

REASON WITHOUT REASONS FOR

MICHAEL G. TITELBAUM

Draft(v4.1). Please do not circulate or cite without author's permission.¹

1

1.1. Shyam Nair's (2016) begins with an example of a familiar normative picture:

Suppose for example that there is a movie theater and a restaurant across town. And suppose that in order to get to that side of town I must cross a bridge that has a \$25 toll. The toll is a reason not to cross the bridge. The movie is a reason to cross the bridge and the restaurant is also a reason to cross the bridge. It may be that if there were just the movie to see, it wouldn't be worth it to pay the toll and if there were just the restaurant, it wouldn't be worth it to pay the toll. But given that there is both the movie and the restaurant, it is worth it to pay the toll.

In this case, the reason to cross the bridge provided by the movie and the reason to cross the bridge provided by the restaurant are individually worse than the reason to not cross the bridge provided by the toll. But together these reasons are better than the reason to not cross the bridge so that you ought to cross the bridge. More generally, cases like this show that we can have two reasons to do an act that are individually worse than a reason to not do that act, but the reasons together are better than the reason to not do that act. (56–7)

This example illustrates what I'll call the **valence account** of reasons. According to the valence account, a set of considerations makes it the case that an agent ought to perform some action or adopt some attitude by virtue of the reasons *for* or *against* that action/attitude contained within the set.

Valenced analyses of a normative situation usually begin with *qualitative* notions: we seek reasons for and reasons against a particular conclusion (also known as *pro tanto* reasons). Sometimes these qualitative items suffice to determine where all-things-considered reason lies. For instance, many valence theorists would agree that if a set contains at least one fact that is a reason for a conclusion, but no facts that are reasons against that conclusion, then the set provides all-things-considered reason to draw the conclusion. But often the set of relevant facts contains both reasons for and reasons against. To handle that type of case, the valence account may introduce *comparative* notions: some reasons tell more strongly for or against a particular conclusion than others. These comparisons allow such principles as:

¹Thanks to Elizabeth Bell, Jonathan Weisberg, Eric Sampson, Sophie Horowitz, an audience at the 2016 Iowa Philosophical Society conference, and the folks I button-holed between sessions at the 2016 Chapel Hill Metaethics Workshop.

when a set contains only one reason against C , and at least one stronger reason for C , then the set provides all-things-considered reason to conclude C . But even this may not be enough. What about cases (such as Nair’s example) in which the set contains multiple reasons for C , only one reason against, but the reason against is stronger than any of the reasons for? At this point the valence account may introduce a *quantitative* notion of the weight of reasons, and some mechanism for adding those weights up. It may turn out that the reasons for, though individually weak, add up to outweigh the single reason against C .

In addition to these qualitative, comparative, and quantitative notions of reasons,² the valence account acknowledges that not every fact bearing on a conclusion is a reason for or against that conclusion. Valence theorists typically recognize some form of what Errol Lord and Barry Maguire (2016) call “conditions” and “modifiers”. Conditions are facts that, while not necessarily reasons for or against the conclusion, affect whether *other* facts count as reasons for or against. Modifiers may also not be reasons for or against; they affect the *weight* of other reasons for or against. For example, the condition *that I like to drive fast* makes the fact *that this car goes fast* a reason for me to buy it. The modifier *that I’m moving to a country with no speed limits* makes the fact *that this car goes fast* an even stronger reason to buy it. There may also be conditions and modifiers on conditions or modifiers. *That I am purchasing a car for my wife and will never get to drive it* means that *I like to drive fast* no longer makes *this car goes fast* a reason to buy the car.

While the terminology varies among valence theorists, I hope the picture I’m painting is generally familiar. And what is it meant to be a picture of? In the first instance, the valence account is an account of the normative landscape. R. Jay Wallace provides a nice summary:

A typical view of this kind is that of Joseph Raz, who has written that “[t]he normativity of all that is normative consists in the way it is, or provides, or is otherwise related to reasons.” Reasons, in turn, are commonly taken to be considerations that recommend or speak for (and against) prospective attitudes or responses on the part of a person or agent. (2012, p. 18)³

Lord and Maguire put the point even more starkly:

It is . . . increasingly common to think that the weighted [normative] notions will explain the non-weighted, and that within the weighted notions, reasons will explain the rest. (2016, p. 4)

The valence account holds that in order for a set of facts to provide all-things-considered reason for a conclusion, at least one of their number must be a reason for that conclusion. When a fact bears on a conclusion, it does so either by being a reason for or against that conclusion, or by being a condition or modifier on a reason for or against, or by being a condition/modifier on a condition/modifier, or . . .⁴

But the valence account is not just a theory of *reasons*; it also has implications for *reasoning*. Again, Wallace neatly presents the view: “We resolve the question of what we are to think or to do by . . . reflecting on the considerations that count in

²I’ve adapted the tripartite distinction from Carnap’s (1950) famous division of scientific concepts into “classificatory”, “comparative”, and “quantitative”.

³The Raz reference is to (Raz 1999, p. 67).

