which bear a constitutive relation to one another, e.g. knowing that p and believing that p.

46

employing a single concept. (Any theory of concepts which precludes this possibility is the worse off for it.)

The proposal is that Mary's case is an instance of this: she receives old information via an old concept and yet it strikes her as new. Of course, this proposal has the burden of explaining (Strike) and hence why Mary is a tolerable exception to (B). To see that explanation, consider explanations of the analogous facts about Sosseh and Alex. The explanation for Sosseh is clear: she received the information at very different times and didn't immediately call to mind what she had learned years ago. The explanation for Alex is similarly clear: he is not yet completely fluent in Spanish and though he understood what he was told, he did not immediately translate it to his native language.

These explanations appeal to the agents being cognitively imperfect or lacking some information. The explanation of (Strike) cannot make a similar appeal as Mary does have all the relevant information and is cognitively perfect. Rather, the explanation of why, upon seeing the apple, the information Mary receives strikes her as new is that Mary is not used to receiving information via color experience.

Mary is much like Alex, who is told the relevant information in a language with which he is only competent. Indeed, Alex might well have known ahead of time what he was about to be told in Spanish and yet still had the information strike him as new. He had to bring to bear knowledge of the relation between English and Spanish to see that it was old information. Likewise, upon seeing the apple, the information Mary receives strikes her as new

because she receives the information in a manner that she is not used to — via color experience. Granted, being fully informed and cognitively ideal she knew all along that she would receive old information. Yet still, she has to bring to bear what she learned previously to relate the information she receives upon seeing the apple to the rest of the information she has. This is because she has not yet gotten used to receiving information via color experience. This is the simple explanation of (Strike).

2.3   Comparing Explanations

We are now in a position to evaluate the advantages and disadvantages of the simple view. I'll focus on a comparison with the phenomenal concept view.

The simple view has one clear advantage over the phenomenal concept view: it does not incur the burden of articulating a theory of concepts which one cannot possess except by having certain experience — which are anathema to our best theories of concepts.

Yet one might worry that the simple view also falters for not positing such concepts. For instance, one benefit of the phenomenal concept view is that it can explain other apparent facts about Mary beyond (Strike). For instance, it can explain the apparent fact that Mary gains new knowledge when she leaves the room. On a standard extension of the phenomenal concept view, one can come to know information one already knows so long as one does so via a new concept. Since upon seeing the apple Mary receives information via a

new concept, on this proposal she can come to know that information despite already knowing it. In a similar fashion, the phenomenal concept view can capture other alleged facts about Mary, including those about her learning, making a discovery, and making epistemic progress.

On a straightforward version of the simple view, Mary does not gain new knowledge upon seeing the apple. She likewise does not learn anything or make a discovery. Yet the advocate of the simple view has a clear explanation of this: Mary is just like Sosseh and Alex: they may initially take themselves to gain new knowledge but neither of them actually do so. If we are ignorant of what concepts, say, Sosseh uses, we might reasonably take her to gain new knowledge upon overhearing the second stranger. But once we learn which concepts she employs our reasons for thinking she gains new knowledge are undermined. She acts as if she gains new knowledge but only because of unusual circumstance. Likewise, she does not learn or make a discovery. (Whether she makes epistemic progress is a bit trickier: she acquires more evidence for the claim that Jon's birthday is December 4th. In that respect she makes epistemic progress.)

Mary is the same: before one knows that she uses the same concept one might reasonably think that she gains new knowledge upon seeing the apple. But once we learn which concepts our reasons for thinking she gains new knowledge are undermined. Mary acts as if she gains new knowledge but only because of unusual circumstance. So, I submit, Mary does not gain new knowledge, learn anything new, or make any discovery when she sees the apple

(though she may, like Sosseh, make epistemic progress).

Though the simple view can tackle Mary's case, one might still worry that it does not have the same theoretical import to other problems for physicalism that the phenomenal concept view has. For instance, phenomenal concepts are invoked to account for the existence of an explanatory gap.[6] The simple view appears unable to offer a similar theoretical upshot.

On the contrary, the simple view meshes well with alternative reply to other problems for physicalism. The explanatory gap, for instance, arises because we fail to recognize that information we acquire from, say, the lab and from experience is the very same information. Such a failure of recognition admits of similar explanations to the one given for Mary: we receive the information from radically different sources. No differences in concepts needed.

---

[6]And relatedly, the non-*a priori* deducibility of phenomenal facts from physical facts.

Chapter 3

Experience Over Time

3.1 Introduction

One currently widely held view is that how things are phenomenally for a subject at a time is constitutively independent of how things are phenomenally for them at other times, and how things are phenomenally with a subject over a period is exhausted by and less fundamental than how things are phenomenally with them at each point during that period. The idea is that what it is like for me right now and what it was like for me at any other time bear no constitutive relation to each other. Moreover, what it is like for me over the last five minutes just reduces to what it was like for me at each point during those five minutes.

There are at least three central motivations for such a view. It is supported by a version of temporal internalism about the phenomenal as well as (at least for intentionalists) a Geachean observation about the non-unfolding nature of mental representation. It is also supported by an emphasis on visual experience in contemporary philosophy of perception together with an assimilation of visual representation to that of pictures or film. But regardless of why it is held, the view often forms part of the background theory against which discussions of experience and experiential content take place.

I will argue here that this view is false. Experience is not static and momentary in this way. Rather, experience is at least sometimes irreducibly temporally extended — sometimes undergoing experience over a period is more fundamental than than undergoing experience over parts of that period. I will show that this view dissolves a certain longstanding philosophical puzzle about

temporal experience and allows one to more clearly theorize about certain categories of the mind.

In the first half of the paper, §3.2, I argue for the proposed view by considering a puzzle concerning experience of dynamic and temporal features. In §3.2.1, I present that puzzle. By way of resolving it I introduce and regiment, in §3.2.2 and §3.2.3 respectively, a distinction between ways of being that essentially take some time to go on and those that do not. I then show, in §3.2.4, that this distinction reveals an equivocation in the original puzzle and that if some experience is of the former sort, the puzzle dissolves.

In the second half of the paper I elaborate the resulting view and respond to potential objections. In §3.3.1, I respond to the claim that on the proposed view, the temporal features of experience and those presented in experience must match in certain ways special among all sensory qualities. In §3.3.2, I respond to to the claim that the view is inconsistent with a Geachean observation to the effect that mental representation does not unfold successively. In §3.3.3, I respond to the worry that on this view phenomenal experience is (temporally) externally determined in certain respects. In §3.4, I gesture at one upshot of the proposed view, that it can be used to clarify several distinctions in the philosophy of mind, including one between occurrent and non-occurrent mental phenomena.

