

Appeared in *Philosophical Studies* 105 (2001), pp. 59-105.

Russellianism and Prediction

David Braun

University of Rochester

Attitude ascriptions can be used to predict behavior. Consider, for example, the attitude ascriptions in (1).

- (1) Lucy wants Twain to autograph her book, and she believes that if she waves, then Twain will autograph her book.

If (1) is true, then Lucy is quite likely to wave. If I believe that (1) is true, and predict (on that basis) that Lucy will wave, then my prediction is quite likely to be right.

But some philosophers, including Michael Devitt (1996) and Mark Richard (1997a), think that none of this would be so if a certain theory of the semantics of attitude ascriptions were true, namely the theory that I shall call *Russellianism* (also sometimes called 'neo-Russellianism' and 'the naive theory'). For suppose that Lucy sincerely utters (2).

- (2) If I wave, then Clemens will autograph my book, and Twain won't. But I want Twain to autograph my book, and not Clemens.

Everyone agrees that if Lucy sincerely utters (2), then she believes that if she waves then *Clemens* will autograph her book. But Russellianism entails that, in this circumstance, Lucy also believes that if she waves then *Twain* will autograph her book, despite what she says using 'Twain'. So Russellianism says that (1) is true, in this circumstance. But clearly Lucy will *not* wave, in this circumstance. Thus, some critics say that if Russellianism is true, then (1) cannot

be used to predict behavior. But (1) clearly can be so used, so Russellianism is false. Call objections to Russellianism that take this sort of line *Prediction Objections*.

In this paper, I present several different Prediction Objections, and defend Russellianism from each. The objections differ (mostly) over the sense in which they say that attitude ascriptions can be used to predict behavior. I argue that, under Russellianism, attitude ascriptions can, in each sense, be used to predict behavior. I further argue that, even *if they cannot* be used to predict behavior under Russellianism, it would still *appear* to ordinary speakers that they can. Therefore, Russellians can either accept or explain away the intuition that attitude ascriptions can be used to predict behavior.

Philosophers have lodged several complaints against Russellianism that are closely related to the Prediction Objections. Mark Crimmins (1992) has claimed that ordinary psychological generalizations would be false under Russellianism. Devitt (1996) and Richard (1990, 1997) have argued that if Russellianism were true, then attitude ascriptions would not *explain* behavior. In fact, Devitt and Richard often present their complaint about explanation "in the same breath" as their complaint about prediction. (We will see one reason why they might do so below.) These issues are indeed closely related; consequently, my replies to the Prediction Objections make use of some of my previous work on Russellianism, explanation, and psychological generalizations (Braun 2000, forthcoming). But the Prediction Objections raise distinct issues that require separate treatment.¹

1. Russellianism

The view that I wish to defend says that (among) the objects of attitudes such as

believing, asserting, and desiring are *Russellian propositions*: structured entities whose constituents are individuals, properties, and relations. These propositions are also the *semantic contents* (or simply *contents*) of sentences with respect to (or *in*) contexts. Sentences *semantically express* Russellian propositions in contexts. The words that appear in sentences also have contents in contexts, which appear as constituents of the propositions that those sentences express in those contexts. The content of a predicate, in a context, is a property or relation. The content of a proper name, in a context, is the individual to which it refers in that context; similarly for simple indexicals like 'I' and 'that'. For example, the proposition expressed by 'Twain is an author', in any context, is a proposition whose constituents are Mark Twain (the man himself) and the property of being-an-author as constituents. We can represent this proposition with the ordered pair

<Twain, being-an-author>.

Since 'Clemens' refers to the same person as 'Twain', the sentence 'Clemens is an author' expresses exactly the same proposition, in any context.²

Russellianism further says that if *S* is a sentence, then the content and referent of the phrase 'that *S*', in a context, is the proposition expressed by *S*, in that context. The content of the predicate 'believes' is the binary relation of believing. Thus the proposition expressed by (3) can be represented by (4).

(3) Karen believes that Twain is an author.

(4) <Karen, <Twain, being-an-author>, believing>

Similar points hold for desire ascriptions like (5).

(5) Karen wants Twain to smile.

The content of the infinitival clause 'Twain to smile' is simply the proposition <Twain, smiling>.³

The content of 'wants' is the binary wanting relation. Thus sentence (5) expresses proposition (6).

(6) <Karen, <Twain, smiling>, wanting>

Russellianism has a number of well known virtues. The arguments of Keith Donnellan, Saul Kripke, David Kaplan, and others against descriptivist theories of proper names and indexicals naturally suggest the view and lend some support to it. The theory easily accounts for the use of indexicals in complement clauses of attitude ascriptions, and for quantification into such clauses. The view also avoids many of the epistemological and semantic problems of its Fregean rivals. (For details, see Salmon 1986, 1989; and Soames 1988, 1995.) But it also has a well known unintuitive consequence. According to it, sentences (7) and (8) express the same proposition, namely (9), and so these sentences must have the same truth value.

(7) Karen believes that Twain is an author.

(8) Karen believes that Clemens is an author.

(9) <Karen, <Twain, being-an-author>, believing>

But many philosophers think it is obvious that (7) and (8) can differ in truth value, and so conclude that Russellianism is false.

Russellians have replied at length to this *Substitution Objection*. According to Nathan Salmon (1986, 1989), Scott Soames (1988, 1995), and Thomas McKay (1979), (7) and (8) really do semantically express the same proposition, and so must have the same truth value; but utterances of them *pragmatically convey* different propositions. The propositions pragmatically conveyed by utterances of (7) can differ in truth value from those conveyed by utterances of (8);

moreover, speakers often confuse the proposition semantically expressed by utterances of (7) and (8) with the propositions they pragmatically convey, and so come to think that the sentences themselves differ in truth value. I have argued elsewhere (Braun 1998) that, whether or not the above claims about the pragmatics of (7) and (8) are correct, a rational speaker can believe both proposition (9) and its negation; and if he does so in the right *ways*, then he may also believe that (7) is true and (8) is false.

This last remark brings me to some metaphysical claims that are part of the theory that I wish to defend here. The above Russellian *semantics* says that the contents of 'believe' and 'desire' are binary relations that hold between agents and propositions. But many Russellians also accept a *metaphysics* of these relations that says that they are *mediated*: agents believe and desire propositions in virtue of standing in certain relations to intermediaries that determine the propositions believed and desired. These intermediaries are *ways of taking propositions*; they are also sometimes called 'propositional guises' or 'modes of presentation'; I will sometimes call them 'ways of believing' and 'ways of desiring' propositions. According to this metaphysical theory, an agent can believe a single proposition in more than one way ("twice over" or more, so to speak). An agent can also believe a proposition in one way, while failing to believe it in another way. Most importantly, a rational agent can believe a proposition in one way while believing the negation of that same proposition in a suitably different way.⁴

To illustrate, let's consider sentences (10) and (11).

(10) Twain is an author.

(11) Clemens is an author.

According to the above Russellian semantics and metaphysics, these sentences express the same

proposition, but that single proposition can be believed in distinct ways. An agent who believes that proposition in a certain way will be disposed to assent to (10) and think that (10) is true.⁵ An agent who believes that proposition in a certain other way will be disposed to assent to (11) and think that (11) is true. An agent who believes the proposition in the first way, but not in the second, will think that (10) is true, but may have no opinion about (11). There is also a way of believing the *negation* of the proposition expressed by (10) and (11) such that an agent who believes that proposition in that way will think that (11) is false. A rational agent could believe the proposition expressed by (10) in the first way, while believing the negation of that same proposition in the third way. Such an agent would believe both a proposition and its negation, in different ways that are suitable for his being rational.

This theory can be fleshed out by saying more about the nature of ways of taking propositions. One common view of them identifies them with mental representations (see Richard 1990 and Crimmins 1992). On this view, mental representations have propositional contents. An agent believes a proposition by having a mental representation in her head (in the right manner) whose content is that proposition. On this view, an agent could believe the proposition expressed by (10) and (11) in one way by having a certain mental representation in her head; she could believe the proposition in another way by having a different mental representation in her head; and she could believe its negation by having yet another mental representation in her head. If the mental representation that allows her to believe the proposition expressed by (10) is suitably different from the mental representation that allows her to believe the negation of that proposition, then the agent can rationally believe both propositions. Similar points hold for mental representations and desiring propositions.

A similar situation can arise with attitude ascriptions such as (7) and (8).

(7) Karen believes that Twain is an author.

(8) Karen believes that Clemens is an author.

Russellian semantics says that these express the same proposition, but (typical) Russellian metaphysics says that this single proposition can be believed in distinct ways. A rational agent could believe the proposition expressed by (7) and (8) in one way, while also believing the negation of that proposition in a suitably distinct way. Such an agent could think that (7) is true and (8) is false, and be inclined to assent to (7) and dissent from (8).

If this Russellian metaphysics is correct, then the way in which an agent believes a proposition can make a difference to behavior. For example, if Karen believes the proposition expressed by (10) in a certain way, she will assent to it, but if she does not, then she may not assent. Yet according to Russellian *semantics*, the proposition expressed by belief sentence (7) says nothing about the way in which Karen believes the proposition expressed by (10). That is, the proposition semantically expressed by (7) does not have a way of taking a proposition as a constituent. Thus one might suspect that if Russellianism were true, then (7) could not be used to predict Karen's assent and dissent behavior towards (10). Similarly for other attitude ascriptions and bits of behavior.

By contrast, typical anti-Russellian views, such as those of Richard (1990), Devitt (1996), Crimmins (1992), and Forbes (1990), say that the proposition expressed by (7) does semantically provide information about the way in which Karen believes the proposition.⁶ Thus, many anti-Russellians argue that attitude ascriptions could be used to predict behavior, if their views were correct, but could not be so used, if Russellianism were correct.

2. Mismatch in Ways

Let's examine the case of Lucy a bit more closely from a Russellian perspective, while keeping the Russellian metaphysics in mind. From here on, let's imagine that *in the actual world*, Lucy sincerely utters (12).

- (12) I want Twain to autograph my book, and if I wave, then Twain will autograph my book.

Russellians, and nearly everyone else, will agree that in the actual world Lucy believes that if she waves then Twain will autograph her book, and that she also wants Twain to autograph her book. In short, (1) is true.

- (1) Lucy wants Twain to autograph her book, and she believes that if she waves, then Twain will autograph her book.

Russellians would add that in the actual world, Lucy believes and desires these propositions in *matching ways*. She believes the proposition that (if she waves then Twain will autograph her book), in a "Twainish" way. She also wants (the proposition) that Twain autographs her book, in a "Twainish" way.⁷

Now consider another, nomologically possible world, in which Lucy sincerely denies (12) and sincerely utters (2).

- (2) If I wave, then Clemens will autograph my book, and Twain won't. But I want Twain to autograph my book, and not Clemens.

Everyone would agree that, in this world, Lucy believes that if she waves then Clemens will autograph her book. Russellians hold that she also believes that if she waves then *Twain* will autograph her book, despite her utterances of the 'Twain' sentences in (2). But in this world she

believes that proposition in a "Clemensish" way, and not in a "Twainish" way.