⁴The position could easily be given a recursive definition.

favor of holding certain beliefs or acting in certain ways.” (2012, p. 19) According to the valence account, good reasoning assesses a conclusion by taking a set of facts identified as bearing on that conclusion, determining which of them are reasons for or against the conclusion (perhaps with the help of various conditions among the set), measuring the weight of those reasons (perhaps in light of various modifiers in the set), then weighing up the reasons to ascertain whether the set as a whole provides all-things-considered reason for or against.⁵

While the valence account is familiar, my hope here is to deepen our understanding of it by questioning some of its commitments. My main question concerns its *generality*—how universally is the valence account meant to apply? Is it always the case that when a set of facts together supports some conclusion, it does so by virtue of the presence of reasons for the conclusion within that set? For the most part I will tackle this generality question by exploring the valence account’s consequences for reasoning. I will consider a number of cases that I take to involve paradigmatic examples of correct reasoning. In each example, there will be a set of facts that together support drawing some conclusion, but none of the facts considered singly looks like a reason for that conclusion (even in the presence of the others). Of course looks may be deceiving, and there will be strategies the valence account can employ to try to bring those cases into the fold. But as the cases proceed, these strategies will come to seem increasingly desperate. Ultimately I will question whether the valence account offers a good general picture of the activity of reasoning, and of the normative landscape that lies beneath.

1.2. Before we get to business, a few technicalities. First, I will assume (along with much of the literature) that reasons are facts. This assumption is merely a convenience; nothing important in my argument hangs on it.

Second, in discussing a piece of reasoning, I will talk as if its conclusion is a proposition *C*. I would like my discussion to apply to both theoretical and practical reasoning, and I realize it’s controversial what philosophical category conclusions of reasoning occupy. For episodes of theoretical reasoning, the conclusion is usually a proposition, or perhaps a propositional attitude. If it’s the latter, we can think of the proposition *C* as the content of the concluded attitude.⁶ In cases of practical reasoning, the conclusion may be an intention, in which case *C* is again an attitude content. Or the conclusion of practical reasoning may be an action, in which case we may take *C* to characterize that action. (*C* might be the proposition *that you dance like nobody’s watching*, or whatever.) Again, I don’t think anything important in

⁵Not all reasoning assesses a conclusion in light of facts identified as bearing on that conclusion. Sometimes reasoning is about generating new options, new conclusions to be assessed. And other times we reason about which facts bear on the conclusion of interest. But I will focus on reasoning that assesses a pre-identified conclusion in light of a pre-collected set of facts.

⁶The valence account may acknowledge some special cases of propositional attitudes as results of theoretical reasoning. For example, suppose I consider all the facts in my possession that bear on proposition *C*, and conclude by suspending judgment in *C*. The valence account may admit that none of the facts that led me to suspend judgment is itself a *reason for suspending judgment in C*. Still, the valence account will probably offer some analysis of why suspension is appropriate in terms of reasons for and against the proposition *C* (or the attitude of believing *C*). To avoid delving into such analyses here, I will set aside theoretical reasoning that results in suspension of judgment, and whatever the practical analogue of that may be (if there is one). I will focus on cases in which a set of facts bearing on proposition *C* together provides some all-things-considered reason to believe *C*.

my argument hangs on precisely what form the conclusions of reasoning take. So I will describe the conclusion of a piece of reasoning as a proposition C , just to give myself a way to talk about it.

Third, I should comment on the phrase “bears on”. There are two ways of reading my main question, but ultimately they strike me as notational variants of each other, with the difference lying in how one uses the word “reason”. One reading treats all the facts an agent properly takes into account in concluding C as “reasons” relevant to C . On that reading, I am questioning whether such a set of reasons always includes reasons *for or against* C . On the other reading, “reasons” must by definition be either for or against. (So, for instance, being a condition or modifier on a reason for/against C doesn’t suffice to be a reason relevant to C .) This seems to be Kieran Setiya’s reading in the following passage:

Reasons for acting—by which I mean *good* or *normative* or *justifying* reasons—are considerations that count in favor of doing something, as for instance the fact that the weather is beautiful is a reason for me to quit working and go outside, while the fact of an impending deadline is a reason to get the job done. As these examples indicate, good reasons are always *facts*; they are features of the world that should go in the pro column if we were to list the pros and cons of someone’s pursuing a course of action. (2003, p. 346, emphases in original)

On this reading, my question is whether, when an agent properly reasons from a set of facts to a conclusion C , one of those facts must be a reason relevant to C . I don’t feel a need to choose between the two readings here.⁷ So I will describe facts that an agent properly takes into account in reasoning to conclusion C as facts that “bear on” C , thereby remaining neutral on whether such facts always count as reasons.⁸

Fourth, I should clarify what I take myself to be doing when I challenge the generality of an account of reasoning. We’re all familiar with the image of making a decision by “weighing up the pros and cons,” and the valence theory is usually offered as a philosophical account of such weighing. I started questioning that account when I noticed that it almost never seems to me like I’m drawing a conclusion by tallying reasons for and against it. But psychologists repeatedly warn against reading the structure of reasoning off of what it seems like you’re doing when you’re reasoning.⁹ (For one thing, we sometimes reason so quickly and automatically that there’s nothing that the process “seems like” at all!) Moreover, we may often reason using heuristics that are faster or less cognitively taxing than carefully weighing up the pros and cons.