3.2 Experience and Temporal Representation

Described most neutrally, the phenomena of interest in discussions of phenomenal experience are properties of a certain sort — properties of subjects — that I'll call **phenomenal properties**.[1] One way to pick out such properties is as the maximal determinates of the property of being phenomenally conscious. Put briefly, each phenomenal property is the salient property someone instantiates in all and only those cases in which things are phenomenally a certain way with them. To undergo experience is to instantiate a phenomenal property. For various reasons, I employ a notion of *experiential representation* in discussing phenomenal experience. For clarity, I do so against the background of intentionalism. On intentionalist views, phenomenal properties are intentional properties, that of experientially representing such-and-such (i.e. each phenomenal property is the property of bearing an experiential attitude to a content). This assumption is independent of the main thesis of the paper.[2]

3.2.1 The Puzzle

Temporal features are represented in experience. This much is hard to deny. Even setting aside temporal relations like succession and simultaneity, it is quite plausible for dynamic features like movement, crescendo, and throbbing. There are various ways one might get a grip on the class of sensible qualities, including as those that: figure into the accuracy conditions of

---

[1]See Nida-Rümelin 2007; Pautz 2009; Speaks 2015 for some discussion, though many philosophers employ a similar conception (sometimes calling them *experiential* properties).

[2]As shown in §3.2.4

experience, figure into true appears/looks reports of a certain 'phenomenal' variety, could be instantiated by an experience itself, could be instantiated by sense-data (as classically conceived), or 'immediately available to introspection.' On all of these ways of getting an initial grip on which features are sensible qualities, temporal features count among them.

Moreover, there are pairs of otherwise similar experiences which clearly differ with respect to the representation of temporal features. For instance, when one watches a series of film frames slowly succeeding one another (each shown for, say, 500 ms), one experiences a series of snapshots. But if the time of each frame is decreased (to, say, 30 ms), the items depicted seem to come to life and one can no longer visually identify individual frames. One straightforward interpretation of the difference between these two cases is that in the second but not the first, the viewer perceptually represents motion.[3]

Yet the claim that experience represents temporal features of some sort generates an apparent puzzle. Many have thought that if there is some period of time during which one represents temporal features (or states of affairs concerning such features — I'll move between the two) then there is some point during that period at which one does so. After all, this follows from the apparently obvious general principle that if something is F during some period of time then it is F at some instant during that period. But it isn't true in every case of temporal experience that there is a single point at which

---

[3]Such cases exhibit what is called the *phi-phenomenon*. See, for example, Wertheimer 1961.

one experientially represents the relevant temporal state of affairs.[4]

To see this more clearly, consider a case paradigm case of temporal experience, one in which Susan hears a sharp ping and almost immediately afterward a dull bang. Plausibly, she thereby represents the state of affairs *there was a ping and then a bang* (for short, *ping then bang*).[5] If she did so then there was some period of time (in fact, many) during which she did so. Call one such duration d. The following three jointly inconsistent claims are each plausible:

(S$_1$) During d, Susan experientially represents *ping then bang.*

(S$_2$) If, during d, Susan experientially represents *ping then bang* then there is some instant at which Susan experientially represents *ping then bang.*

(S$_3$) There is no instant at which Susan experientially represents *ping then bang.*[6]

The case is designed to be one in which **(S$_1$)** is true and we've have good reason to think that there are such cases. **(S$_2$)** follows from the general principle mentioned above. Yet **(S$_3$)** is also quite plausible. If it is false, there is

---

[4]Unlike for perception, as one might suddenly see the bird fly off.

[5]An analogous example with dynamic properties would work as well (when relevant, I treat dynamic features separately), as would an analogous example with singular contents.

[6]An instant is a duration-less point in time. I assume that there are instants but not that time is dense. The main argument of this paper goes through without this assumption. For simplicity I also assume eternalism.

an instant at which *ping then bang* is experientially represented. That, in turn, would be an instant at which both a ping and a bang are experientially represented since in order to represent a relational state of affairs one must represent its relata. But plausibly, there is no one instant at which Susan experientially represents both a ping and a bang. There was no instant at which a ping and a bang were simultaneously present in experience, not even when Susan heard the bang.

Note that these considerations are not as plausible for all cases of experiential temporal representation: perhaps one can imagine a temporal state of affairs at an instant.Nevertheless, in many cases, the considerations in support of **(S$_3$)** are powerful. We can generalize this puzzle as follows:

(1) Over some durations, subjects experientially represent temporal states of affairs.

(2) Whenever, over a duration, subjects experientially represent temporal states of affairs, they do so at some instants (during those durations).

(3) Subjects do not always experientially represent at instants temporal states of affairs.[7]

One central debate regarding temporal experience consists in philosophers leveraging two of these principles against the third. Antirealists employ **(2)**

---

[7]As above, in all three principles it is the representing that is over durations or at an instant, not (necessarily) what is represented.

and **(3)** to reject **(1)**, hence arguing that no claims like **(S₁)** are ever true.[8] It is important to keep in mind that many sensible qualities unfold over time — throbbing, vibrato, stabbing pain — and hence generate similar puzzles. Those who deny **(1)** must accept that none of these features are represented in such cases.

Others motivate **(1)** and **(2)** against **(3)**, claiming that whenever one experientially represents temporal states of affairs, one does so at an instant (a claim known as the **principle of simultaneous awareness**).[9] Such philosophers stress a distinction between the time over which one experiences something to occur — the duration of the activity of representing — and the time one experiences it to occur over — the duration specified by the content of the experience.[10]

One aim of this section is to show that this puzzle — and hence much of the motivation for the principle of simultaneous awareness — rests on an equivocation in **(2)** and **(3)**. Few have leveraged **(1)** and **(3)** against **(2)**, in part because **(2)** is taken to be an obvious truth.[11] Suitably disambiguated, it is clear how unextrordinary it is for there to be temporal experience without

---

[8]See, for example, Chuard 2011; Podevin 2007; Reid 1785–2002.

[9]Philosophers who accept this principle include Brentano 1988; Broad 1925–2008; Foster 1991; Grush 2005; Husserl 1905–1964; Lee 2014b; Strawson 2009; Tye 2003.

[10]See, for instance, Lee 2014a, p. 149. Some such philosophers advocate for *memory-retention* views according to which experientially representing temporal states of affairs is partly constituted by remembering events that just happened (in Susan's case, the ping). Others defend so-called *specious present* views according to which experience, independent of any memorial representation, represents a temporal duration at an instant.