Thus in this world there is a *mismatch* between the ways that Lucy believes and desires the propositions we've mentioned. She believes the first proposition in a "Clemensish" way, but not a "Twainish" way. She desires the second proposition in a "Twainish" way, but not a "Clemensish" way. In such a mismatch world, she won't wave. But if her ways of believing and desiring matched, as they do in the actual world, then (other things being equal) she would wave. Again, we see that the ways in which an agent believes and desires propositions "make a difference" to how the agent behaves.

3. Predicting, Predictions, and Using Sentences to Predict

To understand the upcoming Prediction Objections, we need to understand what it means to use a sentence (such as an attitude ascription) to predict an event (such as a bit of behavior). We also need to think more generally about the nature of predicting and prediction.

The predicate 'predict' appears to express a binary relation that can hold either between a person and a proposition, or between a person and a (type of) event. Consider, for example, (13) and (14).

(13) Joe predicted that Mary would wave.

(14) Joe predicted Mary's wave.

(14) relates Joe to the denotation of the description 'Mary's wave', which is an event (or type of event).⁸ (13) relates Joe to the proposition expressed by its 'that'-clause.⁹ (13) also relates Joe indirectly to a (type of) event, for the proposition expressed by 'Mary will wave' is such that, necessarily, it is true iff a waving event by Mary occurs. Let's say that the proposition that Mary

will wave is *made true* by events of this sort. And let's say that this proposition, and utterances that express it, *describe* such events.¹⁰

The noun 'prediction' exhibits an "act-object" ambiguity: it can be applied either to an *act* of predicting, or to an *object* that is predicted in an act of predicting. ('Belief' is ambiguous in a similar way.) When Joe predicts that Mary will wave, a certain sort of event occurs, an event that involves (at least) Joe, the proposition that Mary will wave, and the predicting relation. This event might be called 'a prediction'. I will say that it is an *act of predicting*. This event may be purely mental, or may be both mental and linguistic, for instance, an intentional utterance of the sentence 'Mary will wave'. The proposition to which Joe is related during the act of predicting is a *predicted proposition*.¹¹ Predictions, in the sense of "objects of acts of prediction", may be either propositions or (types of) events.¹²

Typically, a person engages in an act of predicting an event when the event lies in the future of the act, as yet unobserved by the predictor. The *basis* for her act of predicting is the set of propositions from which she infers that the event will occur. Typically, her basis includes propositions about such matters as (a) the events that have occurred, or are now occurring, or will occur, up until the time *t* of the to-be-predicted event, (b) information about other matters of particular fact holding before *t*, and (c) various laws, or other generalizations, relating types of events and matters of particular fact. Propositions of these sorts are *admissible* into a basis for predicting the event. (Some propositions are *inadmissible* into a basis for predicting an event: for example, the proposition that the event occurs, the proposition that Smith knows that the event occurs, and any proposition about times after *t*).¹³ The predictor's actual basis may or may not include propositions of all the sorts mentioned above.

An act of uttering a prediction is a result, or causal consequence, of some sort of inference. But a more immediate result of such inferences is a believing-event with a propositional content. Thus it's reasonable to think that whenever an act of predicting an event occurs, there also occurs an act of predicting a proposition, for instance, the proposition that the event will occur, or some other proposition describing the event.

These reflections suggest the following tentative analyses. If A is an agent and P is a proposition, then A predicts P iff: there is a set of propositions B such that (i) B contains only propositions that are admissible with respect to P ; (ii) A believes all propositions in B ; and (iii) A infers P from B , and consequently comes to believe P .¹⁴ If E is an event, then A predicts E iff A predicts a proposition P that describes E .¹⁵

Notice that this tentative analysis does *not* require that the act of predicting occur *before* the predicted event. A physics student can be said to predict a past event if he infers that it occurs using only admissible information, e.g., lawful generalizations and information about particular matters of fact that held before the event occurred.¹⁶

We need to consider at least two more notions before we can proceed to the Prediction Objections. The first, and easier, of these is the notion of an agent's using a sentence to predict an event (or a proposition describing the event). We can explicate this idea as follows. A uses sentence S to predict P iff: (i) A predicts P ; (ii) A thinks that S is true and believes the proposition expressed by S ; (iii) the proposition expressed by S is a member of A 's basis for predicting P .¹⁷

4. "Can Be Used to Predict"

The last notion we need to understand before turning to the Prediction Objections is the most crucial: sentence S can be used to predict P . (Here S might be an attitude ascription, and P a proposition describing a bit of behavior.) If we paid attention only to the linguistic form of our analysandum, we might be tempted simply to modalize and existentially generalize on the earlier analysis of $\leq A$ uses S to predict $P \leq$. The result would be: S can be used to predict P iff *possibly*, *some agent* uses S to predict P . But under this analysis, it's obviously true that (1) can be used to predict Lucy's wave, whether or not Russellianism is true. For surely in some possible world some agent uses ascription (1) to predict that Lucy will wave, whether or not Russellianism is true in that world (or ours), and whether or not that prediction is "wise". In fact, in some possible world very like ours, some dim-witted agent uses ascription (1) to predict that Lucy will *fly*, and also uses the sentence 'David Braun sneezed on July 1, 2000' to predict that the temperature in Rochester, NY at noon on January 1, 2100 is 300 degrees Fahrenheit.

Clearly the notion of "can be used" that critics have in mind is stronger than this one. It concerns not so much whether some possible agent uses the sentence to predict P , but rather whether the resulting prediction is *true* or *reliable*.¹⁸ I shall try to motivate some analyses of this notion that these critics might accept by considering an example that does not involve attitude ascriptions.

Consider (15).

(15) Carl is a car whose ignition key is being turned to its start position.

(15) can be used to predict that Carl's engine will start. This claim seems correct, but what does it mean? A first guess might be: given the laws of nature, the turning of Carl's ignition key is

sufficient for the starting of Carl's engine. Another way to put it: it's *nomologically necessary* that if Carl's ignition key is turned to its start position, then Carl's engine will start. But this can't be quite right, for it's just not true that such a turning nomologically necessitates a starting. There may, after all, be some abnormal, interfering factors present; for example, Carl's gas tank may be empty, or there may be moisture in Carl's electrical system that prevents sufficient spark from occurring. But under more favorable conditions, and in the absence of interfering conditions, the engine will start. So we might think that, given the laws, the turning of Carl's key is sufficient, *other things being equal*, for Carl's engine to start. Rephrasing, we get (16).

- (16) It's nomologically necessary that: if Carl is a car whose ignition key is being turned to its start position, then, other things being equal, Carl's engine will start.

The conditional embedded inside (16) is commonly called a '*ceteris paribus* conditional'.¹⁹

These reflections motivate the following attempt to explain "can be used to predict", which I call the 'Nomological Sufficiency Analysis'.

- (17) *The Nomological Sufficiency Analysis*

S can be used to predict P iff: it's nomologically necessary that if S is true, then, other things being equal, P is true.

Using this analysis, we can infer that (15) can be used to predict that Carl's engine will start iff (16) is true.²⁰

There are phrases besides 'other things are equal' that are commonly used in *ceteris paribus* conditionals, for example, 'typically', 'usually', and 'it is quite likely that'. Thus we could rephrase the conditional embedded in (16) with (18).

- (18) If Carl is a car whose ignition key is being turned to its start position, then it is quite likely that Carl's engine will start.

At the beginning of this paper, I used the phrase "quite likely that" to paraphrase the claim that (1) can be used to predict that Lucy will wave: I said that if (1) is true, then Lucy is quite likely to wave. I suspect that you accepted this paraphrase without hesitation. I take this to be some confirmation of the analysis given in (17).

There's another closely related sense in which (15) can be used to predict that Carl's engine will start, a sense which some critics of Russellianism might have in mind. In the last section, we found that an agent predicts P only if he infers P from a basis of admissible propositions. If the agent is reasonable, then it should be possible to construct a reasonable *argument* from the propositions in the basis to the conclusion P . Typically, the argument would include some generalizations, probably generalizations that the agent takes to be laws or "near laws". Scientists sometimes explicitly formulate such predictive arguments. The arguments that justify their predictions could be viewed as a kind of ideal for predictive arguments, and for acts of predicting.

Reflections like these might motivate the following idea of an *ideal predictive argument*. An ideal predictive argument for P is a sequence of sentences that has as its conclusion a sentence expressing P . It is valid, or in some other way strongly *supports* its conclusion. Its premises are all admissible with respect to P . At least one of the premises is a law, or a law-like generalization. All of the premises are essential in the following sense: if one of the premises were deleted, the remaining premises would no longer logically entail or support the conclusion. An ideal predictive argument whose premises are true is *correct*.

Sentence (15) could serve as the sole particular-fact premise in an ideal predictive argument for the conclusion that Carl's engine will start. Consider argument (19).

- (19) a. Carl is a car whose ignition key is being turned to its start position.
b. For all x , if x is a car whose ignition key is being turned to its start position, then, other things being equal, x 's engine will start.
c. Therefore, Carl's engine will start.

(19a), which is just (15), is the only premise concerning particular fact in argument (19). (19b) is a *ceteris paribus* generalization. If (15) is true, and (19b) is a true law-like generalization, then (19) is a *correct* ideal predictive argument, whose conclusion is (therefore) likely to be true (in some sense).

A theorist might think that an *act* of predicting is ideally commendable only if it is the result of an inference that (in some way) parallels an ideal predictive argument. He might therefore think that (15) can be used to predict that Carl's engine will start only if (15) figures in an ideal predictive argument whose other premises are true. These ideas might encourage a theorist to propose the following analysis of "can be used to predict", which I call the 'Predictive Argument Analysis'.

(20) *The Predictive Argument Analysis*

S can be used to predict P iff: S is the sole particular-fact premise in some ideal predictive argument whose general-premises are true and whose conclusion expresses P .

Those who are familiar with *covering-law* theories of *explanation* may notice that the above ideal predictive arguments strongly resemble the *ideal explanations* (or ideal explanatory

arguments) of such covering-law theories. In fact, many who accept covering-law theories of explanation hold that explanations and predictions are deeply similar. For example, Carl Hempel says that the difference between an explanation and a prediction is merely pragmatic: an explanation of an event is an argument given after the event occurs (and after the event is known to have occurred) whereas predictions are (typically) made before the event occurs, with the help of a supporting argument. Thus Hempel once proposed that all and only ideal explanations are ideal predictive arguments. (See Hempel 1965b; and Hempel and Oppenheim 1965.) Those who accept covering-law theories of explanation might similarly think that ideal predictive arguments are somehow central to any theory of (reliable, reasonable) prediction, and so might well be tempted to accept something like the Predictive Argument Analysis.

The critics of Russellianism who are most explicit about presenting Prediction Objections to Russellianism, namely Devitt and Richard, seem to accept some sort of covering-law theory of *explanation*. (They do not explicitly endorse any such theory, but there are many signs in their work that they tacitly presuppose such a theory.) And when they mention Russellianism's (alleged) problems with prediction, they always do so "in the same breath" in which they mention its (alleged) problems with explanation. Thus when they say that attitude ascriptions could not be used to predict behavior, if Russellianism were true, they may well be using the phrase "can be used" in a sense like that given by the Predictive Argument Analysis.²¹

There is clearly a close relationship between the Nomological Sufficiency Analysis and the Predictive Argument Analysis. If *S* figures as the sole particular-fact premise in an ideal predictive argument, then the truth of it must be nomologically sufficient, under the law-like generalizations mentioned in the argument, for the truth of *P*, if other things are equal. Thus, if

the Predictive Argument Analysis entails that S can be used to predict P , then it's nomologically necessary that if S is true then P is true, other things being equal, and so the Nomological Sufficiency Analysis is also satisfied. Conversely, if the Nomological Sufficiency Analysis is satisfied, then there are law-like generalizations that, together with the antecedent of the relevant *ceteris paribus* conditional, are sufficient to formulate an ideal predictive argument whose conclusion is the consequent of the conditional.²²

5. Two Prediction Objections

I am now ready to present two Prediction Objections, corresponding to the two analyses of "can be used to predict" that I presented above. I call these the *Predictive Argument Objection* and the *Nomological Sufficiency Objection*. These objections are inspired by Devitt and Richard, but I hesitate to attribute the objections to them, for their explicit statements of their objections are not as detailed as those that follow. Thus I shall attribute them to an imaginary critic.²³

Consider ascription (1) again.