So I won’t assess the valence account as a psychological theory of how reasoning always actually goes, or seems to go. The best way to read the valence account as fully general is to interpret it as a theory of how reasoning always *could* go: In every

⁷For what it’s worth, it seems to me that the second reading better tracks most natural uses of “reason” in English. But I could easily be wrong about that.

⁸If it helps, think of the facts that “bear on” C as the facts that are relevant to C . That seems a perfectly fine equivalence to me, given how we use “relevant” in everyday language. Unfortunately, “relevant” has also come to have a technical meaning, and for reasons I’ll explain later I don’t want to endorse the thesis that every fact that bears on C is relevant to C in the technical sense.

⁹See, for example, (Nisbett and Wilson 1977).

case in which an agent considers a set of facts, and (correctly) determines that they together provide all-things-considered reason to draw some conclusion, it's always *possible* that if the agent slowed down and made all her reasoning explicit, she could tell a compositional story about how that all-things-considered feature of the set arises from individual reasons for and against within the set. Reasoners don't always do this explicitly; sometimes the weighing happens subconsciously, while other times heuristics are used to get similar results in a more efficient fashion. But the ultimate story about what's going on (or what could be going on), and why it's correct, can always be given in terms of reasons for and against. *That's* the version of the valence account I intend to challenge.

Finally, I should clarify that while holists about normativity will reject the entire valence project on grounds of its compositionality, I am not a holist. I believe there is a bottom-up, atomistic, compositional process by which facts combine to support conclusions.¹⁰ I want to know whether it's always the case when a conclusion is supported that at least one of the inputs to that compositional process is a reason for or against.

2

2.1. My first example is inspired by one from Richard C. Jeffrey's *Logic of Decision* (1983). You've been invited to a dinner party, and tasked with bringing the wine. You must choose among a red, a white, and a rosé. Before you arrive, the hosts will choose between two entrées, steak or a delicate fish. The pairing of your wine with their entrée yields the following possible utility payoffs:

utilities	steak	delicate fish
red	5	-10
white	-10	5
rosé	-1	-1

You are equally confident that the hosts will serve steak or fish. Which wine should you bring?

When I reason about this case, I conclude that you should bring the rosé, on decision-theoretic grounds. Nowhere in applying the decision theory do I search for reasons for or against bringing rosé among the stated facts of the case. Reasons for are no part of the machinery of decision theory.¹¹

As I noted above, the valence theorist need not accept my reasoning at face value. First, she may respond that whenever a conclusion is drawn, multiple ways of reasoning to it are available; just because one way of reasoning to rosé contains no reasons for or against doesn't mean there aren't other ways of reasoning that do. Or, the valence theorist may contend that just because I don't explicitly think

¹⁰I also tend to think that the compositional process is systematic, and might be expressible in general principles. But I will not assume that here, so my approach should be compatible with various forms of particularism.

¹¹Dietrich and List (2013) note that "reasons are largely absent from modern rational choice theory." (p. 105) They then propose a reason-based theory of rational choice, but their theory largely uses reasons for and against to determine an agent's utility distribution over possible worlds. When it comes time to combine utilities and probabilities, Dietrich and List do so in the traditional fashion, which does not count either the utility facts or the credence facts as reasons for performing a particular act. (Thanks to Reuben Stern for bringing the Dietrich and List to my attention.)

about reasons for or against while doing decision theory, that doesn't mean they aren't in there somewhere.

All right, then: Of the facts provided in the story above, which could be considered a reason for bringing rosé? Is it that rosé yields an expected utility of -1 when steak is served? But red yields an even higher utility in that case! Similarly for the utility fact about rosé and fish. What about the fact that you divide equal confidence between the two possible entrées? Well, that doesn't seem to be a reason to bring one wine rather than another. . . .

So it's not obvious which fact is meant to be a reason for rosé. But the valence account has some strategies available here. First, the valence theorist could argue that the equal confidence split is in fact a reason to bring rosé. If that were the only fact in the story it would not act as a reason for rosé, but the valence theorist might tell some story about how the other facts act as conditions and modifiers to make it so. (Maybe there's something in there about radical entrée uncertainty making it best to bring the safest wine.)

A second strategy for the valence account would be Reasons Contrastivism. (See especially (Snedegar 2015).) On a contrastive view, a fact is a reason for a particular conclusion only relative to specified other options. The contrastivist might say that the fact that rosé yields a higher utility than white when steak is served is a reason for bringing rosé *rather than* white in case of steak. Similarly, the fact that rosé yields a higher utility than red when fish is served is a reason for bringing rosé *rather than* red. Now the contrastivist still has some work to do: we need to combine these reasons that obtain only on particular entrée selections into an overall decision in light of entrée ignorance. But at least the contrastivist has identified some of the facts in our story as reasons (in some sense) to bring rosé.