[11]Recently some have denied the principle of simultaneous awareness, including Dainton 2000; Grube 2013; O'Shaughnessy 2000; Phillips 2008; Rashbrook 2013; Soteriou 2013.

there being any instant at which we represent temporal states of affairs in their entirety. The upshot of that discussion turns on the basic metaphysical nature of experience. To begin, however, we must consider a distinction among conditions and properties.

3.2.2 Two Kinds of Properties

When we use expressions like *so-and-so experienced such-and-such*, *had a such-and-such experience*, or *had an experience as of such-and-such*, we report a condition of a subject. Conditions, and reports of them, are classically divided into two categories: states and occurrences.[12] States include being angry, brightly colored, divorced, and 74° as well as hating Jon and believing that the end is nigh. Occurrences include processes like running, singing in a choir, growing old, dissolving, melting, and digesting one's breakfast as well as events like having climbed the mountain, rolled to the bottom of the hill, and run to the store. We can model conditions as properties, in part to meet up this discussion with that of phenomenal properties.[13]

One crucial difference between states and occurrences is that generally occurrences but not states are essentially *ongoing*. To see this distinction in

---

[12]This taxonomy originates from Kenny 1963; Vendler 1957 and the terminology from Mourelatos 1978. The occurrence/state distinction helps us get a grip on the ongoing/non-ongoing distinction, but the two may not perfectly match up. I set aside here what are sometimes called *achievements* or *instantaneous events*, such as winning the race, snapping, and popping.

[13]This is not a commitment to a metaphysical analysis of the three kinds of phenomena as properties but merely a way of modeling them. That being said, such analyses are commonplace.

more detail, consider the state of being spherical. Suppose a lump of clay is formed into a sphere at 2:30 and then reformed into another shape at 3:30, having not changed shape in the interim. It is true that the clay was spherical for that hour. Moreover, for each five minute period during that hour, the clay was spherical during that period. The same holds for every five second period, as well as every five millisecond period. Indeed, for every instant during that hour, it is true that the clay was spherical at that instant.

The same holds for processes like running and moving. Suppose Jon went for a run at 2:30pm and ran constantly for an hour.[14] For every five minute, or even five second, period during that hour, Jon was running during that period. Moreover, for every quarter-second period during the hour, Jon was running during that period (though of course he failed to complete even a stride in that time — see below). Indeed, for every instant during that hour, Jon was running at that instant. If Jon ran constantly from 2:30pm to 3:30pm then at the instant the clock began to chime three o'clock, Jon was running. In these respects, being spherical and running are alike.

If the lump of clay is spherical for some period, how the world is in other respects during that period is at least nomologically sufficient for the clay to be spherical during that period.[15] Moreover, for every instant at which the clay is spherical, how the world is in other respects at that instant is alone

---

[14]One can run for an hour without running constantly for an hour, if for instance one stops briefly at a crosswalk.

[15]Perhaps how the clay is intrinsically is sufficient for it to be spherical, or perhaps it depends on the structure of space-time during that period.

sufficient for the clay to be spherical at that instant. That is, being a sphere at a time nomologically supervenes on the state of the world at that time.

Yet analogous claims do not hold for running. For some spans of time it is true both that Jon was running during that period and that the state of the world in other respects during that period was alone insufficient for Jon to be running during that period. Consider some quarter-second period during Jon's run. Jon was running during that quarter second. Nevertheless, he failed to complete even a single stride in that time. The activity during that quarter second was not by itself sufficient for Jon to run, despite the fact that he was running during that period.

The same holds for instants. It is true both that Jon was running at exactly three o'clock and that the state of the world in other respects at exactly three o'clock was alone insufficient for Jon to be running at that time. At exactly three o'clock, Jon was positioned in some way. But no matter how he was positioned, his being so positioned is alone insufficient for him to be running at that time. No matter how quickly one can take a stride, it takes some time to run.[16]

The same applies to movement. If the ball is moving at exactly three o'clock, it is not solely in virtue of how the world is in other respects at just that time. Likewise for humming: one can be humming at exactly three o'clock,

---

[16]Perhaps there is no duration within which running cannot occur (e.g. The Flash can always run that quickly). Nevertheless, no matter how fast he is, it takes The Flash some time to run.

though it takes time to hum and hence how things are at exactly three o'clock do not alone make it true that one is then humming.

Crucially, the difference between running and being spherical characterized here is not between whether something can only be F over sufficiently long periods of time. Rather, the difference lies in what makes it the case that one is F over relatively short periods of time.[17] To be running over a quarter-second period one must be in the midst of a stride during that period. In this respect, running requires some temporal duration to occur. It is an ongoing property. Many properties are essentially ongoing and instantiated over time in this way. Running, melting, rolling, humming, walking with a group, burning, moving and digesting are all of this sort. Others are not, though of course they maybe instantiated for some time. Call the former **durative** and the latter **momentary**.

3.2.3 An Account of the Momentary/Durative Distinction

Let us spend some time clarifying the momentary/durative distinction before employing it to resolve our puzzle. Everything said in this section is independent of the general upshot of this paper, but it will help make clear the commitments (and non-commitments) of the resulting view. Moreover, it provides an account of the distinction that is theoretically valuable beyond its application to temporal experience.

---

[17]Cf. Rothstein 2004, p. 14.

One potential way to cash out the distinction is in terms of necessary conditions concerning other times on the property being instantiated at or over some time. On this approach, however, the qualifying necessary conditions must be restricted in some way, as there are trivial necessary conditions of this sort for all properties. Even besides those, many momentary properties require for their instantiation non-trivial facts concerning other times. Being married at t requires having participated in a marriage ceremony at some time before t. Yet being married is not an ongoing activity. Essential properties are also incorrectly classified as durative on this approach.

One might restrict the relevant necessary conditions by putting restrictions on what they concern or at which times they obtain. However, even the strictest restrictions will not rule out all false positives. Suppose we require that the necessary condition concern the very same property and concern the times in the immediate temporal vicinity of the relevant instant. On this view, a property F is durative iff necessarily, if o is F at an instant t then o is F at t', for some t' in the immediate temporal vicinity of t.

On this view, the property of being married is misclassified as durative, since one cannot be married for just an instant. If someone is married at all then they are married for a period of time (however short). Still, being married is not an ongoing activity like running or moving, which essentially take place over time. It is due, rather, to features of how one can enter into and cease to be married.