- (1) Lucy wants Twain to autograph her book, and she believes that if she waves, then Twain will autograph her book.

Our imaginary critic claims that (1) can be used to predict Lucy's behavior. Using the Predictive Argument Analysis, our critic claims that if (1) can be used to predict that Lucy will wave, then (1) appears as the sole particular-fact premise in some ideal predictive argument whose general premises are true. Now any such ideal predictive argument will contain at least one *ordinary psychological generalization*, by which I mean a generalization that contains attitude ascriptions that do not explicitly mention ways of taking propositions. The simplest such argument is (21),

but there are more complicated ones like (22).

- (21) a. Lucy wants Twain to autograph her book, and she believes that if she waves, then Twain will autograph her book.
 - b. If a person wants Twain to autograph her book, and she believes that (if she waves then Twain will autograph her book), then, other things being equal, she will wave.
 - c. Therefore, Lucy will wave.
- (22) a. Lucy wants Twain to autograph her book, and she believes that if she waves, then Twain will autograph her book.
 - b. If a person wants Twain to autograph her book, and she believes that (if she waves then Twain will autograph her book), then, other things being equal, she will *intend* to wave.
 - c. If a person intends to wave, then, other things being equal, she will wave.
 - d. Therefore, Lucy will wave.

But, the critic maintains, if Russellianism is true, then no ordinary psychological generalization is true. Suppose, for example, that *Mary* (actually) believes and desires the propositions mentioned in (21b) in mismatching ways. Suppose that *Mary's* hands are not tied down, that she is not suffering from paralysis, and so on. All other things are equal. And yet she doesn't wave. So *Mary* shows that (21b) is false, if Russellianism is true. *Mary* similarly shows that (22b) is false, under Russellianism, for she satisfies the antecedent, and other things are equal, but she doesn't form the intention to wave. In fact, if Russellianism is true, then *no* ordinary psychological generalization is true, because of mismatch cases. But if this is so, then no ideal predictive

argument whose sole particular fact premise is (1) has true general premises. Thus (1) cannot be used to predict Lucy's wave, according to the Predictive Argument Analysis. But (1) can be used to predict Lucy's way, so Russellianism is not true. Rearranging and summarizing, we get argument (23).

(23) *The Predictive Argument Objection*

- a. (1) can be used to predict Lucy's wave.
- b. If (1) can be used to predict Lucy's wave, then some ordinary psychological generalization is true.
- c. Therefore, some ordinary psychological generalization is true.
- d. If Russellianism is true, then no ordinary psychological generalization is true.
- e. Therefore, Russellianism is not true.

The Nomological Sufficiency Objection is similar in outline. As before, the critic claims that it's obvious that (1) can be used to predict Lucy's wave. But according to the Nomological Sufficiency Analysis, if this is so, then *ceteris paribus* conditional (24) is true.

- (24) It's nomologically necessary that: if Lucy wants Twain to autograph her book, and she believes that (if she waves then Twain will autograph her book), then, other things being equal, she will wave.

But (24) is false, if Russellianism is true. For if Russellianism is true, then there is a nomologically possible world in which all other things are equal, and Lucy sincerely assents to (2), and thus believes and desires the mentioned propositions in mismatching ways. According to Russellianism, the antecedent is true in such a world. And yet Lucy fails to wave there. So if

Russellianism is true, (24) is false, and (1) cannot be used to predict Lucy's behavior. But (1) can be so used. So Russellianism is not true. Rearranging and summarizing, we get argument (25).

(25) *The Nomological Sufficiency Objection*

- a. (1) can be used to predict Lucy's wave.
- b. If (1) can be used to predict Lucy's wave, then (24) is true.
- c. Therefore, (24) is true.
- d. If Russellianism is true, then (24) is not true.
- e. Therefore, Russellianism is not true.

6. An Initial Reply

Both of the objections say that some important claim containing a *ceteris paribus* phrase is false under Russellianism. The Nomological Sufficiency Objection claims that (24) is shown false, under Russellianism, by what happens to Lucy in another nomologically possible world. The Predictive Argument Objection claims that all ordinary psychological generalizations are false under Russellianism, and all such generalizations contain *ceteris paribus* phrases.

The obvious Russellian reply is to claim that other things are *not* equal in the cases we considered above. For example, Mary believes and desires the relevant propositions in *mismatching ways*. This is sufficient for other things *not* to be equal. Thus Mary is no counterexample to generalization (21b) under Russellianism. Similarly, for Lucy: in the nomologically possible world in which she believes and desires the propositions in mismatching ways, other things are not equal simply because of this fact.²⁴

Of course, I need to defend the claim that other things are not equal in Lucy's

counterfactual case and Mary's actual case. To do this, it will be useful to think more generally about *ceteris paribus* conditionals and generalizations, starting with some that do not contain attitude ascriptions.

7. *Ceteris Paribus* and Context

Let's return to Carl, and consider the *ceteris paribus* conditional embedded inside (16), namely sentence (26).²⁵

(26) If Carl is a car whose ignition key is being turned to its start position, then, other things being equal, Carl's engine will start.

Let's say that *suitable conditions* hold when all other things are equal. Then we might paraphrase (26) as follows: if it were the case that (Carl is a car whose ignition key is being turned and suitable conditions hold), then it would be the case that Carl's engine will start. Carl himself shows that (26) is false if his ignition key is being turned and suitable conditions hold and yet his engine doesn't start. Thus to determine whether a conditional (or generalization) is false in a situation, we need to determine whether suitable conditions hold in it. But when we do so, we need to keep in mind that the suitable conditions associated with a *ceteris paribus* conditional vary from context to context. Let me explain.

Ceteris paribus conditionals are context-sensitive: their truth conditions (and contents) vary from context to context. In most contexts, (26) would seem to us to be true (keeping the world of evaluation fixed). But imagine the following scenario: Carl is my car. As I walk to the parking lot in the afternoon, I worry about whether I left Carl's lights on in the morning, and whether Carl's battery is dead. In such a context, (26) might seem to me to be false, but (27)

might seem to be true.

- (27) If Carl is a car whose ignition key is being turned to its start position, *and Carl's battery is not dead*, then, other things being equal, Carl's engine will start.

This judgment about the truth values of (26) and (27) seems reasonable, and is good evidence that (26) is context-sensitive: in some contexts, (26) expresses a true proposition, in others a false one. A reasonable explanation of this context-sensitivity is that the suitable conditions associated with a *ceteris paribus* conditional vary from context to context. In most contexts, the condition of Carl's battery's being live seems to be included in the suitable conditions associated with the conditional (just as Carl's having some gas in his tank, and his starter motor's being in working order, seem to be included in the associated suitable conditions for *both* contexts we considered). But in other contexts, the battery's being live seems not to be included in the suitable conditions associated with the conditional. In such contexts, the condition that the battery be live must be explicitly added to the antecedent to get a *ceteris paribus* conditional that is true in the context.

The most obvious way to account for this context-sensitivity is to assume that the phrase 'other things are equal' is context-sensitive. Or if the phrase 'other things being equal' forms a single connective with the 'if...then' construction (a plausible hypothesis), then we can say that this *ceteris paribus* conditional connective is context-sensitive. Let's say that the phrase 'other things being equal', or the *ceteris paribus* connective, determines a suitable condition *with respect to a context*. The phrase(s) may determine different suitable conditions with respect to different contexts.²⁶

Speakers' thoughts and intentions clearly play a large role in determining the suitable

conditions associated with a *ceteris paribus* conditional in a context. But speakers in a context usually lack detailed knowledge of the suitable conditions determined by their utterances of *ceteris paribus* conditionals. For example, Carl's engine will start when his ignition key is turned only if Carl's solenoid is attached to the starter motor. So this condition is part of the suitable conditions determined by typical utterances of (26). But many speakers do not know that cars have solenoids.

If *ceteris paribus* conditionals are context-sensitive in this way, then so are *ceteris paribus* generalizations. For example, generalization (19b) is context-sensitive.

(19b) For all x , if x is a car whose ignition key is being turned to its start position, then, other things being equal, x 's engine will start.

Similar points about context-sensitivity also hold for conditionals and generalizations containing attitude ascriptions. Consider (28).

(28) If *Nancy* wants Twain to autograph her book, and she believes that (if she waves then Twain will autograph her book), then, other things being equal, she will wave.

Imagine that Nancy believes and desires the relevant propositions in matching ways (so issues about Russellianism are irrelevant). Nevertheless, (28) might be true in some contexts and false in others. Imagine a context in which the speakers suspect that Nancy is considering employing a method other than waving for attracting Twain's attention. Perhaps Nancy is considering yelling "Please sign my book, Mr. Twain!" instead of waving, and our speakers know this. They think that Nancy prefers waving over yelling, but may not be sure. In such a context, they might judge that (28) is false but (29) is true.

- (29) If Nancy wants Twain to autograph her book, and she believes that (if she waves then Twain will autograph her book), *and she does not prefer an alternative way to catch Twain's attention*, then, other things being equal, she will wave.

Thus it seems that (28) is context-sensitive: it is true with respect to some contexts, false with respect to others, even keeping all facts about Nancy fixed.

8. Modifying the Prediction Objections

The Nomological Sufficiency Analysis and the Predictive Argument Analysis used *ceteris paribus* conditionals and generalizations (the latter in ideal predictive arguments) to analyze "can be used to predict", but they did not take into account the context-sensitivity of those sentences. Consider the following instantiation of the Nomological Sufficiency Analysis to sentence (1).

- (1) can be used to predict that Lucy will wave iff: it's nomologically necessary that if (1) is true, then, other things being equal, Lucy will wave.

The *ceteris paribus* conditional on the right-hand side varies in truth value from context to context, even keeping the facts about Lucy fixed. And so, according to the analysis, the truth value of the left-hand side varies from context to context, even if the facts about Lucy are kept fixed. Hence this analysis is "contextually shifty" in a way similar to the following (silly) analysis of intelligence.

- x is intelligent iff x 's IQ is as high or higher than my IQ.

Different theorists who employ this last "analysis", without realizing the context-sensitivity of the right-hand side, will come to different conclusions as to whether 'Lucy is intelligent' is true.

Similar phenomena arise for the Predictive Argument Analysis: the general-premises of an ideal predictive argument like (21) vary in truth value from context to context. So the analysis entails that '(1) can be used to predict that Lucy will wave' is true in some contexts, false in others.