Finally, the valence theorist could adopt what I call a Summary Fact strategy. The idea here is to identify some other fact in the vicinity that summarizes the relevant property of the facts presented in the story, then treat that summary fact as a reason (even if the other facts aren't). For instance, one could suggest that the fact *that bringing rosé maximizes expected utility* is your reason for bringing rosé.¹²

The Summary Fact strategy will often be tempting in the examples to come. It suffers from one universal problem, and one problem particular to this case. The universal problem is that we will always need to know how the agent is supposed to reason to the summary fact. In the present example, you're meant to take the facts given in the case, reason to the conclusion that rosé maximizes expected utility, and then treat that fact as your reason to bring the rosé. But this just pushes the problem back to the reasoning that gets you the conclusion that rosé maximizes expected utility. If the valence account offers a compositional account of *all* reasoning, then the utility calculations that get you this conclusion should also operate with reasons for and against. And again, I don't see which of the facts outlined in the original story is meant to be a reason for the conclusion that rosé maximizes expected utility.

Again, a version of this problem will come up every time we consider the Summary Fact strategy below. However, the Summary Fact approach has a particular problem in this example, having to do with how we understand utility in the case. Sometimes utility is understood as a natural entity or property, which our actions may produce in greater or lesser quantities. (Canonical examples include pleasure,

¹²Thanks to Eric Sampson, among others, for this suggestion.

happiness, etc.) If we understand utility this way, then even once an agent has determined what maximizes utility (or expected utility), there can always be a further question whether utility in something she has all-things-considered reason to maximize. I take it that on the proposed interpretation of the wine case, the agent calculates that rosé maximizes expected utility, then some other facts in the vicinity make it the case that she has reason to choose the act which maximizes expected utility. These further facts are then conditions which convert the fact that bringing rosé would maximize expected utility into a reason for bringing rosé.

Yet that's not how I meant utility to be understood in setting up the case. I meant utility to be a direct measure of value—genuinely *normative* value—not a measure of some sort of natural product that might or might not be valuable.¹³ You're not meant in the example to first determine what would maximize expected utility, then ask the further question what you have all-things-considered reason to do. The expected utility calculation is meant to be your means of calculating what you have all-things-considered reason to do. Put another way, calculating expected utilities is how you determine whether the facts in the story give you all-things-considered reason to bring rosé. On this reading of the utilities in the story, it's much more difficult to get the proposed Summary Fact response going.

2.2. While I doubt the Summary Fact strategy can be made to work for the decision-theoretic case, the other two valenced strategies might have some hope. So let's move on to a second example, in which neither of those strategies fares as well.

In the days leading up to the 2016 presidential election, the RealClearPolitics website offered daily averages of the most recent polls. At 9:30am on the morning of November 7, they reported an average lead for Hillary Clinton of 2.5%. They also reported the individual poll results that factored into their average; for instance, a CBS News/New York Times poll showed Clinton leading by 4%. While some of the individual polls were closer to 2.5% than that, and some were farther, none of the polls reported exactly a 2.5% Clinton lead.

The folks at RealClearPolitics began with a set of facts about the results of individual polls, and reasoned their way to a conclusion about what they should report. So which of those facts was a reason for the conclusion that they should report an average of exactly 2.5%? The CBS/NYT result clearly bears on that conclusion, but is a 4% poll a reason to report a 2.5% average? Was any of the individual poll results a reason to report that average?

Here our valence strategies from above look hopeless: We can't pick one of the poll results and argue that that one (in the context of all the others) is the reason to report an average of 2.5%. Which one would we pick? We could try going contrastivist again. Perhaps a 4% poll result is a reason to report an average of 2.5% rather than, say, an average of 50%. But if that's right,¹⁴ then a 4% result is also a reason to report an average of 3.9% rather than 2.5%. And now all the interesting, important compositional work is going to be done by whatever story combines these reasons relative to *pairs* of possible averages into a selection amongst the full

¹³I take it this is how many contemporary ethical utilitarians interpret the utilities appearing in their theories.

¹⁴And I'm not entirely sure that it is. Depending on what other polls are in the collection, a 4% result might be a stronger reason to report an average of 50% than 2.5%.

set of averages available.¹⁵ Perhaps we have found some (contrastive) reasons for reporting 2.5% in the mix, but they look awfully trivial and inconsequential.

Is a summary fact available? Perhaps RealClearPolitics' reason for *reporting* an average Clinton lead of 2.5% is that, given all the individual poll results, the mathematically-calculated mean *is* 2.5%. This mathematical fact supplies a reason for reporting 2.5% (in light of various other conditions making it a bad idea to try to mislead the American people about basic mathematics). But now we have the same universal problem as the first time we considered a Summary Fact strategy: Shift focus to the intermediate conclusion. Before the RCP folks reasoned to the intention or action of *reporting* an average of 2.5%, they had to reason from the individual polls to a *belief* that the average at that time was 2.5%. That is, they had to do a bit of theoretical reasoning in order to do their practical reasoning. The conclusion of that theoretical reasoning was the intermediate conclusion *that the current polling average is* 2.5%, and the individual poll results were the relevant facts bearing on that conclusion. So what fact was the reason for that conclusion? Here it won't do to invoke the summary fact *that the current polling average is* 2.5% as a reason for that conclusion—we cannot reason from a fact to itself!