In formulating an account of durative properties we should keep in mind

the intuitive characterization of them as properties for which being instantiated over time is importantly prior to being instantiated at an instant. Consider movement again: the fact that o is moving at instant t is not merely a matter of how things are at that instant. What it is for o to move at t is for o to move over some duration d that includes t (unlike marriage, for which such durational facts are guaranteed but not explanatorily prior). It is hard to see how mere necessary conditions could make sense of this distinction. To do so we need a relation stronger than mere modal dependence - one of metaphysical dependence. One such relation is grounding, a metaphysical explanatory relation that (I'll assume) holds between facts.[18]

A simple grounding account has it that a property is durative just in case facts about its instantiation at a time are partly grounded in facts about other times. Yet this also misclassifies some momentary properties as durative. The fact that o is a heart at some point is partially grounded in facts about the evolutionary history of the species of the creature that possesses o. Nevertheless, being a heart is not an ongoing activity (though hearts engage in ongoing activities). Restricting the range of times to those just around t won't help either. So a simple account of that sort will not suffice. Rather, we must heed the observation above that for durative properties, durational fact are explanatorily prior. On such a view, a property is durative just in

---

[18]See, for instance, Fine 2012; Rosen 2010; Schaffer 2009. It is controversial whether metaphysical grounding is a relation between facts.

A formulation of the distinction in term of metaphysical grounding will provide insight into how the distinction might be made in terms of other dependence relations, such as property realization, essential properties, and fundamentality relations among events.

case facts about it being instantiated over a period partly ground facts about it being instantiated at instants during that period.[19] More precisely,

> **Durative Grounding (DG)** a property F is durative iff some fact of the form $[F(o)$ at $t]$ is partly grounded in a fact of the form $[F(o)$ for $d]$ for some duration $d \ni t$.[20]

On this account, the fact that Jon was running at 3:30 is grounded in the fact that Jon was running over some period that includes 3:30.[21] Compare this to a view on which the instantaneous facts are grounded in facts about some other activity around that time. On such a view, the fact that Jon was running at t is grounded in the fact that Jon was, for instance, positioned in certain ways around and at t. On such a view there are no fundamental durative properties — properties the facts about the instantiation of which are ungrounded. DG allows for such properties since it is silent on what, if anything, grounds the durational facts (i.e. the facts of the form $[F(o)$ over $d]$). This is a virtue of DG, if for no other reason than that it does not preclude theories on which experience is a fundamental property.

---

[19]This does not require diacronic grounding relations but only syncronic relations between facts concerning different times.

[20]A related view has it that a plurality of facts of the form 'o is F at t_i' (for all t_i in the relevant duration) together ground the fact that o is F at t. For various reasons, I prefer the present principle.

[21]For properties like movement any duration will do. For those like running, the period must be one in which some amount of activity took place. For many durative properties the duration must just contain a sufficient amount of some activity. If Jon was running from 3:00 to 4:00, the fact that he was running at exactly 3:30 isn't grounded in the fact that he was running for the entire hour. Rather, it is grounded in a fact about him running around that time.

This allows us to distinguish two senses in which durative properties can be instantiated. In some cases a durative property is instantiated at/over some time in virtue of its being instantiated over some longer, encompassing period. In other cases it is instantiated over a period not in virtue of being instantiated over an encompassing period. Let us say that in the former case the property is **partly** instantiated over that period while in the latter case it is **wholly** instantiated over the period. Jon *wholly* runs over a five minute period but *partly* runs over a .005 second period, and only *partly* runs at exactly three o'clock. More precisely,

> **Partial Instantiation (PI)** o is partly F over d/at t iff $[F(o)$ over d/at $t]$ is grounded in a fact of the form $[F(o)$ over $d']$ for some period $d' \ni d/t$.

Otherwise o is wholly F over d/at t. Momentary properties are always wholly and never partly instantiated while durative properties are sometimes wholly and sometimes partly instantiated, depending on the period in question.

3.2.4   A Solution

Applying this distinction to our original puzzle, there are two ways to take the claim that S experientially represents that p: as the claim that S wholly experientially represents that p and as the claim that S partly experientially represents that p. Consequently there is an ambiguity in our original

puzzle. **(1)** clearly means *wholly experientially represents*, but **(2)** and **(3)** can be taken in two ways, resulting in two inconsistent triads:

> **(1_w)** Over some durations, subjects **wholly** experientially represent temporal states of affairs.
>
> **(2_{ww})** Whenever, over a duration, subjects **wholly** experientially represent temporal states of affairs, they **wholly** do so at some instants (during those durations).
>
> **(3_w)** Subjects do not always **wholly** experientially represent temporal states of affairs at instants.[22]

and

> **(1_w)** Over some durations, subjects **wholly** experientially represent temporal states of affairs.
>
> **(2_{wp})** Whenever, over a duration, subjects **wholly** experientially represent temporal states of affairs, they **partly** do so at some instants (during those durations).
>
> **(3_p)** Subjects do not always **partly** experientially represent temporal states of affairs at instants.

---

[22]As before, the duration here is that of the representing, not of what is represented.

If experience is a momentary property, as many have assumed, then by PI there is no partial experiential representation; if Susan ever represents *ping then bang* then she wholly does so. That is, if experientially representing is always momentary then experientially representing that p is wholly experientially representing that p, $(\mathbf{2}_{ww})$ is true and original dilemma results.[23] §§3.3.1-3.3.3 confront arguments for the view that all experiential properties are momentary.

If, however, experience is a durative property of subjects then $(\mathbf{2}_{ww})$ is false. It isn't always true that if you wholly experientially represent that p over a duration then you *wholly* experientially represent that p at some instant during that time. There is thus no pressure to deny $(\mathbf{3}_w)$ in the first puzzle. On the other hand, $(\mathbf{2}_{wp})$ is true (and is the intuitive principle we observed in §3.2.1). However, $(\mathbf{3}_p)$ can be denied without the same consequences as denying $(\mathbf{3}_w)$. Recall that one reason given above for accepting $(\mathbf{3})$ was that if it is false then when Susan experientially represents *ping then bang* she does so at an instant and hence experientially represents a ping and a bang at that instant, which is not the case. However, this inference only holds for *wholly* experientially representing that p and hence is only a consideration against $(\mathbf{3}_w)$. That $(\mathbf{3}_p)$ is false does not entail that there is an instant at which Susan experiences both a ping and a bang.