All of this seems contrary to the intent of the critics of Russellianism. To do justice to them, we should try to modify the analyses of S can be used to predict P so that they are not themselves context-sensitive, while at the same time taking the context-sensitivity of *ceteris paribus* conditionals and generalizations into account. The best way to do this (I believe) is to existentially generalize over contexts to get analyses of the following form.

S can be used to predict P iff: *there is a context c such that. . . S . . . c *

I propose the following modified analyses.

(30) *The Modified Nomological Sufficiency Analysis*

S can be used to predict P iff: there is a context c and a sentence S such that (i) S expresses P in c , and (ii) it's nomologically necessary that if S , then, other things being equal, P is true in c .

(31) *The Modified Predictive Argument Analysis*

S can be used to predict P iff: there is some context c such that S is the sole particular-fact premise in some ideal predictive argument whose general-premises are true in c and whose conclusion expresses P in c .

The main ideas behind these modifications can be seen by applying them to (1) and Lucy's wave. (24) is true in some contexts, namely those in which the suitable conditions associated with (24) are strong enough to assure that if Lucy were to believe and desire the mentioned propositions, and those suitable conditions held, then she would wave. (Those suitable conditions include the

conditions that Lucy not prefer alternative ways to catch Twain's attention, that she not be paralyzed, and so on.) Thus, if (1) is also true in such a context, then Lucy will wave. So the existence of such a context seems to be a reasonable sufficient condition for the claim that (1) can be used to predict Lucy's wave. And anyone who found the original Nomological Sufficiency Analysis attractive should admit that the Modified version states a necessary condition. Thus analysis (30) seems reasonable as an analysis of "can be used to predict". Similarly, (21b) is true in some contexts, namely those where the *ceteris paribus* phrase of generalization (21b) excludes all interferences to the waving of an agent who believes and desires the mentioned propositions. Thus if (1) is also true in such a context, then Lucy will wave. So the above Modified analyses preserve the spirit of their originals, but also take context-sensitivity into account.

Critics of Russellianism who accept the Modified Analyses might want to add that there are *many* contexts that, under these Modified Analyses, suffice to show that (1) can be used to predict Lucy's wave. They might claim (for instance) that there are many contexts in which (24) is true, and many contexts in which the general premises of predictive argument (21) are true. They might claim that among these contexts are those in which ordinary speakers *think* that (24) is true, or think that the generalizations that appear in simple ideal predictive arguments are true; and also contexts in which some ordinary speaker uses (1) to predict that Lucy will wave.²⁷

Given these Modified Analyses, the most straightforward modifications of the earlier objections are the following.

(32) *The Modified Predictive Argument Objection*

- a. (1) can be used to predict Lucy's wave.

- b. If (1) can be used to predict Lucy's wave, then some ordinary psychological generalization is true *in some context*.
- c. Therefore, some ordinary psychological generalization is true in some context.
- d. If Russellianism is true, then no ordinary psychological generalization is true *in any context*.
- e. Therefore, Russellianism is not true.

(33) *The Modified Nomological Sufficiency Objection*

- a. (1) can be used to predict Lucy's wave.
- b. If (1) can be used to predict Lucy's wave, then (24) is true *in some context*.
- c. Therefore, (24) is true in some context.
- d. If Russellianism is true, then (24) is not true *in any context*.
- e. Therefore, Russellianism is not true.

9. Russellian Replies to the Modified Objections

In reply to these Modified Objections, I deny (32d) and (33d).²⁸ Some ordinary psychological generalization is true in some context, even if Russellianism is true; in fact, (21b) is true in some contexts. (24) is true in some contexts, even if Russellianism is true. In fact, I claim that (24) and (21b) are true in *many* contexts under Russellianism, including (nearly) all of those contexts in which the critics think they are true.

To see that this reply is plausible, let's begin with (24).

(24) It's nomologically necessary that: if Lucy wants Twain to autograph her book, and

she believes that (if she waves then Twain will autograph her book), then, other things being equal, she will wave.

Consider a context c in which the speakers are *not* considering the possibility that Lucy might become suddenly paralyzed, or that she might prefer some way of catching Twain's attention other than waving, and so on. So the suitable conditions associated with (24) in such a context include the conditions that Lucy not be paralyzed, not prefer alternative means, and so on. The critics of Russellianism would agree that (24) is true in such a context.

Now let's suppose further (at least for the moment) that the suitable conditions associated with (24), in this context c , also require that Lucy believe and desire in *matching ways*. Call this the *matching-ways condition*. Now consider a nomologically possible world in which Lucy believes and desires in *mismatching ways*, and consider the truth value of the conditional embedded in (24) with respect to this world and context c . Under Russellianism, the antecedent of the conditional is true in this world, with respect to c , but she does not wave. Yet the conditional is not false, with respect to this world and context c , because Lucy does not believe and desire in matching ways at this world, and so does not satisfy the suitable conditions associated with (24) in c . In fact, the conditional is true, with respect to c , in any nomologically possible world in which Lucy believes and desires the propositions in matching ways, and all other suitable conditions associated with (24) in c hold; for in all such worlds, she waves. Therefore, *if* there are contexts in which the suitable conditions associated with (24) include the matching-ways condition, then (24) is true, with respect to these contexts, even if Russellianism is correct.

Now I say that the matching-ways condition *is* part of the suitable conditions associated

with (24) in some contexts. It is part of the suitable conditions associated with (24) in (nearly) all contexts in which ordinary, non-Russellian, non-philosophical speakers consider (24), including most contexts in which the critics think that (24) is true. I have three reasons for thinking this is so.

First, the matching-ways condition typically holds in cases that ordinary speakers think satisfy the antecedent of (24). But speakers' dispositions to think of certain sorts of cases as "typical" or "normal", in a context, seem to help determine what counts as suitable conditions, in a context. So we have reason to think that the suitable conditions associated with (24) in ordinary contexts includes the matching-ways condition. Now a critic might object that ordinary speakers are ignorant of ways of taking propositions, and their matching and mismatching, and so might claim that the matching-ways condition cannot be part of the suitable conditions associated with (24), even if it holds in most cases that they think satisfy the antecedent of (24). But ignorance of this sort cannot prevent a condition from being part of the suitable conditions associated with a conditional in a context. Recall that most speakers in most contexts don't know that cars have solenoids, and yet the connection of a solenoid with a starter is part of the suitable conditions associated with (26) in many contexts containing such ignorant speakers.

Second, the matching-ways condition strongly resembles other conditions that are clearly a part of the suitable conditions associated with (24) in many ordinary contexts. Consider the requirement that Lucy is performing means-end reasoning in a normal way. This requirement seems to be part of the suitable conditions associated with (24) in many contexts. But the matching-ways condition is similar to this requirement; both concern conditions that allow reasoning to proceed in typical ways.

Third, if the matching-ways condition is part of the suitable conditions in *any* context with ordinary speakers, then it is unlikely to be excluded from the suitable conditions in other contexts containing only ordinary speakers. For recall that a condition that is part of the suitable conditions associated with a conditional in one context is (usually) excluded from the suitable conditions in another context when the speakers in the second context are considering the possibility that the condition doesn't hold. (This was the point of my example of my walking to the parking lot while considering whether Carl's battery is dead.) Now hardly any ordinary speakers think about ways of believing and desiring propositions. So there are (virtually) no contexts in which ordinary speakers are considering (24) and considering the possibility that Lucy might believe and desire the relevant propositions in mismatched ways. So if the matching-ways condition is part of the suitable conditions in at least one context with only ordinary speakers, it's very likely to be so in all such contexts.

Thus it's quite likely that the suitable conditions associated with (24) in many ordinary contexts includes the matching-ways condition. But as I argued above, if this is so, then (24) is true in some contexts, even if Russellianism is true.²⁹

Similar points hold about generalizations like (21b).

(21b) If a person wants Twain to autograph her book, and she believes that (if she waves then Twain will autograph her book), then, other things being equal, she will wave.

In many contexts, the suitable conditions associated with (21b) include the requirement that the agent believe and desire the propositions in matching ways. Mary (who actually believes and desires the propositions in mismatching ways) does not satisfy the suitable conditions associated

with (21) in such contexts; therefore, with respect to such contexts, Mary does not show that (21b) is false. In fact, the generalization is true with respect to such ordinary contexts.

Therefore, the matching-ways condition is a part of the suitable conditions associated with (24) and (21b) in some ordinary contexts. So (24) and (21b) are true with respect to some contexts, even if Russellianism is true. Thus premises (32d) and (33d) of the Modified Predictive Argument Objection and the Modified Nomological Sufficiency Objection are false, and the objections fail.

The context-sensitivity of *ceteris paribus* phrases may also help explain why anti-Russellians *think* that (24) and (21b) are false in all contexts, under Russellianism. There are some contexts in which the matching-ways condition is *not* part of the suitable conditions associated with (24) or (21b): for example, contexts in which philosophically sophisticated anti-Russellians are arguing against Russellianism. When anti-Russellians argue against Russellianism, they usually raise the possibility that an agent believes and desires the relevant propositions in mismatching ways. They thus create a "sophisticated" context in which the suitable conditions associated with (24) no longer include the requirement that Lucy believe and desire the relevant propositions in matching ways. A nomologically possible world in which Lucy believes and desires the propositions in mismatching ways is a world that *does* satisfy the suitable conditions determined by (24), with respect to such a philosophically sophisticated context. Thus (24) is false with respect to such a sophisticated, philosophical context. So (33d) may seem true to anti-Russellians who are thinking about mismatching ways while they are objecting to Russellianism, and are not paying attention to the context-sensitivity of *ceteris paribus* phrases. Similar points hold for generalization (21b) and premise (32d).

10. The Predictive Proposition Objection

Some readers who sympathize with the Prediction Objections may now be thinking that I got off on the wrong track a long time ago. These readers may be convinced that my discussion of *ceteris paribus* conditionals and generalizations is largely irrelevant to the anti-Russellians' real concerns. After all, anti-Russellians are worried about *attitude ascriptions* under Russellianism. Whatever their worries amount to, they surely have little or nothing to do with the semantic peculiarities of *other* English sentences, such as *ceteris paribus* conditionals and generalizations. The strange semantic properties of these latter sentences should not show that the critics' concerns about prediction and attitude ascriptions, under Russellianism, are incorrect.

In reply to these readers' concerns, I want first to point out that I was led into a discussion of *ceteris paribus* conditionals and generalizations by my efforts to provide charitable interpretations of the anti-Russellians' objections, and what they might mean by "can be used to predict". I believe that many anti-Russellians would make use of *ceteris paribus* sentences, and predictive arguments containing them, to explain and bolster their Prediction Objections. So I could not have simply ignored these sentences. Nevertheless, I strongly sympathize with the above worries: it seems to me that there should be some issue about Russellianism, attitude ascriptions, and prediction that is independent of the semantic peculiarities of *ceteris paribus* sentences. I shall now try to formulate such a worry.

According to typical *anti*-Russellian views, sentence (1) expresses a proposition that is partly about the *ways* in which Lucy believes and desires the relevant propositions.

- (1) Lucy wants Twain to autograph her book, and she believes that if she waves, then Twain will autograph her book.