Yet a new strategy is available to the valence account in this case.¹⁶ Suppose that prior to receiving any of the polling data, the RCP folks had some opinions about what the average polling gap between Clinton and Trump might be. Perhaps they even went so far as to assign degrees of confidence across all the available numerical values. When the 4% CBS News result comes in, it might increase their confidence that the ultimate polling average will be 2.5%. That doesn't mean that their confidence in a 2.5% average given a single 4% result need be very high. But it will have increased, while other possible averages (such as a 2.5% Clinton *deficit*) will have decreased. Focusing on the deviation the CBS News result causes from the pollsters' default confidence in an average of 2.5%, the 4% result counts as a reason in favor of reporting 2.5%.

This Default-Comparison strategy is also appealing because it meshes with theories of reasons available in the literature. To this point I have relied on intuitive judgments about whether a particular fact counts as a reason for a particular conclusion. For theorists who take reasons to be basic, there may not be too much more to say. But some philosophers have offered detailed theories about when a fact counts as a reason for a conclusion. Many of these theories (such as those of Schroeder (2007), Finlay (2006), and Kearns and Star (2009)) involve confidence- or probability-raising of some sort. They differ on precisely what proposition needs its probability raised, and raised relative to what. But the general theme that a reason raises the probability of some proposition relative to a default value is there. So it's a strategy worth considering as we go forward.

2.3. Let's test that strategy on another case. Suppose that initially I know nothing about Tiresias. Then I learn the following facts:

All men are mortal, and all women are immortal.
Only women serve Hera.
Tiresias serves Hera.

¹⁵In Snedegar's terminology, the reasons relative to normative questions Q containing only two alternatives have to be combined into an ought relative to the question Q containing *all* the averages RealClearPolitics might report.

¹⁶Thanks to Sophie Horowitz for this suggestion.

From these three facts I reason to the conclusion that Tiresias is immortal. Which of the facts I learned is a reason for that conclusion? Take a minute and consider each of those facts one at a time. Does any of them look like a reason to conclude that Tiresias is immortal?

The valence theorist will grant that of course none of the facts is a reason on its own for the conclusion that Tiresias is immortal—the reason for that conclusion only becomes a reason for in light of the other two. But which one is that one? Can the valence theorist provide any non-arbitrary grounds for designating one as the reason and the others as conditions or modifiers?

In case the reasons theorist thinks she can, consider the example of these three facts:

1. $(A \& \sim B \& C) \vee (A \& \sim B \& \sim C) \vee (\sim A \& B \& C) \vee (\sim A \& B \& \sim C)$
2. $(A \& B \& \sim C) \vee (A \& \sim B \& \sim C) \vee (\sim A \& B \& C) \vee (\sim A \& \sim B \& C)$
3. $(A \& B \& C) \vee (A \& \sim B \& \sim C) \vee (\sim A \& B \& \sim C) \vee (\sim A \& \sim B \& C)$

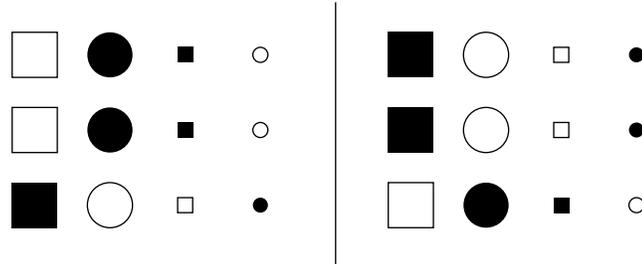
An agent wondering whether proposition A is true ought, in light of these three facts, conclude that it is. Now which of these facts is the reason for A , made such by the presence of the other two? The lesson is the same for both the Tiresias example and this more abstract one: *When multiple premises jointly entail a conclusion, none of them needs be a reason for that conclusion.*

I think it's with these logical cases that the valenced accommodation strategies we've tried so far begin to look truly desperate. There's no non-arbitrary way to argue for one of the premises as a reason for and the others as conditions or modifiers. Contrastivism is also useless here: In the Tiresias deduction, the clear contrast is between the conclusion that Tiresias is immortal and the conclusion that Tiresias is mortal; in the more abstract deduction it's between A and $\sim A$. In neither case does one of the premises look like a reason for the former conclusion over the latter.

What about the Summary Fact strategy? The valence theorist could take the fact *that the conclusion is logically entailed by the premises* as a summary fact providing reason for the ultimate conclusion. But then she'll run into versions of both the problems the Summary Fact strategy had in the decision-theoretic case. The universal problem, once again, will be that we need to know how the compositional reasoning from the premises to the summary fact goes, and which of the premises is a reason for that fact. In deductive cases all the really interesting reasoning happens at the stage where the agent determines that the conclusion is entailed; the valence account needs something to say about that reasoning.