Consider an analogy with cinematic representation. Films represent temporal features - they no more solely represent snapshot-like states of af-

---

[23]The momentary theorist will of course deny both $(\mathbf{2}_{wp})$ and $(\mathbf{3}_p)$ in the second puzzle.

fairs than dot-matrix images solely represent shading at points. Yet the analog of $(\mathbf{3}_w)$ for film is false. For instance, a scene in which a dog barks and then, a moment later, a cat meows represents *bark then meow*. But no single instant during that scene is one at which *bark then meow* is wholly represented. Rather, there is a period of time over which *bark then meow* is wholly represented and hence at some instants during that time, *bark then meow* is partially represented. So the analog of $(\mathbf{3}_p)$ is true for film. Yet it does not follow from this, the fact that the film partly represents *bark then meow* at a time, that it represents a bark and a meow at that time.

Take another example of a temporally distributed representation: morse code. Suppose some signals sent over a duration d represent that the battle is won. Some portions of d won't contain activity that by itself represents that the battle is won, so the signals do not wholly represent that the battle over those portions. In fact, some portions of d won't contain activity that by itself represents anything at all, so the signals do not wholly represent anything over those portions. Nevertheless, the activity during all portions of d are parts of the total activity over d (which does wholly represent that the battle is won) so the signal does partly represent that the battle is won over those portions of d.

Durative representation can also be understood by analogy with some forms of spatial representation. Portraits represent spatial relations between at least some of the things they represent. For instance, a portrait might represent

that a mouth is below a nose.[24] So there is a spatial state of affairs it represents, *mouth below nose*. Moreover, some parts of the canvas represent *mouth below nose*, including some contiguous portions of the canvas that encompasses the portion that represents the mouth and the portion that represents the nose. However, at least some parts of such contiguous portions do not by themselves represent *mouth below nose*, though they do contribute to the whole portion so representing.

I advocate the following:

The Durative View (DV) Some phenomenal properties are durative.

On this view, sometimes one experientially represents a state of affairs over a period of time without wholly representing it at any point during that time. DV can straightforwardly explain the representation of any feature, dynamic or otherwise, that generates a puzzle analogous to the one presented here, including those regarding change in volume, pain intensity, and flavor.

So far I have avoided using the count-noun 'experiences.' We can, though, express DV in such terms. Experiences are a kind of event (i.e. that of a subject instantiating a phenomenal property). Susan has a half-second experience, an event of representing *ping then bang*. But as with events generally, temporal parts of that experience — especially instantaneous parts —

---

[24]Assuming the content of a portrait is propositional. Analogous claims can be made for non-propositional content.

needn't themselves also be events of experientially representing *ping then bang*. Not all of the temporal parts of the event of Jon's run are themselves runs. Not every temporal part of a car crash is itself a car crash. Likewise, not all temporal parts of Susan's half-second experience of *ping then bang* are themselves experiences of *ping then bang*. Even supposing that the temporal parts of Susan's experience are themselves experiences, it needn't be that those parts are moreover experiences in which *ping then bang* is represented. Experiences have temporal parts but the latter do not always inherit the properties of the former.

Though this discussion has been put in terms of intentionalism, DV is available to non-intentionalists. Theories of experience are theories of the nature and structure of phenomenal properties. There are non-representational durative properties, and it is open to, for instance, naïve realists, enactivists, sense-data theorists, and adverbialists to identify phenomenal properties with them.[25]

3.3   The Durative View

Put briefly, the durative view is the view that experience is at least sometimes irreducibly temporally extended — that the phenomenal properties one instantiates over a period do not always reduce to the phenomenal

---

[25]In fact, enactivists often frame their views as chiefly opposed to a so-called *snapshot* view of experience, similar in many ways to the view opposed here. See, for instance, Hurley 1998; Nöe 2004; O'Regan and Nöe 2001.

properties they instantiate at each point during that period. In §3.2 I presented an argument for this view on the basis of its ability to dissolve a puzzle about temporal experience. DV is also relevant to more general issues regarding experience, as seen in §3.4. Below I consider three objections to DV. But first I address a few potential misunderstandings of the view.

To begin, recall the definition of durative properties established earlier:

DG a property F is durative iff some fact of the form $[F(o)$ at $t]$ is partly grounded in a fact of the form $[F(o)$ for $d]$ for some duration $d \ni t$.

First recall that DG is only a claim of partial grounding. It is consistent with a property being durative that other facts ground facts about its instantiation. For instance, in order to be running at a time one must have certain intentions at or around that time. If so, facts about having such intentions plausibly also ground the fact that one is running at a time.

Since this view says nothing about what grounds durational facts it is not committed to durative properties being dynamic, i.e. requiring some change in their object. Some durative properties are not dynamic: constant humming and droning are both durative but instantiating them does not necessitate change in an object. DV does not entail that experience is always changing.

In fact, DV does not even entail that *all* phenomenal properties are durative, only that some are. Consider an analogous question for other durative

representational kinds. Many have thought that only some cinematic representational properties are durative and that whatever else they represent, for the brief period during which a single frame is displayed, a film represents what that frame represents and does so like a picture (that is, via a momentary property). On this view, cinematic representation consists at least partly in a succession of snapshot-like momentary representation.

Other representational kinds do not ever represent momentarily in this way. Transmitting information via morse code requires producing a representational signal. Yet there are no snapshot-like representational properties of the signal. All of its representational features are durative. Morse code is an irreducibly durative representational medium.

Correspondingly, one version of DV is the **pure durative view** according to which all phenomenal properties are durative. Granted, if phenomenal properties are not fundamental then facts about their being instantiated over a duration must be grounded in some facts or other. But they needn't be grounded in facts about momentary *representational* properties. Another version of DV is the **impure durative view** according to which some phenomenal properties are momentary. One natural version of such a view has it that experience consists in a sequence of momentary representational properties which ground durative representational properties, much like how on some views the frames of a film give rise to the film's additional representational features. Another has it that experience consists in both both durative and momentary

representation but that there is no grounding relation between the two.[26]

One might claim that if we experientially represent momentary properties then we do so via momentary representational properties and hence not all phenomenal properties are durative. But this does not follow. One represents the smoothness of a surface by moving a part of one's body over it. The property represented is momentary but it is represented via a durative phenomenal property.

These have been mostly clarifications of the view. Let us more now to more substantive problems.

3.3.1 Temporal Mirroring

An objection one might raise against DV is that it entails a false principle about the relationship between the temporal features of experience and the temporal features represented in experience. In general, a representation of F need not itself be F. A representation of a dog, lime green, or 75° need not itself be a dog, lime green, or 75°. Yet some philosophers argue that experiences of duration themselves must have duration. Moreover, they claim that the temporal features represented in experience match the temporal features instantiated by experience.[27] For instance, if one experiences a particular pe-

---

[26]There is even a version of impure DV that allows for both momentary and durative representation of temporal features in experience. In fact, there is reason to think that we represent temporal features in two ways, corresponding to drop-offs in preservation of temporal information beyond 150-300ms and beyond 2-3 seconds. A combined view is well placed to explain this.