On these views, the proposition expressed by (1) says (or entails) that the ways in which Lucy believes and desires match. By contrast, Russellianism says that (1) does *not* express any proposition about ways of believing and desiring. Now take some typical anti-Russellian view, and the proposition (1) semantically expresses according to it; call this proposition 'AR' (for 'anti-Russellian'). Take the proposition that (1) semantically expresses according to Russellianism, and call it 'R'. Intuitively speaking, proposition AR contains *more information* than proposition R, namely information about ways of believing and desiring. This extra information is nomologically relevant to whether Lucy will wave, for as everyone concedes, if she believes and desires in matching ways, she will wave (or is "likely" to), but if not, not.

A critic might claim that some propositions contain enough information that is nomologically relevant to Lucy's wave to count as *predictors* for Lucy's wave. AR is one of these. Others do not contain enough information that is nomologically relevant to Lucy's wave to count as predictors of her wave. One obvious example is the proposition that Bill Clinton sneezes at time *t* (assuming that Clinton's sneeze is causally isolated from Lucy's wave). But R is another such proposition, claims the critic. A *sentence* can be used to predict Lucy's wave iff it semantically expresses a predictor (a predictive proposition) for Lucy's wave. Sentence (1) expresses a predictive proposition for Lucy's wave if anti-Russellian views are correct, whereas (1) does not express a predictive proposition, if Russellian views are correct. So (1) can be used to predict Lucy's wave if anti-Russellian views are correct, but cannot be so used if Russellian views are correct. But (1) clearly can be so used, so Russellianism is false. Call this the *Predictive Proposition Objection*.

I believe that this objection contains some useful suggestions about "can be used to

predict". But it says little to support the claim that *R* is not a predictor of Lucy's wave. Our two earlier objections tried (at least in effect) to distinguish between predictors and non-predictors by appealing to *ceteris paribus* sentences; the basic idea was that predictors can be combined with the (propositions expressed by) *ceteris paribus* conditionals and generalizations to support the conclusion that Lucy waves. But we saw that the context-sensitivity of *ceteris paribus* phrases makes this way of distinguishing predictors from non-predictors more complicated than it initially appeared; and there is reason to doubt that the facts about these sentences favor the anti-Russellians. So to help the anti-Russellians, we need to find another method for distinguishing predictors from non-predictors.

But in fact, any reasonable method for distinguishing predictors from non-predictors will entail that *R* is a perfectly good predictor for Lucy's wave. To see this, it will be useful to return (again) to Carl, and to think more generally about *predictive information*.³⁰

11. Predictive Information

Consider the nomologically possible worlds in which Carl's engine starts. In each of these worlds, the starting of Carl's engine has a causal history; at any time before the starting, there is a class of *all* of the events that occur at that time that are causes of the engine's starting, in that world.³¹ Given the occurrence of the events in any such class, the starting of Carl's engine is nomologically inevitable.³² (Notice that I do *not* need to add a *ceteris paribus* phrase here; all of the "other things" that "need to be equal" are already members of the class.³³) Now consider a proposition that describes every event in one such class of events, and which is entirely about those events. Such a proposition is admissible with respect to the proposition that Carl's engine

starts. Moreover, for any nomologically possible world, if such a proposition is true at that world, then Carl's engine starts at that world. Such a proposition is *ideal for predicting* the proposition that Carl's engine will start. More generally:

- (34) If P and Q are propositions, and Q describes an event, then P is *ideal for predicting* Q iff: (i) P is a true proposition that is admissible for predicting Q ; (ii) for any nomologically possible world w , if P is true in w then Q is true in w .

We can also say that P is an *ideal predictor* for Q , or an *ideal predictive proposition* with respect to Q .³⁴ If E is an event that Q describes, then we can also say that P is ideal for predicting event E .

Some of the propositions that are ideal for predicting a starting of Carl's engine describe a turning of Carl's ignition key to its start position. Thus the proposition that Carl is a car whose ignition key is being turned to its start position *contains some, but not all, of the information contained* in these ideal predictive propositions. So the proposition that Carl is a car whose ignition key is being turned to its start position *provides some predictive information* concerning the starting. More generally,

- (35) If P and Q are propositions, and Q describes an event, then P *provides predictive information concerning* Q iff: P contains some of the information contained in some proposition that is ideal for predicting Q .

We can similarly say that P provides predictive information concerning an event E that Q describes. We can also say something similar about *sentences*, as opposed to propositions. Sentence (15) semantically expresses a proposition that provides predictive information concerning the proposition that Carl's engine will start. Thus we could say that sentence (15)

itself provides predictive information. More generally,

- (36) If S is a sentence and Q is a proposition that describes an event, then S provides predictive information concerning Q iff: S semantically expresses a proposition that provides predictive information concerning Q .

It's intuitively clear that some sentences provide more predictive information concerning a particular event than others. Compare (15) and (37).

- (15) Carl is a car whose ignition key is being turned to its start position.
- (37) Carl is a car whose ignition key is being turned to its start position and whose battery is live.

At least one ideal predictor for the starting describes both the turning and the live state of the battery. (37) provides more of the information contained in this proposition than does (15).³⁵

Can any finite English sentence express *all* of the information contained in a proposition ideal for predicting that Carl's engine will start? Perhaps sentences that contain phrases like 'other things being equal' can. For example, 'Carl is a car whose ignition key is being turned to its start position and all other things are equal' may, in some context, express a proposition that is ideal for predicting a starting of Carl's engine. But if it does, it expresses such a proposition in an "inarticulate" and context-sensitive way. Sentences that lack a *ceteris paribus* phrase are likely to provide only *some* of the information contained in an ideal predictive proposition.

In light of these considerations, how should we analyze "S can be used to predict P"? The most straightforward proposal is (38).

- (38) S can be used to predict P iff: S provides predictive information concerning P .

(15) provides predictive information concerning the proposition that Carl's engine will start, so

(38) entails that (15) can be used to predict that Carl's engine will start.

Analysis (38), however, is quite liberal. According to it, any sentence describing any cause of an event, no matter how remote or "partial" the cause, counts as a sentence that can be used to predict that the event will occur. Consider, for example, the sentence 'The solar system formed'. It expresses a proposition about one of the distal causes of the engine's starting. There is a true, admissible proposition that describes the formation, and all the other events that occurred at that time that are in the causal history of the engine's starting and are (together) nomologically sufficient to cause the (eventual) starting. This latter proposition is ideal for predicting the starting. The sentence about the formation provides some information contained in this ideal proposition. So analysis (38) entails that it can be used to predict that Carl's engine will start. Similar points hold for sentences that describe more proximal causes, but give little information about them, for instance, 'Carl has a solenoid' and 'Carl has a metallic part'.

We might try to craft a more restrictive informational analysis of "can be used to predict" in terms of *how much* predictive information the sentence provides. Presumably, the analysis would have a form like the following:

(39) S can be used to predict P iff: S provides at least amount M (or percentage X) of predictive information contained in some proposition ideal for predicting P .

But I suspect that precisely measuring amounts of predictive information will be extremely difficult or impossible. Furthermore, even if it is possible, I doubt that any choice of M will seem intuitively correct. After all, for any choice of M there will be some proposition that provides a little less than M amount of predictive information about P , but which will still look rather predictive. Thus I am inclined to think that we should not strive for an informational analysis of

"can be used to predict" that is more restrictive than (38). We should accept (38), and settle for saying that some sentences provide more predictive information, while others provide less (with respect to a given P).

However, I need not rely on the liberal standards of (38) in my defense of Russellianism from the Predictive Proposition Objection. I can instead rely on an analysis that (I think) will be more acceptable to those who find (38) too liberal. Intuitively, it seems that (15) provides a *substantial* amount of predictive information regarding the starting, whereas 'The solar system formed' does not. Thus a person who rejects the liberal analysis might find (40) more acceptable.

(40) *The Substantial Predictive Information Analysis*

S can be used to predict P iff: S provides a *substantial* amount of the information contained in some proposition that is ideal for predicting P .

We can say that the *proposition* expressed by S is a *predictor* (or a *predictive proposition*) for P (or for some event E that P describes) iff S provides a substantial amount of the information contained in some propositions that is ideal for predicting P .

The expression 'substantial' in (40) is vague, of course. I can make it a bit more precise by stipulating that (as I use the term) (15) does provide a substantial amount of information contained in some proposition ideal for predicting that Carl's engine starts, whereas the sentences 'The solar system formed' and 'Carl has a metallic part' do not. Furthermore, suppose that (15) provides a certain fraction of the information contained in some proposition ideal for predicting the starting; then any sentence that provides the same fraction of information from some proposition that is ideal for predicting P shall also count as providing a substantial amount of the information from that ideally predictive proposition. These stipulations should seem reasonable,

given our judgments about the predictive status of (15).

12. Reply to the Predictive Proposition Objection

With these considerations and analyses in mind, let's return to attitude ascriptions and the imaginary critic's objection to Russellianism.³⁶ Consider again the propositions that (1) expresses according to Russellianism and a representative anti-Russellian theory, namely *R* and *AR*, respectively.

- (1) Lucy wants Twain to autograph her book, and she believes that if she waves, then Twain will autograph her book.

The critic claims that *AR* provides enough predictive information to be a predictor for Lucy's wave, whereas *R* does not. Let's examine this claim in light of the above considerations.

Neither *R* nor *AR* is ideal for predicting Lucy's wave, for there are nomologically possible worlds in which *R* is true and Lucy fails to wave, and worlds in which *AR* is true and Lucy fails to wave. Nevertheless, both contain some of the information contained in some proposition that is ideal for predicting Lucy's wave; so both provide some predictive information. *AR* contains a bit more predictive information than *R*. Thus sentence (1) provides a bit more predictive information if anti-Russellianism is true than if Russellianism is true. But according to our Substantial Predictive Information Analysis of "can be used to predict", (1) can be used to predict that Lucy will wave iff it provides a substantial amount of the information contained in some proposition that is ideal for predicting that Lucy will wave. Thus, according to this analysis, the critic is right iff *AR* provides a substantial amount of information contained in some proposition ideal for predicting that Lucy will wave, whereas *R* does not.

Let's first consider the claim that *R* does not provide a substantial amount of the information contained in some proposition that is ideal for predicting that Lucy will wave. There's very little to be said in favor of this claim. Perhaps the best way to see how little merit it has is to compare ascription (1) with sentence (15). (15), everyone agrees, can be used to predict that Carl's engine will start. Therefore, we can use the proposition it expresses as a very rough standard for "measuring" whether other propositions provide substantial predictive information (as, indeed, I did in crafting our latest analysis). But proposition *R* and the proposition expressed by (15) (call it '*G*' for 'ignition') are informationally similar in many seemingly important respects. Proposition *G* describes a significant proximal cause of the starting and describes some significant properties of this cause. *R* also describes a significant proximal cause of Lucy's wave and some significant properties of it.³⁷ *R* and *G* are also similar in the ways in which they fall short of expressing ideally predictive propositions. *G*, for instance, fails to rule out the possibility of a dead battery; *R* fails to rule out the possibility that Lucy believes and desires in mismatching ways. But *R* does not seem to fall further short of its ideal predictive proposition(s) than does *G* from its. Overall, then, we can conclude that if *G* provides a substantial amount of predictive information, then so does *R*. Since *G* does provide a substantial amount of predictive information, so does *R*, and the critic's claim that *R* does not is incorrect.

The critic also claims that *R* fails to provide substantial predictive information, whereas *AR* does provide substantial predictive information. This conjunction of claims is also implausible. The difference in the information that *R* and *AR* provide is simply not significant enough to warrant the claim. The claim is no more plausible than the claim that (37) provides substantial predictive information about Carl's starting, whereas (15) does not.