But then there's also the problem that this Summary Fact approach misreads the significance of the logical entailment in this case. Logical entailment is the relation between the set of premises and the conclusion *in virtue of which* that set provides all-things-considered reason to believe the conclusion. (Many other such relations are possible; we'll see one in our next example.) An agent doesn't first determine that a set of premises entails a conclusion, then add the fact that she has reason to believe what's entailed by her premises, then as a result decide she has all-things-considered reason to believe the conclusion. Why not? Because that bit of reasoning itself involves a logical entailment. ("If this set of premises entails this conclusion then I have reason to believe the conclusion," "This set of premises

FIGURE 1



entails this conclusion,” so. . . .) If every logical entailment needed a further fact to provide a reason, we would be off on an Achilles/Tortoise regress. (Carroll 1895)

Which brings us back to the Default-Comparison strategy. We can’t apply that strategy to these examples yet, because we don’t know what the default confidences/probabilities are supposed to be going into our deductions. So let’s make some reasonable ones up. In the more abstract deduction case, each of the conjunctive disjuncts appearing in the premises is what’s known as a “state-description”. (In a language with only the three atomic sentences A , B , and C , these state-descriptions give the fullest description available of the state of a possible world.) With three atomic sentences there are eight possible state-descriptions. So let’s suppose that before any of the premises is given, each of the eight state-descriptions (each possible arrangement of the world in terms of those atomic sentences) is equally probable. In that case, the default probability of A is $1/2$. With a bit of calculation, you should be able to determine that *none* of the premises in the argument raises the probability of A . That is, if you start with no information, and take any one of those premises by itself, the probability of A in light of that premise is still $1/2$. (In fact, that remains true of any *pair* of premises within the set; it’s only once all three premises are added that the probability of A increases above $1/2$.) So even if we think of a reason for a conclusion as a fact that raises the probability of that conclusion above some default value, none of the premises in the abstract deductive argument is a reason for that argument’s conclusion.

2.4. This last effect may be easier to grasp in an example of inductive theoretical reasoning.¹⁷ Consider the twenty-four shapes in Figure 1. Each is either large or small, black or white, a circle or a square. Suppose I tell you that I have selected one of them at random, with each shape having an equal chance of being selected. I then provide you with the following facts:

My shape is large.
My shape is a square.
My shape is white.

I ask you to consider whether my shape is to the left or the right of the dividing line. (Feel free to think about it for a moment, and do some counting.)

¹⁷Inductive examples may also be preferable to deductive because logical entailments are sometimes viewed as special limit cases of more commonplace reasoning practices. To the extent I understand the legal aphorism, developing a theory of reasoning by attending primarily to deductive cases may be a situation in which “hard cases make bad law.”

I submit that taken together, those three facts give you some all-things-considered reason to conclude that my shape is to the left of the line. The reason is not conclusive, but it is there. Now if the valence account applies to this case, at least one of the three facts must be a reason for that conclusion, either considered in isolation or in light of the other two. So which one(s) is it?¹⁸

Let's work through our valenced strategies in turn. First, there doesn't seem to be any way to single out one of the facts as the genuine reason for the conclusion in the context of the other two. The three facts are perfectly symmetrically situated relative to the conclusion that my shape is on the lefthand side. Second, a contrastive approach will not help, as we are already considering a clear contrast of conclusions between left and right. Third, the relevant summary fact in the vicinity (that the set raises the probability of the conclusion) can be dismissed for reasons we've already covered.

Fourth and finally, what about comparison to defaults? Given the background information about how I selected my shape, each shape should clearly start with a probability of $1/24$. Since there are twelve shapes to the left of the line, the left conclusion has a default confidence of $1/2$. Now take any of the three facts above and add it to your evidence. Your confidence in left should remain $1/2$. (For example, twelve of the shapes are squares, and six squares lie left of the line.) None of the three facts increases the probability of left above its initial value. So on a probability-raising theory of reasons, none of those facts is a reason for the conclusion. In fact, to make matters worse, no conjunction of two of the facts raises the probability of left above $1/2$; put another way, assuming one of the facts doesn't make another of them a probability-raiser. (For example, the number of *large* squares is the same on each side of the line.) The probability that my shape is left of the line goes up to $2/3$ only once *all three* facts are incorporated.

This illustrates something interesting about a cluster of Bayesian notions one might have thought akin to reasons for. On the Bayesian picture, evidence E confirms hypothesis H just in case the probability of H given E is greater than the default probability of H . This condition is also known as *positive relevance*, or *positive correlation*; Bayesians often take it to be necessary and sufficient for E 's constituting *evidence for H*. If that's right, then in the present example a set of facts provides evidence for a conclusion even though no fact within the set is evidence for that conclusion. Put another way, an entire set may be relevant to a conclusion even if no fact in that set is relevant.¹⁹ This is no strike against Bayesian confirmation theory; it simply shows that the Bayesian approach to evidence doesn't satisfy platitudes suggested by thinking in terms of reasons for.

To sum up: In the shapes example there are no grounds to claim that any of the facts provided is a reason for the conclusion that my shape lies to the left of the line. Yet together those facts provide all-things-considered reason to draw that conclusion. The valence account looks to have failed again.

¹⁸I find that whenever I present a statistical example, some readers have a tendency to fight against the example (perhaps suggesting that there isn't really any all-things-considered support in this case). Rest assured that if I wanted, I could tweak the numbers here so that the ratio between large white squares on the left and right sides is as high as you like. I could also arbitrarily increase the number of facts provided as evidence, while keeping intact all the relevant symmetries. I simply presented a case with small numbers to simplify visual counting.