[27]This must be modified in various ways. See Phillips 2014; Soteriou 2010.

riod of time to pass then it took just that long to do so. Some examples used above may suggest that proponents of DV are committed to such a claim.

DV does not entail a mirroring thesis. Consider again film. The temporal features of a film need not match the temporal features represented in the film. Films include jumps in time, flashbacks, slow-motion and time-reversal sequences, as well as periods during which no temporal passage is represented at all (as when things are represented as frozen in time). These are all cases where the temporal features of a film come apart from the temporal features they represent. In the case of morse code this separation is even more apparent. The temporal features represented by a morse code signal rarely have any interesting connection with the temporal features of the signal. As far as DV goes, experiential representation might be like film or morse code in this respect.

Insofar as all representation takes place in time, there is a general connection between vehicle and content: in order to represent a temporal property something must instantiate some temporal property or other. But the same holds for spatial representation; insofar as all representation takes place in space, in order to represent a spatial property something must instantiate some spatial property or other. These principles reflect general facts about what is required to exist in a spatial or temporal world, not anything in particular about the nature of representation.

3.3.2 Stream of Consciousness

It is sometimes argued that mental representation is not *continuous* in a certain respect, and such complaints might be extended to show that no mental representation is durative. Discussions of this issue generally center around an objection from Geach against the idea of a Jamesean stream of consciousness.[28] Geach is specifically concerned with thought but his points generalize to other forms of mental representation.

One of Geach's main claims is that one does not mentally represent that p over a period by mentally representing parts of the content p over parts of that period. For instance, in thinking that Fido chased Sookie one does not first represent Fido, and then chasing, and then Sookie. Applied to experience (and formulated in terms of grounding), Geach claims

> **Geach's Constraint (GC)** experientially representing that p is not a property F such that at least some facts of the form $[F(o)$ over $d]$ are grounded in facts of the form $[F'(o)$ over $d']$

(where $d' \in d$ and F' is the property of representing a part of what F is a property of representing).

Let us suppose that Geach is right that experiential representation is not successive in this way. DV does not by itself conflict with this observation. It is consistent with DV that durational facts about experience are not

---

[28]Geach 1969, p. 34, James 1890.

grounded in facts about successive representing, and thus it isn't true that one represents that p over d in virtue of representing logical parts of the content p over parts of d.[29]

Putting too much stock in the analogy between durative representation and film may lead one to think that all durative representation must be successive in the way GC rules out. Yet other forms of durative representation are clearly not this way, as the previous examples of morse code should make clear. Or consider another case. Suppose that Noe asks Greg whether the two of them can go to the beach tomorrow, to which Greg replies *sure*. In saying this, Greg represented that they can go to the beach tomorrow and did so via a durative property. It essentially takes some time for him to say that they can go to the beach, and the entire (albeit short) time he was speaking, he was so representing. Yet Greg did not represent that they can go by successively representing parts of that proposition. There were no additional representational properties successively instantiated when he said that they can go. There was no shorter, more basic representing that gave rise to Greg representing that they can go. Durative experiential representational properties might be like this: they are essentially instantiated only over durations but not in virtue of other representational properties being instantiated successively over those durations. This would preclude the more natural version of the impure durative view sketched above.

---

[29]Cf. Soteriou 2013.

3.3.3   Phenomenal Non-localism

DV is compatible with phenomenal internalism as the latter is ordinarily formulated. On the standard formulation, phenomenal internalism is the claim that phenomenal properties supervene on intrinsic properties of subjects. Many durative properties are internal in this respect. Digesting, for instance, is both an intrinsic and a durative property of subjects.

However, DV is incompatible with a temporal phenomenal internalist thesis such as

> **Phenomenal Temporal Localism (PTL)** the phenomenal properties instantiated by a subject at a time supervene on the intrinsic properties of the world at that time.

Pairs of cases will exhibit a failure of supervenience of this sort if the two match up to some time at which one is cut short before enough activity has taken place for the durative property to be instantiated. For example, two runners spring from the blocks but one trips half way through her first stride. She was never running, though her twin was running as soon as she sprang from the block. A bit more abstractly, there are also so-called **Russell World Scenarios** in which the world sprang into existence part-way through when a given durative property would otherwise have been instantiated.[30]

---

[30]See Phillips 2008.

Fortunately, PTL is false.[31] Experience is realized by mental *activity* and mental activity (indeed, all activity) is durative. This alone does not guarantee that experience is durative — let us grant that a momentary property can be realized by a durative one.[32] Still, it does guarantee that whether one is experientially representing that p at a time at least requires that the mental activity that realizes such representation was going on just before or just after that time.and in such a way that if that activity hadn't occurred then the representational property would not have been instantiated. So PTL is false regardless of whether DV is true.

We might explain away the temptation to accept PTL by pointing out that a related claim holds, namely that someone can identify which phenomenal properties they instantiate at a time without knowing what occurred at other times.[33] I can know that I am digesting without first knowing anything about the state of my digestive system in the past or future. DV is compatible with this claim.

Some intentionalists have an additional reason to accept at least the possibility of durative experiential representation. It is empirically possible that the information output by the relevant perceptual processing mechanisms

---

[31]As are analogous claims for non-mental and non-phenomenal mental forms of representation.Note that non-mental forms of representation (e.g. speech, film, morse code) clearly violate it as do non-phenomenal mental states (e.g. calculating the tip, thinking through a proof, running through a mental list).

[32]This would be possible if facts about realized properties don't inherit the grounds of facts about their realizers.

[33]See Pautz 2014 for a similar discussion.

is temporally encoded. Many intentionalists are committed to the claim that such information is (or is intimately related to) experiential content. So such theorists should accept that DV is an empirical possibility.[34]

3.4   Conclusion

My main goal has been to argue for a view according to which experience is sometimes an irreducibly temporally extended phenomenon — that how things are for a subject over a period of time does not reduce to how things are for them at each point during that period. I hope to have shown here that this often overlooked view is theoretically viable and resolves certain longstanding philosophical problems and confusions.

I want to end by gesturing at an additional upshot of the momentary/ durative distinction. Philosophers of mind are especially concerned with so-called *personal-level* mental phenomena, those which are properly attributed to subjects rather than their parts (e.g. their cognitive systems). Personal-level mental phenomena include believing that the end is nigh, hoping for a brighter future, feeling a shooting pain up the thigh, loving Leonard Cohen, tasting ripe pineapple, fearing for one's life, and listening to crickets.