(15) Carl is a car whose ignition key is being turned to its start position.

(37) Carl is a car whose ignition key is being turned to its start position and whose battery is live.

I conclude that the Predictive Proposition Objection fails to make a convincing case against Russellianism.

13. An Alternative Russellian Reply

Each Prediction Objection has claimed that (1) can be used to predict Lucy's wave, though different objections have relied on different analyses of, or different senses of, the phrase "can be used to predict". In my replies, I have accepted the claim that (1) can be used to predict Lucy's wave, in each of these senses. I then tried to show that some other premise of each objection is false. Along the way, I argued that attitude ascriptions can be used to predict behavior, in each of these senses, even if Russellianism is true.

My replies have relied on claims that might be reasonably doubted. For instance, I claimed that the suitable conditions associated with a *ceteris paribus* conditional like (24), in a context, sometimes include the condition that Lucy believe and desire the propositions in matching ways. I claimed that (1) provides a substantial amount of the predictive information contained in some proposition that is ideal for predicting Lucy's wave. These claims, I think, are highly plausible, but they are not indubitable. Therefore, I would like to consider whether Russellians can reply in any other way to these objections, should my claims about suitable conditions, etc., turn out to be wrong.

Let's suppose, then, that there are no contexts in which the suitable conditions associated

with a *ceteris paribus* conditional or generalization require that an agent believe and desire propositions in matching ways. Let's suppose furthermore that under Russellianism (1) falls short of providing a substantial amount of predictive information. Thus, the standards for "can be used to predict" are higher than I've been assuming thus far; consequently, my attempts to criticize the Prediction Objections have failed (thus far).

Supposing all of this, I think that there is one more plausible reply that a Russellian can make: he can deny that (1) can be used to predict Lucy's wave, in any of the ("high standard") senses explained above. Instead, it merely *appears* that (1) can be used to predict Lucy's wave, in these senses. Of course, a Russellian is obligated to explain how it can appear that (1) can be used to predict Lucy's wave, in the above senses, even though (1) cannot, in fact, be so used. But I think Russellians can do this (if need be). The explanation of the appearance will rely heavily on the fact that ordinary speakers, and other non-Russellians, think that (1) is false when Lucy believes and desires the propositions in mismatching ways. Let me explain.

According to the Modified Nomological Sufficiency Analysis of "can be used to predict", (1) can be used to predict Lucy's wave only if (24) is true in some context.

(24) It's nomologically necessary that: if Lucy wants Twain to autograph her book, and she believes that (if she waves then Twain will autograph her book), then, other things being equal, she will wave.

Now we are supposing that (24) is false in every context if Russellianism is true. The reason is that if Russellianism is true, then there is some nomologically possible world in which Lucy believes and desires the relevant propositions in mismatching ways, and the suitable conditions of the context hold (despite the fact that Lucy believes and desire in mismatching ways), and yet

the consequent is false.

But notice that even if Russellianism is true, and (24) really is false in every context, nevertheless ordinary speakers (and other non-Russellians) would tend to think that (24) is true, in some contexts. For if they consider a world (a "situation") in which Lucy assents to (2), and so believes and desires the propositions in mismatching ways, they would think that the antecedent of (24) is *false* in this situation (in their context).

- (2) If I wave, then Clemens will autograph my book, and Twain won't. I want Twain to autograph my book, and not Clemens.

Thus the falsity of the consequent in the situation wouldn't convince them that the whole conditional is false. They might think: "In this situation, Lucy doesn't wave. But neither does she believe and desire the propositions mentioned in the antecedent of (24). It's still true that *if* she believed and desired those propositions (which she doesn't in this situation), then, other things being equal, she would wave." In short, even if Russellianism were true, and (24) were false in all contexts, nevertheless ordinary speakers (and other non-Russellians) would tend to think that (24) is true (at least in some contexts). Therefore, if the Modified Nomological Sufficiency Analysis really captures the sense in which ordinary speakers (and non-Russellians) use the phrase "can be used to predict", then such speakers would think that (1) can be used to predict Lucy's wave, even if Russellianism is true and (1) *cannot* be so used.

Similarly for ordinary psychological generalizations like (21b), and the Modified Predictive Argument Analysis of "can be used to predict". Even if generalization (21b) were false (in all contexts) under Russellianism, ordinary speakers would tend to think that it is true (in some contexts), for they would not recognize that counterexamples to it are really

counterexamples. Ordinary speakers would think that Mary (who actually believes the propositions in mismatching ways) fails to satisfy the antecedent of (21b). So they wouldn't think that she is a counterexample to the generalization, even if she were. Furthermore, the people that they think satisfy the antecedent are people who believe and desire the propositions in matching ways. Those people tend to wave, other things being equal. So ordinary speakers and other non-Russellians (in many contexts) would tend to think that generalization (21b) is true, even if Russellianism is true and (21b) is false. Therefore, if the Modified Predictive Argument Analysis captures the sense in which ordinary speakers (and other non-Russellians) use the phrase "can be used to predict", then they would tend to think that (1) can be used to predict that Lucy will wave, even if (1) cannot be so used.

According to the Substantial Predictive Information Analysis, (1) can be used to predict that Lucy will wave iff (1) provides a substantial amount of the information contained in some proposition that is ideal for predicting that Lucy will wave. I am supposing now that (1) falls short of providing a substantial amount of such predictive information, if Russellianism is true, because the proposition it expresses fails to provide any information about the ways in which Lucy's believes and desires the propositions. But ordinary speakers have a strong tendency to think that (1) is true only when Lucy believes and desires the relevant propositions in matching ways. Therefore, ordinary speakers have a strong tendency to think that (1) is true only when proposition *AR* is true. *AR* does provide substantial predictive information, according to anti-Russellians. But if *AR* is true whenever ordinary speakers think that (1) is true, and *AR* does provide substantial predictive information, then it's easy to see how ordinary speakers might mistakenly come to think that (1) itself provides substantial predictive information. Therefore, if

the Substantial Predictive Information Analysis captures ordinary speakers' way of using "can be used to predict", then ordinary speakers will tend to think that (1) can be used to predict Lucy's wave, even though (1) itself does not provide substantial predictive information.

All of these replies leave Russellians with an obligation to explain why ordinary speakers think that (1) is false in mismatch circumstances, when Russellianism says that (1) is, in fact, true in such circumstances. But Russellians have already offered explanations of such ordinary judgments. According to many Russellians, an ordinary speaker who is informed that Lucy is disposed to assent to (2) will (typically) come to believe the *negation* of the proposition that (1) expresses, in a way that corresponds to the negation of sentence (1). His believing this proposition in this way explains why he thinks that sentence (1) is false. (See Braun 1998.) Salmon, Soames, and McKay might add that utterances of (1) pragmatically convey the proposition that Lucy would assent to "If I wave, then Twain will autograph my book". Our ordinary speaker believes that this conveyed proposition is false. Moreover, he confuses it with the proposition expressed by (1), and so judges that (1) itself is false.³⁸

14. Conclusion

I intend this paper to be a contribution to the long-term project of assessing Russellianism (a project involving many philosophers besides myself).³⁹ One important part of this project is the assessment of objections to Russellianism, along with replies to those objections. I here considered objections that say (roughly) that if Russellianism were true, then attitude ascriptions could not be used to predict behavior. I have formulated what I take to be the strongest objections of this type, using as my starting points the brief remarks by critics of Russellianism

who mention prediction. I have tried to present the strongest replies to these objections. In my opinion, the replies are more persuasive than the objections. Thus, in my opinion, the Prediction Objections do not constitute a convincing case against Russellianism.⁴⁰

References

Braun, David. 1998. "Understanding Belief Reports." *Philosophical Review* 107, pp. 555-595.

Braun, David. 2000. "Russellianism and Psychological Generalizations." *Nous* 34, pp. 203-36.

Braun, David. Forthcoming. "Russellianism and Explanation." *Philosophical Perspectives*.

Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Oxford University Press.

Crimmins, Mark. 1992. *Talk About Beliefs*. Cambridge, MA: MIT Press.

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

Devitt, Michael. 1997. "Meanings and Psychology: A Response to Mark Richard." *Nous* 31, pp. 115-31.

Fodor, Jerry. 1991. "You can Fool Some of hte People All of the Time, Everything Else Being Equal: Hedged Laws and Psychological Explanation." *Mind* 100, pp. 19-33.

Forbes, Graeme. 1990. "The Indispensability of Sinn." *Philosophical Review* 99, pp. 337-72.

Hempel, Carl G. 1965a. *Aspects of Scientific Explanation and Other Essays in Philosophy of Science*. New York: Free Press.

Hempel, Carl G. 1965b. "Aspects of Scientific Explanation." In Hempel 1965a, pp. 331-496. Portions reprinted in Ruben 1993, pp. 42-55.

Hempel, Carl G., and Oppenheim, Paul. 1965. "Studies in the Logic of Explanation." In Hempel 1965a, pp. 245-290. Originally appeared in *Philosophy of Science* 15 (1948), pp. 135-75.

Laymon, R. 1989. "Cartwright and the Lying Laws of Physics." *Journal of Philosophy* 86, pp. 353-372.

Lewis, David. 1986. "Causal Explanation." In his *Philosophical Papers, Volume II*, pp. 214-40. Oxford: Oxford University Press. Reprinted in Ruben 1993, pp. 182-206.

McKay, Thomas. 1979. "On Proper Names in Belief Ascriptions." *Philosophical Studies* 39, pp. 287-303.

Montague, Richard. 1974. "The Proper Treatment of Quantification in Ordinary English." In Montague (Richmond Thomason, ed.), *Formal Philosophy*, pp. 247-270. New Haven: Yale University Press. Originally appeared in J. Hintikka, J. Moravcsik, and P. Suppes (eds.),

Approaches to Natural Language, pp. 221-42. Dordrecht: Reidel, 1973.

Morreau, Michael. 1997. "Fainthearted Conditionals." *Journal of Philosophy* 94, pp. 187-211.

Morreau, Michael. 1999. "Other Things Being Equal." *Philosophical Studies* 96, pp. 163-182.

Pietroski, Paul and Rey, Georges. 1995. "When Other Things Aren't Equal: Saving *Ceteris Paribus* Laws from Vacuity." *British Journal of Philosophy of Science* 46, pp. 81-110.

Railton, Peter. 1981. "Probability, Explanation, and Information." *Synthese* 48, pp. 233-56.
Reprinted in Ruben 1993, pp. 160-181.

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

Richard, Mark. 1997a. "Propositional Attitudes." In Bob Hale and Crispin Wright (eds.), *A Companion to Philosophy of Language*, pp. 197-226. Oxford: Blackwell.

Richard, Mark. 1997b. "What Does Commonsense Psychology Tell Us About Meaning?" *Nous* 31, pp. 87-114.

Ruben, David-Hillel (ed.). 1993. *Explanation*. Oxford: Oxford University Press.

Salmon, Nathan. 1986. *Frege's Puzzle*. Cambridge, MA: MIT Press.

Salmon, Nathan. 1989. "Illogical Belief." *Philosophical Perspectives* 3, pp. 243-285.

Salmon, Nathan. 1992. "Reflections on Reflexivity." *Linguistics and Philosophy* 15, pp. 53-63.