¹⁹This is why I opted earlier for "bears on" rather than "relevance" talk in my discussion—see note 8 above.

3

3.1. The valence theorist has one last, Hail Mary strategy available that could be applied across the board. In each of the cases I’ve offered, the valence account could maintain that the *conjunction* of the facts bearing on the conclusion counts as a reason for that conclusion, with the individual facts counting as reasons for that conjunction. This strategy has the advantage that at least in the cases of theoretical reasoning, it meshes well with a probability-raising account of reasons for. A standard result in probability theory demonstrates that as long as a conjunction and each of its conjuncts has non-extreme probability, each of the conjuncts individually raises the probability of the conjunction.²⁰ And at least in the deductive and inductive reasoning cases I offered above, the conjunction of the pieces of evidence described does in fact raise the probability of the conclusion.²¹

Does this strategy save the valence account’s claim to full generality? I think the answer is subtle, and depends on what exactly the account is trying to achieve.

First, it’s important to see that when a fact in a set is a reason for a conjunction, and that conjunction is a reason for some conclusion, this does *not* entail that the fact is a reason for that conclusion. On a probability-raising account of reasons for, the reasons-for relation need not be transitive. For example, the fact that I have brown hair raises the probability of the conjunction that I have brown hair and the number of stars in the universe is even. That conjunction raises the probability that the number of stars in the universe is even. But that doesn’t mean the fact that I have brown hair is a reason for the conclusion that the number of stars in the universe is even.²²

So does the conjunctive strategy establish the valence account’s full generality? On the one hand, we’ve found a way that whenever an agent has all-things-considered reason to draw some conclusion, the valence theorist can offer a kind of explanation of this fact in terms of reasons for. When an agent has all-things-considered reason to conclude that C , she possesses a set of facts that each provides reason for some conjunction, and that conjunction provides reason for C . On the other hand, the initial aspiration of the valence account was to explain why a particular set of facts provides all-things-considered reason to conclude C in terms of facts within the set that are reasons for C . The failure of reasons-for transitivity shows that we can’t achieve this aspiration just by showing that facts within the set are reasons for some conjunction that is reason for C .

And there’s another, related problem here. Showing that individual conjuncts raise the probability of their conjunction explains why an agent in possession of those conjuncts has reasons to believe the conjunction is *true*. But we need something more explained: We need an explanation why that conjunction is a reason for

²⁰The “non-extreme” qualifier is required to make that standard result work. It means that a particular class of cases—cases in which the facts bearing on a conclusion and the conclusion itself are all mathematical theorems—cannot be covered by the strategy under consideration. But cases in which reasons for a conclusion and the conclusion itself all have probability 1 were never going to be amenable to a probability-raising account of reasons for, due to the general “logical omniscience” problem in probability theory. So this isn’t a *particular* problem for the conjunctive strategy I’m considering at this point in the text.

²¹In the abstract deductive case the conjunction is equivalent to $A \& \sim B \& \sim C$, which raises the probability of A from $1/2$ to 1. In the inductive case the conjunction says that my shape is large, square, and white, which raises the probability of left from $1/2$ to $2/3$.

²²This “problem of irrelevant conjunction” is well-known to formal epistemologists.

the conclusion C . In the theoretical reasoning cases above, this further explanation may be provided by the probability-raising account of reasons for. But in other cases, we have no reasons for-based explanation of why the conjunction of the facts on offer is a reason for the conclusion of interest. Why is the conjunction of the values in the utility table plus the fact that you're equally confident of steak and fish a reason to bring rosé? (Of course I can offer a decision-theoretic explanation of why that conjunction provides a reason to bring rosé, but this explanation doesn't employ reasons for.)

The valence theorist may object that her account doesn't owe us an explanation of why particular facts are reasons for particular conclusions. Many valence theorists consider the reason-for relation basic; it's supposed to ground all other normative relations. But this reply commits the valence theorist to treating like cases in unlike fashion. Return to Nair's toll bridge example. The valence theorist may say that the conjunction of facts about the toll price, the movie, and the restaurant provides a reason for crossing the bridge. But the whole point of Nair's discussion is to explain *why* this conjunction is a reason for—it's because the movie and restaurant provide independent reasons for crossing that together outweigh the reason against provided by the toll. When the valence account works, it explains why a conjunction is reason for a conclusion by pointing to reasons for that conclusion among the conjuncts.

In the dinner party case we know full well why the conjunction of utility and confidence facts provides a reason to bring the rosé. The explanation runs through decision theory, a genuinely normative theory that provides substantive standards for reasoning from sets of facts to actions (or intentions). It's just that this explanation doesn't work by identifying reasons for or against. It makes no sense to maintain that the reason-for status of the conjunction in Nair's example can be explained in terms of normative properties of the conjuncts involved, while the reason-for status of the conjunction in the decision theory example admits of no genuinely normative explanation. That is, it makes no sense unless one is antecedently committed to a theory on which all normative explanations must bottom out in reasons for.