Personal level mental phenomena are often divided into so-called *occurrent* and *standing* kinds: believing that the end is nigh and loving Leonard Cohen are standing while feeling a shooting pain in the thigh and tasting

---

[34]See, for instance, Dainton 2014; Lee 2014b.

ripe pineapple are occurrent. Personal-level mental phenomena are also often distinguished into phenomenal and non-phenomenal kinds: feeling a shooting pain up the thigh and imagining a brilliant sunset are phenomenal while believing that the end is nigh and loving gin are not.

I have introduced a third distinction, between momentary and durative mental phenomena. As we saw above, mental phenomena like believing that the end is nigh and loving Leonard Cohen are momentary properties. I have argued that at least some phenomenal properties are durative and similar arguments could be extended to almost all phenomenal properties, including non-perceptual experience like inner monologue and imagination.

The paradigmatic examples on either side of these three distinctions are the same. This suggests a simple view, namely that among personal-level mental states the occurrent/standing distinction, the phenomenal/non-phenomenal distinction, and the durative/momentary distinction line up.[35]

---

[35]Crane 2013; Soteriou 2013 suggest related claims. Cf. Bartlett 2017.

Bibliography

Arjo, Dennis (1996). "Sticking up for Oedipus. Fodor on Intentional Generalizations and Broad Content." In: *Mind & Language* 11.3, pp. 231–245.

Aydede, Murat (1997). "Has Fodor Really Changed His Mind on Narrow Content?" In: *Mind and Language* 12.3-4, pp. 422–458.

— (1998). "Fodor on Concepts and Frege Puzzles." In: *Pacific Philosophical Quarterly* 79.4, pp. 289–294.

Aydede, Murat and Philip Robbins (2001). "Are Frege Cases Exceptions to Intentional Generalizations?" In: *Canadian Journal of Philosophy* 31.1, pp. 1–22.

Båve, Arvid (2008). "A Pragmatic Defense of Millianism." In: *Philosophical Studies* 138.2, pp. 271–289.

Bartlett, Gary (2017). "Occurrent States." In: *Canadian Journal of Philosophy*.

Block, Ned (1986). "Advertisement for a Semantics for Psychology." In: *Midwest Studies in Philosophy* 10, pp. 615–678.

Braun, David (1998). "Understanding Belief Reports." In: *Philosophical Review*, pp. 555–595.

— (2000). "Russellianism and Psychological Generalizations." In: *Noûs*, pp. 203–236.

— (2001). "Russellianism and Explanation." In: *Philosophical Perspectives*, pp. 253–289.

— (2002). "Cognitive Significance, Attitude Ascriptions, and Ways of Believing." In: *Philosophical Studies*, pp. 65–81.

— (2004). "Review: Conscious and Cognition." In: *Philosophy and Phenomenological Research* 69.2, pp. 484–491.

Braun, David and Jennifer Saul (2002). "Simple Sentences, Substitutions, and Mistaken Evaluations." In: *Philosophical Studies* 111.1, pp. 1–41.

Brentano, Franz (1988). *Philosophical Investigations on Space, Time, and the Continuum.* Trans. by Barry Smith. Croom Helm.

Broad, C.D. (1925–2008). *The Mind and Its Place in Nature.* Ed. by T.K. Paul. Routledge.

Broome, John (1999). "Normative Requirements." In: *Ratio* 12.4.

— (2013). *Rationality Through Reasoning*. Wiley-Blackwell.

Cappelen, Herman and Josh Dever (2018). *Puzzles of Reference*. OUP.

Chuard, Philippe (2011). "Temporal Experiences and Their Parts." In: *Philosophers' Imprint* 11.11, pp. 1–28.

Crane, Tim (2013). "Unconscious Belief and Conscious Thought." In: *Phenomenal Intentionality*. Ed. by Uriah Kriegel. Oxford University Press.

Crawford, Sean (2004). "Pure Russellianism." In: *Philosophical Papers* 33.2, pp. 171–202.

Dainton, Barry (2000). *Stream of Consciousness. Unity and Continuity in Conscious Experience*. Routledge.

— (2014). "Flows, Repetitions, and Symmetries. Replies to Lee and Pelczar." In: *Debates in the Metaphysics of Time*. Ed. by L. Nathan Oaklander, pp. 175–212.

Devitt, Michael (1996). *Coming to Our Senses. A Naturalistic Program for Semantic Localism*. Cambridge University Press.

Evans, Gareth (1982). *The Varieties of Reference*. OUP.

Fine, Kit (2012). "Guide to Ground." In: *Metaphysical Grounding*. Ed. by Fabrice Correia and Benjamin Schnieder. Cambridge University Press, pp. 37–80.

Fodor, Jerry (1990). "Substitution Arguments and the Individuation of Beliefs." In: MIT press.

— (1994). *The Elm and the Expert. Mentalese and its Semantics*. MIT Press.

Foster, John (1991). *The Immaterial Self. A Defence of the Cartesian Dualist Conception of the Mind*. Routledge.

Frances, Bryan (1998). "Defending Millian Theories." In: *Mind* 107, pp. 703–27.

Frege, Gottlob (1892). "On Sense and Reference." In: *Translations from the Philosophical Writings of Gottlob Frege*. Ed. by Peter Geach and Max Black. Trans. by Max Black. Basil Blackwell, Oxford, pp. 56–78.

Geach, Peter (1969). "What do we Think With?" In: *God and the Soul*. St. Augustine's Press.

Grube, Enrico (2013). "The Process of Experience." Dissertation. The University of Texas at Austin.

Grush, Rick (2005). "Internal Models and the Construction of Time. Generalizing from State Estimation to Trajectory estimation to Address

Temporal Features of Perception, including Temporal Illusions." In: *Journal of neural Engineering* 2.3.

Harman, Gilbert (1986). *Change in View.* MIT press.

Hawthorne, John and David Manley (2012). *The Reference Book.* OUP.

Heck, Richard Kimberly (1995). "The sense of Communication." In: *Mind* 104, pp. 79–106. (Originally published under the name "Richard G. Heck, Jr").

— (2002). "Do demonstratives have senses?" In: *Philosophers' Imprint* 2.2, pp. 1–33. (Originally published under the name "Richard G. Heck, Jr").

— (2012). "Solving Frege's Puzzle." In: *Journal of Philosophy* 109, pp. 132–174. (Originally published under the name "Richard G. Heck, Jr").

Hurley, Susan (1998). *Consciousness in Action.* Harvard University Press.

Husserl, Edmund (1905–1964). *The Phenomenology of Internal Time-Consciousness.* Ed. by Martin Heidegger. Trans. by James Churchill. Indiana University Press.