Schiffer, Stephen. 1991. "Ceteris Paribus Laws." *Mind* 100, pp. 1-18.

Silverberg, Arnold. 1996. "Psychological Laws and Non-Monotonic Logic." *Erkenntnis* 44, pp. 199-224.

Soames, Scott. 1988. "Direct Reference, Propositional Attitudes, and Semantic Content." In Nathan Salmon and Scott Soames (Eds.), *Propositions and Attitudes*, pp. 197-239. Originally appeared in *Philosophical Topics* 15 (1987), pp. 47-87.

Soames, Scott. 1995. "Beyond Singular Propositions?" *Canadian Journal of Philosophy* 25, pp. 515-550.

Soames, Scott. 1994. "Attitudes and Anaphora." *Philosophical Perspectives* 8, pp. 251-72.

Notes

1. Here's one reason to think that the issues concerning prediction require separate treatment from those concerning explanation. In Braun (forthcoming), I argue that S explains event E if S describes some cause of E , or S provides a (substantial) portion of information from an ideal explanatory argument that explains E . This theory turns out to be rather liberal about what counts as an explanation of E . Therefore, one might reasonably wonder whether every S that counts as explaining E under this theory could really be used to *predict* E . This at least needs to be argued, but arguing for it requires some discussion of the conditions under which a sentence can be used to predict an event.

2. I am ignoring time and tense as much as possible throughout this paper. A Russellian who took them into account might say that these sentences express a proposition, with respect to a context, that has a time (interval) as a constituent. See Salmon 1986. Many of the sentences I discuss below have no context-sensitive expressions other than tense expressions. Thus I often ignore context-sensitivity below, and speak imply of the proposition expressed by a sentence.

3. Again, I am ignoring time and tense (and tenselessness) here.

4. Thus underlying the binary belief relation there is a ternary relation BEL that holds between agents, propositions, and ways of taking propositions. (See Salmon 1986.) An agent A believes proposition P iff $\pm \Box \text{BEL}(A, P, w)$. More informally: A believes P in some way or other. Similarly, there is a ternary relation DES that can hold between an agent, a proposition, and a way of taking a proposition; A desires P iff $\pm \Box \text{DES}(A, P, w)$. Informally: A desires P in some way or other.

5. I often use numerical indices for sentences as abbreviations for their quotation names, especially in 'that'-clauses. For instance, I use (i) as an abbreviation of (ii).

(i) Karen believes that (10) is true.

(ii) Karen believes that 'Twain is an author' is true.

Strictly speaking, there are different ways in which a single agent, like Karen, can believe that 'Twain is an author' is true. For instance, Karen could have a name for the sentence 'Twain is an author', say 'Fred'. Then (if the content of a quotation-name is simply the expression to which it refers) she could believe that 'Twain is an author' is true, in at least two ways: a "'Twain is an author'" -ish way and a "Freddish" way. Or she could believe the proposition that 'Twain is an author' is true in one way, while failing to believe it in the other way.

6. This is a bit of an oversimplification. On Richard's (1990), Crimmins's (1992), and Forbes's (1990) views, different utterances of the same attitude ascription can express different propositions. On their views, some utterances of (7) express propositions that provide information about the ways in which Karen believes the proposition, whereas others do not. According to Devitt (1996), attitude ascriptions relate agents to properties of mental representations tokens; the ascriptions are ambiguous between transparent and opaque readings. The opaque reading of (7) provides some information about Karen's mental representations that

resembles what I would call "information about the way in which Karen believes the proposition that Twain is an author".

7. I am making the following assumptions about Lucy and the ascriptions in (1). First, if Lucy sincerely asserts a sentence that contains 'I', then she believes the proposition expressed by the sentence in a *first-person way*. For instance, she believes the proposition by having in her head a mental representation containing something like a first-person pronoun. Second, I assume that if she believes, in a first-person way, the proposition that she wants Twain to autograph her book, then she really does want Twain to autograph her book. I make similar assumptions about all of the characters that appear in this paper. I also assume that, under Russellianism, the contents of the (anaphoric) pronouns 'she' and 'her' in (1) are Lucy herself; thus, (1) ascribes to Lucy belief and desire in singular propositions that contain Lucy herself as a constituent. Russellianism is compatible with other analyses of such (anaphoric) pronouns (see Salmon 1992 and Soames 1994), but the issues surrounding these analyses are orthogonal to my concerns here.

8. Sentences like (13), in which the syntactic object of 'predicts' is an NP, raise some difficult semantic problems. For instance, it appears that (13) can be true even if Mary never waves; thus it can be true even if the NP 'Mary's wave' has no denotation. (Thanks to John Bennett for pointing this out to me.) Similarly, a sentence like 'Joe predicted a hurricane in Rochester' can be true when no hurricane in Rochester ever occurs. Moreover, this sentence may be true, while the sentence 'Joe predicted an earthquake in Des Moines' is false, even though the predicates 'hurricane in Rochester' and 'earthquake in Des Moines' have the same (empty) extension. Thus the verb 'predicts' seems to have at least some of the properties typical of *intensional transitive verbs*, such as 'seeks'. The semantics of such verbs is notoriously difficult. Many theorists (most famously, Montague 1974) hold that the NPs that appear in the objection positions of such verbs refer to properties (of either individuals or properties). That is why I occasionally say that such sentences relate agents to *types* of events. But I mostly try to skirt around this difficult semantic problem in what follows. Fortunately, the semantic issues surrounding intensional transitive verbs won't affect the issues I discuss here.

9. Russellians and anti-Russellians would almost certainly disagree over the semantics of prediction-sentences with sentential complements. (If so, they would almost certainly also disagree about the semantics of prediction-sentences with NP objects, but I won't discuss this further here.) Russellians would certainly hold that 'predict' expresses a binary relation; in addition, Russellians who accept the earlier metaphysical theory of belief would surely hold that agents always predict propositions *in some way or other*, via some way of taking the relevant proposition. They should say that a sentence of the form $\langle A \text{ predicts that } S \rangle$ can be metaphysically analyzed as $\langle \lambda w \text{ PRED}(A, \text{that } S, w) \rangle$, where 'PRED' is a ternary relation holding between agents, propositions, and ways of taking propositions. Such a Russellian should say that an agent can rationally predict both a proposition and its negation, as long as he does so in suitably different ways. For instance, an agent might predict that Twain will wave, while taking that proposition in a "Twainish" way, but predict that Twain will not wave, while taking it in a "Clemensish" way. Anti-Russellians might hold that the verb 'predicts' *semantically expresses* the ternary relation PRED; these theorists would dispute my claim that 'predicts' expresses a binary relation. Other

anti-Russellians might hold that the relation is binary, but that the propositional objects of the relation incorporate ways of taking propositions. This dispute between Russellians and anti-Russellians over the semantics of 'predict' won't affect what follows.

10. A person can predict propositions that may not describe events. E.g., Joe may predict that the United Nations won't exist in the year 3000, though it's questionable whether the embedded sentence describes an event. In what follows, all predicted propositions will (purport to) describe events.

11. A proposition is a predictive one iff someone stands in the predicting relation to it. This may sound strange, but compare it with 'belief' as a term for propositions. Not all propositions are beliefs. The beliefs (the believed propositions) are those to which someone bears the believing relation. They are beliefs iff someone stands in the believing relation to them.

12. English grammar allows us to state identities between predictions and propositions (denoted by 'that'-clauses), as in 'Joe's prediction is that Mary will wave'. Formulating a seemingly true identity between a prediction and an event is more difficult. If I say 'Joe's prediction is Mary's wave', I will probably be taken to be (wrongly) identifying Joe's *act* of predicting with Mary's act of waving. I can say 'The event that Joe predicted is Mary's wave', but this does not explicitly use the expression 'prediction'. Nevertheless, I shall say that some predictions (in the sense of "objects of prediction") are identical with certain events.

13. It's difficult to say when a proposition is admissible into a basis for predicting an event. I shall leave the notion unanalyzed. One might worry that I am travelling in a very tight circle here. To determine whether a person is *predicting* P , one must know whether her basis contains only propositions admissible with respect to P , but to know whether those basis propositions are admissible, one must determine whether it is legitimate to predict P using them. I suspect that the circle really is rather tight; nevertheless, I think the analyses I provide below are adequate for my purposes here.

14. Russellians would obviously hold that an agent may believe the propositions in basis B in certain ways but not in other ways. And as I mentioned before, many Russellians would hold that $\langle A \text{ predicts that } S \rangle$ can be metaphysically analyzed as $\langle \pm W \rangle \text{PRED}(A, \text{that } S, w) \langle \pm \rangle$, where PRED is an underlying ternary relation between agent, propositions, and ways of taking propositions. (See note 9.) These Russellians should hold that the inferring relation is a ternary relation that holds between an agent A and propositions P and Q : A infers Q from P . But such a Russellian should metaphysically analyze the ternary relation of inferring into a quintary relation: A infers Q , taken in way W_2 , from P , taken in way W_1 . An agent may infer Q from P , when taking Q and P in certain ways, but fail to infer Q from P , when taking them in certain other ways. For instance, an agent might infer the proposition that Twain will fall, taking that proposition in a "Twainish" way, from the proposition that Twain stepped on a banana peel, taking that proposition also in a "Twainish" way; but he may fail to infer that Twain will fall, taking that proposition in a "Clemensish" way, from the proposition that Twain stepped on a banana peel, taking that latter proposition in a "Twainish" way.

15. I do not intend this *metaphysical* analysis of $\leq A$ predicts E to serve as a *semantic* analysis of English sentences of the form $\leq A$ predicts NP , where NP might be an expression like 'Mary's wave' or 'a hurricane in Rochester'. I hope to skirt around the difficult semantic issues regarding the denotation of NPs in sentences of the form $\leq A$ predicts NP . See note 8.

16. The analysis, however, does *not* allow a person to predict an event E on the basis of propositions about events that occur after E , or any other propositions about times after E . If the physics student who knows that E occurred infers that E had earlier causes of such-and-such type, then he does not *predict* those earlier causing events. We can instead say that he *retrodicted* those causing events.

17. Strictly speaking, a Russellian should probably make some additions to this analysis. First, clause (ii) might be modified to the following effect: A thinks that S is true and believes the proposition expressed by S in an " S "ish way. Clause (iii) might be modified to say something like the following: the proposition expressed by S , taken in an " S "ish way, is a member of A 's basis for predicting P . Fortunately, none of these qualifications will affect what follows below.

David Hunter has pointed out to me that one may test a theory (which we can think of as a conjunction of sentences) by drawing consequences from it and determining whether those consequences are true. A person who does this might be said to use the theory to predict certain propositions, whether or not he believes the theory. Thus Hunter has suggested that my analysis of $\leq A$ uses S to predict P should not include clause (ii). (One might, for related reasons, wonder whether my prior analysis of $\leq A$ predicts P should include the requirement that A believe P and the basis propositions.) I am of two minds about whether to say that a person who tests a theory in this way really uses the theory to predict propositions. I am inclined to say that he merely draws out some consequences of the theory; alternatively, I'd say that there are two senses in which one can use a theory to make predictions, and I am interested in only one of these. Fortunately, this issue (again) does not affect the main issues with which I am concerned.