And really, the same point can be made about the other examples. In the deductive cases, philosophers have a well-developed theory of the relations that must hold among a set of facts in order for their conjunction to entail some conclusion. (We teach it in our introductory logic courses.) Similarly, we have a formal theory of what must be going on with the facts in a set for the conjunction of those facts to increase the probability of a conclusion. But neither of these theories trades in the currency of reasons for.

If the valence theorist simply wants to defend the claim that whenever an agent has all-things-considered reason to draw some conclusion, there exists somewhere in the vicinity a reason for that conclusion, then the conjunctive strategy can take her that far. But if the valence account is supposed to do more—if it's supposed to be a general compositional account of correct reasoning, or a story about the ultimate grounds of normativity—then the conjunctive strategy fails to do the job. Whether we focus on what a set of facts supports, or on what the conjunction of the elements in that set supports, there are cases in which the relevant normative explanation just doesn't work in terms of reasons for and against.

3.2. We're philosophers. There's always another move available. Perhaps the valence account could go epistemicist: in each case above one of the facts is the true reason for the conclusion; it's just that there's no way of knowing which one it is. But it's a bit difficult to make that work. In at least the last two cases (abstract deduction and shapes), the facts are all situated perfectly symmetrically with respect to the conclusion, so it's difficult from a metaphysical point of view to figure out how there could even be a difference that would make it the case that one is a reason while the others aren't. So maybe instead we go supervenientist. . . .

At this point we're trying to save a theory that doesn't want saving. In all of the examples above, everyone knows in virtue of what the facts together support the conclusion. In the first example, we know how to do decision theory; in the second example, we know how to do math; in the entailment examples, we know how to do logic; in the shapes example, we know how to do statistical inference. There is no mystery in any of these cases how to reason from the facts that bear on the conclusion to support for that conclusion. The trouble comes in only when we try to take these well-known reasoning patterns and layer over them a narrative of reasons for, a narrative that isn't native to those ways of reasoning.

So what happens if valence theorists bite the bullet—what if they admit that the valence account is *not* a fully general account of how agents come to have all-things-considered reason? Well, then we'll have to ask if any fully general compositional account is available, which I think is a very open question. I hope to have suggested that if there *is* a fully general theory in the offing, then the valence account—with its tallying up of pros and cons—isn't it. Perhaps in some restricted set of cases, the right way to reason is by considering reasons for and against. Or perhaps valence is sometimes a useful heuristic that approximates what truly makes reasoning correct. (In which case we'll have to treat valence like any other heuristic, and watch out for cases in which it leads us astray.)

There may also be a further consequence here. It's no coincidence that normative explanations in terms of reasons for have recently gained popularity in tandem with realist metaethics. Reasons have a metaphysical profile that fit well with realism; we can imagine there are truths in the universe about which facts are reasons for which conclusions. Yet when we look at the first and last examples above, the norms in those cases are most naturally explained by decision theory and Bayesian confirmation theory. To the extent we understand the normative underpinnings of those theories, they seem best understood in terms of subjective rational consistency, rather than realist-style norms.²³ Now neither of those theories is perfect, and each certainly has its counterexamples. But suppose our account of why certain sets of facts obligate an individual to perform a particular act or adopt a particular attitude moves more in the direction of those sorts of theories, and away from a valenced approach. In the end, this may be bad news for metanormative realism.

²³See, for instance, Niko Kolodny's work on this point. I also take this to be one major upshot of the trend among Bayesians towards Subjective rather than Objective versions of Bayesianism.

REFERENCES

- Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Carroll, L. (1895). What the Tortoise said to Achilles. *Mind* 4, 278–280.
- Dietrich, F. and C. List (2013). A reason-based theory of rational choice. *Noûs* 47, 104–34.
- Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy* 84, 1–20.
- Jeffrey, R. C. (1983). *The Logic of Decision* (2nd ed.). Chicago: University of Chicago Press.
- Kearns, S. and D. Star (2009). Reasons as evidence. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Volume 4, pp. 215–42. Oxford: Oxford University Press.
- Lord, E. and B. Maguire (2016). An opinionated guide to the weight of reasons. In E. Lord and B. Maguire (Eds.), *Weighing Reasons*, pp. 3–26. Oxford: Oxford University Press.
- Nair, S. (2016). How do reasons accrue? In E. Lord and B. Maguire (Eds.), *Weighing Reasons*, pp. 56–73. Oxford: Oxford University Press.
- Nisbett, R. E. and T. D. Wilson (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review* 84, 231–259.
- Raz, J. (1999). Explaining normativity: On rationality and the justification of reason. In *Engaging Reason*, pp. 67–89. Oxford: Oxford University Press.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Setiya, K. (2003). Explaining action. *Philosophical Review* 112, 339–393.
- Snedegar, J. (2015). Contrastivism about reasons and ought. *Philosophy Compass* 10, 379–88.
- Wallace, R. J. (2012). Constructivism about normativity: Some pitfalls. In J. Lenman and Y. Shemmer (Eds.), *Constructivism in Practical Philosophy*, pp. 18–39. Oxford: Oxford University Press.