Jackson, Frank (1982). "Epiphenomenal Qualia." In: *Philosophical Quarterly* 32.April, pp. 127–136.

James, William (1890). *The Principles of Psychology.* Vol. 1. Henry Holt & Co. Dover Publication.

Kaplan, David (1989). "Demonstratives. An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals." In: *Themes from Kaplan.* Ed. by Joseph Almog, John Perry, and Howard Wettstein. Oxford University Press.

Kenny, Anthony (1963). *Actions, Emotions, and Will.* Routledge.

Kripke, Saul (1979). "A Puzzle About Belief." In: *Meaning and Use. Papers Presented at the Second Jerusalem Philosophical Encounter April 1976.* Reidel, pp. 239–83.

Lee, Geoffrey (2014a). "Extensionalism, Atomism, and Continuity." In: *Debates in the Metaphysics of Time.* Ed. by L. Nathan Oaklander, pp. 149–173.

— (Feb. 2014b). "Temporal Experience and the Temporal Structure of Experience." In: *Philosophers' Imprint* 14.3, pp. 1–21.

Loar, Brian (1988). "Social Content and Psychological Content." In: *Contents of Thought.* Ed. by editor. Ed. by Robery H. Grimm. Ed. by Daniel Davey Merrill. University of Arizona Press, pp. 99–110.

McGee, Vann (1985). "A Counterexample to Modus Ponens." In: *The Journal of Philosophy* 82.9, pp. 462–471.

McKay, Tom (1981). "On Proper Names in Belief Ascriptions." In: *Philosophical Studies* 39.3, pp. 287–303.

Millikan, Ruth (1993). "White Queen Psychology. or, the Last Myth of the Given." In: *White Queen Psychology and Other Essays for Alice*. MIT Press, pp. 279–363.

Mourelatos, Alexander (Jan. 1978). "Events, Processes, and States." In: *Linguistics and Philosophy* 2.3, pp. 415–434.

Nida-Rümelin, Martine (2007). "Dualist Emergentism." In: *Contemporary Debates in Philosophy of Mind*. Ed. by Brian McLaughlin and Jonathan Cohen. Blackwell Publishing, pp. 269–86.

Nöe, Alva (2004). *Action in Perception*. MIT Press.

O'Regan, Kevin and Alva Nöe (2001). "What It Is Like to See. A Sensorimotor Theory of Perceptual Experience." In: *Synthese* 129, pp. 79–103.

O'Shaughnessy, Brian (2000). *Consciousness and World*. Oxford University Press.

Pautz, Adam (July 2009). "What are the Contents of Experience?" In: *The Philosophical Quarterly* 59, pp. 483–507.

— (2014). "The Real Trouble with Armchair Arguments against Phenomenal Externalism." In: *New Waves in Philosophy og Mind*. Ed. by Mark Sprevak and Jesper Kallestrup. Palgrave Macmillan UK, pp. 153–181.

Phillips, Ian (2008). "Perceiving Temporal Properties." In: *European Journal of Philosophy* 18.2, pp. 176–202.

— (2014). "The Temporal Structure of Experience." In: *Subjective Time. The Philosophy, Psychology, and Neuroscience of Temporality*. Ed. by Valtteri Arstila and Dan Lloyd. MIT Press, pp. 139–158.

Podevin, Robin Le (2007). *The Images of Time. An Essay on Temporal Representation*. Oxford University Press.

Rashbrook, Oliver (Nov. 2013). "An Appearance of Succession Requires a Succession of Appearances." In: *Philosophy and Phenomenological Research* 87.3, pp. 584–610.

Reid, Thomas (1785–2002). *Essays on the INtellectual Powers of Man*. Ed. by Derek Brookes. Pennsylvania State University Press.

Richard, Mark (1990). *Propositional Attitudes. An Essay on Thoughts and How We Ascribe Them*. Cambridge University Press.

Rosen, Gideon (2010). "Metaphysical Dependence: Grounding and Reduction." In: *Modality: Metaphysics, Logic, and Epistemology.* Ed. by Bob Hale and Aviv Hoffman. Oxford University Press, pp. 109–136.

Rothstein, Susan (2004). *Structuring Events. A Study in the Semantics of Lexical Aspect.* Blackwell.

Salmon, Nathan (1986). *Frege's Puzzle.* Cambridge University Press.

Saul, Jennifer (2007). *Simple Sentences, Substitution, and Intuitions.* Oxford University Press.

Schaffer, Jonathan (2009). "On What Grounds What." In: *Metametaphysics: New Essays on the Foundations of Ontology.* Ed. by David Manley, David Chalmers, and Ryan Wasserman. Oxford University Press, pp. 347–383.

Schiffer, Stephen (1990). "The Mode-of-Presentation Problem." In: *Propositional Attitudes. The Role of Content in Logic, Language, and Mind.* Ed. by Anthony C. Anderson and Joseph Owens. Center for the Study of Language and Information, pp. 249–68.

Schneider, Susan (2005). "Direct Reference, Psychological Explanation, and Frege Cases." In: *Mind & Language* 20.4, pp. 423–447.

Sider, Theodore and David Braun (2006). "Kripke's Revenge." In: *Philosophical Studies* 128, pp. 669–82.

Soames, Scott (1988). "Direct Reference, Propositional Attitudes, and Semantic Content." In: *Propositions and Attitudes.* Ed. by Nathan Salmon and Scott Soames. Oxford University Press, pp. 197–239.

— (2002). *Beyond Rigidity.* Oxford University Press.

— (2006). "Reply to Critics." In: *Philosophical Studies* 128, pp. 711–38.

Soteriou, Matthew (2010). "Perceiving Events." In: *Philosophical Explorations* 13, pp. 223–241.

— (2013). *The Mind's Construction. The Ontology of Mind and Mental Action.*

Speaks, Jeff (2015). *The Phenomenal and the Representational.* Oxford University Press.

Strawson, Galen (2009). *Selves. An Essay in Revisionary Metaphysics.* Oxford University Press.

Thau, Michael (2002). *Consciousness and Cognition.* Oxford University Press.

Tye, Michael (1978). "The puzzle of Hesperus and Phosphorus." In: *Australasian Journal of Philosophy* 56, pp. 219–24.

— (2003). *Consciousness and Persons.* MIT Press.

Vendler, Zeno (Apr. 1957). "Verbs and Times." In: *The Philosophical Review* 66.2, pp. 143–160.

Wertheimer, Max (1961). "Experimental Studies on the Seeing of Motion." In: *Classics in Psychology.* Ed. by Thorne Shipley. Philosophical Library.