18. One might initially think that the critic's notion of "can be used to predict" has something to do with the *reasonableness* of agents who use the sentence to make predictions. In fact, Richard sometimes maintains that if Russellianism were true, then attitude ascriptions "would give us no reason" to believe that an agent will act a certain way (see Richard 1990, p. 176, and Richard 1997, p. 202). So one might propose the following analysis: S can be used to predict P iff it is possible for a reasonable agent to use S to predict P . The problem with this analysis is that agents can reasonably believe falsehoods which they can "combine" with the proposition expressed by S to predict just about any other proposition. For example, it seems to me that, in some possible world, the following generalization is false and yet some agent in that world reasonably believes it: if a person believes that (if she *waves* then Twain will autograph her book), and she wants Twain to autograph her book, then she will *yell*. But clearly (1) cannot be used to predict that Lucy will yell, at least not in the critics' intended sense of "can be used to predict". (Thanks to John Bennett for discussion of these matters.)

19. Notice that *ceteris paribus* conditionals are strong conditionals of some sort. Consider the conditional embedded in (16): neither the falsity of its antecedent at a world nor the truth of its

consequent at a world is sufficient for the truth of the whole at that world.

20. To derive this result, we must allow ourselves to use instances of the following schemas:

'S' is true iff *S*.

The proposition that *S* is true iff *S*.

21. See note 23 for passages in which Devitt and Richard discuss Russellianism's alleged problems with prediction at the same time that they mention its alleged problems with explanation. Devitt quite often mentions the generalizations that (he thinks) support explanations using attitude ascriptions: see Devitt 1996, pp. 174-5, 220-1, and 230-7; and Devitt 1997, pp. 117. Richard says that ordinary psychological explanations of behavior are "underwritten by" or "implicitly invoke" psychological generalizations: Richard 1990, pp. 260-3, and Richard 1997b, p. 90. He also speaks of the premises of belief-desire explanations: Richard 1990, pp. 176, 219. These passages indicate that Devitt and Richard are assuming some sort of covering-law theory of *explanation*. Since they closely associate prediction with explanation (see note 23), they may also hold a covering-law theory of *prediction*, and so they might accept the Predictive Argument Analysis of "can be used to predict".

22. I am assuming that here that the law-like generalizations can be *linguistically formulated* in such a way that they can figure in some ideal predictive argument.

Both the Nomological Sufficiency Analysis and the Predictive Argument Analysis entail that a sentence that expresses information irrelevant to predicting *P* can (nevertheless) be used to predict *P*. For instance, both allow that the sentence 'All whales are mammals and Carl is a car whose ignition key is being turned' can be used to predict that Carl's engine will start. I am not sure whether this is a defect of the analyses; in any case, I shall not try to reformulate the analyses so as to exclude such sentences. This defect in the analyses (if it is one) seems irrelevant to the Prediction Objections that critics lodge against Russellianism; the critics, after all, maintain that attitude ascriptions provide *too little* predictive information to be usable for predicting behavior. (For related discussion, see note 34. Thanks to Gabriel Uzquiano for discussion of this topic.)

There may be yet other senses in which the truth of (15) can be used to predict that Carl's engine will start. It might be possible to use some notion of conditional probability. For instance, one might propose that the subjective probability of 'Carl's engine will start' given (15) is high (or above .5). But subjective probabilities vary from person to person, and it's clear that the critics have in mind something more "objective" when they say "can be used to predict". One might propose some more objective notion of conditional probability, perhaps having to do with frequency, or objective chance, or logical probability. But as far as I can tell, the critics of Russellianism have nothing like this in mind; and since these notions are so difficult and contentious (even more contentious than the notions I introduced above), I shall steer clear of them.

23. As I've mentioned before, Devitt and Richard tend to present their objections concerning explanation and prediction "in the same breath". Devitt (1996, pp. 55-6) claims that attitude ascriptions attribute semantic properties to token mental events (thinkings). He says that one of our primary purposes in attributing attitudes is "to explain and predict the behavior of the subject,

which I shall abbreviate 'to explain behavior'" (1996, pp. 57-8). He again mentions prediction, together with explanation, on p. 151: "Opaque ascriptions serve our purposes of explaining and predicting behavior" (see also p. 153). His main argument against Russellianism occurs on pp. 174-5; he mentions prediction there more often than explanation, though it's hard to attribute any particular significance to this.

As far as I know, Richard explicitly mentions prediction in only one place: Richard 1997a, p. 202. When discussing Russellian claims about the pragmatics of attitude ascriptions, he says "I don't find this line of defense satisfactory, for it forces us to say that attempts to explain or predict behavior by ascribing propositional attitudes, if taken literally, *cannot* be successful" (Richard's emphasis).

In constructing the following Prediction Objections to Russellianism, I have also made use of Devitt's and Richard's objections to Russellianism concerning explanation. For these objections, see: Devitt 1996, pp. 151-3, 182-4, 243, and 304; Richard 1990, pp. 126 and 173-6; and Richard 1997a, p. 202.

24. I shall assume that if generalization (21b) is true under Russellianism, then it is sufficiently law-like to serve as a general premise in ideal predictive arguments. The arguments I present below to show that (24) is true under Russellianism would also suffice to show that (21b) is nomologically necessary under Russellianism.

25. In this section, I borrow and modify material from Braun (2000). In that paper, I present further arguments for the context-sensitivity of *ceteris paribus* phrases, and present a tentative semantic analysis that modifies the theories of Morreau (1997) and Silverberg (1996). Much of the literature on *ceteris paribus* generalizations (including Fodor 1991, Schiffer 1991, Cartwright 1983, Laymon 1989, and Pietroski and Rey 1995) is concerned with whether these generalizations can be laws. This is not my main concern here. Although critics of Russellianism tend to think that *ceteris paribus* psychological generalizations are laws, the Prediction Objections would hardly be affected if the psychological generalizations were not laws. A critic of Russellianism could conceivably hold that they are *not* laws, but that their *truth* is necessary in order for attitude ascriptions to be used to predict behavior; he could then claim that they would be false under Russellianism, and make some rather trivial changes to the objections I present in the main text. I (nonetheless) briefly discuss the debate between Fodor and Schiffer over *ceteris paribus* laws in my (2000). I think that much of the literature on *ceteris paribus* generalizations, and their status as laws, suffers from the fact that it does not sufficiently recognize the context-sensitivity of these generalizations (the above papers by Morreau and Silverberg are exceptions see also Morreau 1999).

26. There is another compelling reason to think that 'other things being equal' must be context-sensitive: its content seems to differ from one conditional to another. The phrase 'other things are equal' seems to "allude" to different suitable conditions in (26) than it does in (i).

(i) If this match is struck, then, other things being equal, it will light.

This should be sufficient to show that the phrase, or connective, is context-sensitive.

27. Thanks to Earl Conee for helpful advice about formulating the Modified Analyses and the Modified Objections (not all of which I have heeded).

28. In this section, I again borrow and modify material from Braun (2000).

29. I present further reasons to think that suitable conditions can include the matching-ways condition in Braun (2000).

30. The following account of predictive information is inspired by Railton's (1981) and Lewis's (1986) accounts of *explanatory* information. I endorsed modified versions of their views on explanation in Braun (forthcoming). My account of predictive information is a more distant descendant of their views on explanatory information.

31. Some of the causes of Carl's starting (for example, Carl's having a battery and Carl's having spark plugs) might be thought to be *states* or *state of affairs* rather than events. I am counting these as events.

32. I am ignoring indeterminism here. Taking it into account would introduce unenlightening complications.

33. To be a bit more precise: the "other things" that "must hold" are included in the class, and the "other things" that "must *not* hold" are nomologically precluded by the members of the class. For instance, in every nomologically possible world in which Carl's engine starts, one cause of the engine's starting is an event (or state) of there being air in Carl's air intake at the relevant time. This event nomologically necessitates that Carl's air intake is *not* full of water at that time. In that sense, it precludes that sort of interfering factor. Other interfering factors are similarly precluded by other causes. (Thanks to David Hunter for discussion of this.)

34. Some propositions that are ideal for predicting Carl's starting may differ from the propositions describing a temporal cross-section of events that cause Carl's engine's starting. For instance, some such propositions may describe such a cross-section *plus* some other causally irrelevant events, or may contain some other causally irrelevant information. Others may describe some events at time t_1 and some other events at time t_2 which together make up a nomologically sufficient set for the occurrence of Carl's engine's starting.

35. There may even be a sense in which the truth of (37) makes the starting "more likely" than does the truth of (15). There are nomologically possible worlds in which (15) is true but Carl's engine does not start. In some of these worlds, the engine does not start because Carl's battery is dead. (37) is not true in these worlds. Thus the truth of (37) "eliminates" some possible interferences to Carl's starting that (15) does not (that is, the set of nomologically possible worlds in which (37) is true is a subset of those in which (15) is true). In some sense, the more possible interferences to an outcome that a sentence eliminates, the "more likely" the truth of the sentence makes the outcome.

36. The following replies to the Predictive Proposition Objection are somewhat similar to my replies to certain explanation objections in Braun (forthcoming).

37. That is, sentence (1), and the proposition it expresses according to Russellianism, namely *R*, describe a certain sort of event involving Lucy, the relations of believing and describing, and certain Russellian propositions. One of the causes of Lucy's wave is an event of this sort. (I think that among its causally relevant properties are that of being-an-event-composed-of-a-believing-and-desiring, and also its being-an-event-composed-of-a-believing-and-desiring in certain Russellian propositions. But I won't argue for this.) This is so even if *R* fails to describe certain other causally relevant features of that event, such as its being (partly) a believing-in-proposition-*P-in-way-W*. (Thanks to David Hunter for discussion of this.)

38. Some Russellians might claim that utterances of (1) *pragmatically convey* proposition *AR*, even though it *semantically expresses* only proposition *R*. Thus they might agree with the critics in holding that (1) does not semantically express a proposition that can be used to predict Lucy's wave, but maintain that (1) does pragmatically convey such a proposition. I rather doubt that utterances of (1) routinely pragmatically convey *AR*, or the metalinguistic proposition mentioned in the text, or any of the other sorts of proposition that Salmon, Soames, and McKay say they do. I also think that pragmatic explanations often fail to explain ordinary speakers' judgments about the truth of (1). See Braun 1998.

This section provides much of the material needed to answer yet another Prediction Objection, one which is a variant on the classic Substitution Objection. Consider (i) (=1)) and (ii).

- (i) Lucy wants Twain to autograph her book, and she believes that if she waves, then *Twain* will autograph her book.
- (ii) Lucy wants Twain to autograph her book, and she believes that if she waves, then *Clemens* will autograph her book.

The objection: if Russellianism is true, then (i) and (ii) express the same proposition. If they express the same proposition, then either both or neither can be used to predict that Lucy will wave. (i) can be so used, but (ii) cannot. Therefore, Russellianism is false. The reply: the third premise is false. (i) and (ii) can both be used to predict Lucy's wave. It merely *appears* to ordinary speakers that (ii) cannot be so used. The reasons why it appears so are similar to those given in the text. I cannot explain any further here (for lack of space), but see my replies to similar variants on the Substitution Objection in Braun 2000 and forthcoming.

39. My other intended contributions to this project (thus far) are Braun 1998, 2000, and forthcoming.

40. Thanks to John G. Bennett, Earl Conee, David Hunter, and Gabriel Uzquiano for many valuable comments and suggestions.