

REASON WITHOUT REASONS FOR

MICHAEL G. TITELBAUM

Draft (v5.2).¹

In 1772, Joseph Priestley was invited by Lord Shelburne to become his librarian. Conflicted about whether to leave behind his happy life in Leeds, Priestley wrote to Benjamin Franklin for advice. Franklin replied as follows:²

In the Affair of so much Importance to you, wherein you ask my Advice, I cannot for want of sufficient Premises, advise you *what* to determine, but if you please I will tell you *how*. When these difficult Cases occur, they are difficult chiefly because while we have them under Consideration all the Reasons *pro* and *con* are not present to the Mind at the same time; but sometimes one Set present themselves, and at other times another, the first being out of Sight. Hence the various Purposes or Inclinations that alternately prevail, and the Uncertainty that perplexes us. To get over this, my Way is, to divide half a Sheet of Paper by a Line into two Columns, writing over the one *Pro*, and over the other *Con*. Then during three or four Days Consideration I put down under the different Heads short Hints of the different Motives that at different Times occur to me for or against the Measure. When I have thus got them all together in one View, I endeavour to estimate their respective Weights; and where I find two, one on each side, that seem equal, I strike them both out: If I find a Reason *pro* equal to some two Reasons *con*, I strike out the three. If I judge some two Reasons *con* equal to some three Reasons *pro*, I strike out the five; and thus proceeding I find at length where the Ballance lies; and if after a Day or two of farther Consideration nothing new that is of Importance occurs on either side, I come to a Determination accordingly. And tho' the Weight of Reasons cannot be taken with the Precision of Algebraic Quantities, yet when each is thus considered separately and comparatively, and the whole lies before me, I think I can judge better, and am less likely to take a rash Step; and in fact I have found great Advantage from this kind of Equation, in what may be called *Moral* or *Prudential Algebra*. (Franklin 1975, emphases in original)

¹Thanks to Michael Barkasi, Jeff Behrends, Elizabeth Bell, Josh DiPaolo, Luca Ferrero, Branden Fitelson, Sophie Horowitz, Ram Neta, Eric Sampson, Justin Snedegar, Roy Sorensen, Reuben Stern, Sergio Tenenbaum, Jonathan Weisberg, audiences at the 2016 Iowa Philosophical Society conference, the 2017 St. Louis Annual Conference on Reasons and Rationality, and the 2017 Bergen conference on the Normativity of Logic, and the folks I button-holed between sessions at the 2016 Chapel Hill Metaethics Workshop.

²Thanks to Roy Sorensen for bringing Franklin's letter to my attention.

The assessment of normative situations by tallying up reasons pro and con has attracted increasing philosophical interest over the past few decades. Philosophers have considered such questions as: What makes a particular consideration a reason for or a reason against a particular decision? What determines the weights of such reasons? How do the weights of reasons interact? and so forth.

I want to ask a slightly different question, one that I do not believe has been adequately addressed (if it has been addressed at all). I want to ask how generally Franklin’s “Algebra” can be applied. Franklin is confident that whatever the details of Priestley’s affairs, his dilemma can be resolved by weighing the pros and cons. Can every normative issue that is resolvable be resolved in this fashion? Whenever we set out to make up our minds, will it always be the case that if a determination can properly be reached, it can be reached by tallying up reasons for and against? Or are there cases in which, while the relevant considerations definitely combine to point in one direction, they do not do so by offering reasons for that conclusion?

In the first section of this essay I will specify more precisely how I understand the claim that every normative case is amenable to something like Franklin’s approach. In the second section, I will offer a series of counterexamples to this claim. In the final section, I will consider the philosophical significance of the fact that one can have all-things-considered *reason* to draw a conclusion without having any individual *reasons* for doing so.

1

1.1. I will label a normative situation that *can* be correctly analyzed by something like Franklin’s method a “Franklin case”. The purpose of this section is to specify more precisely what I mean by a Franklin case. To help us keep a fixed example before our minds, here’s a passage from Shyam Nair:

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. (Nair 2016, pp. 56–7)

This is a canonical example of what I will call a Franklin case. The agent in the example has all-things-considered reason to cross the bridge. This is because the agent’s reasons for crossing (having to do with the movie and the restaurant)

together outweigh the agent's reason against (having to do with the toll). The pros outweigh the cons, so the agent ought to cross.

In Nair's example, a number of considerations bear on the agent's decision: that it costs \$25 to cross the bridge, that the movie theater is across town, that the restaurant is across town. Going forward, I will assume that such normatively relevant considerations are always facts. Here I follow the trend in the metaethical literature of assuming that reasons are facts. Yet I do so as a convenience; I don't think anything important in my argument hangs on this assumption.

I will label an example a Franklin case when it satisfies three conditions:

- (1) If the agent in the case ought to do something (perform an action, form a belief, etc.), then some set of facts together provides that agent with all-things-considered reason to do so.

To use a bit of contemporary jargon, we might say that Franklin cases instantiate a mild, highly general form of "rationalism". Why mild? Because rationalism is usually a thesis about reasons in particular; moral rationalism, for instance, is the thesis that moral oughts always entail the existence of reasons. Condition 1 merely says that when an agent ought to do something, there is some set of facts that together normatively support the agents' doing so. It says nothing about the nature of the facts in the set, or what about them allows them to provide such support.

What kinds of thing might be required by an ought? I am interested in cases of both practical and theoretical support, and it would be convenient to be able to discuss both kinds of case with a unified vocabulary. So (perhaps somewhat artificially) I will talk about a set of facts' together requiring an agent to "conclude C ", where C is some proposition. In theoretical cases, to conclude C will simply be to believe C . In practical reason cases, to conclude C might be to form an intention with the content C . Or it might be to act in the manner described by C . (For example, in Nair's case C is the proposition *that I cross the bridge*.) I realize it's controversial what sorts of entities are supported by reasons; hopefully all of the available approaches can be shoehorned into this way of talking. And again, I don't think anything important in my argument depends on the ontology of the conclusions of reasoning.

The next necessary condition for a Franklin case is:

- (2) If in the case a particular set of facts provides at least some all-things-considered reason to conclude C , then it does so because of local relations among the facts within the set and C . These local relations make it the case that the set provides all-things-considered reason to conclude C , and the local relations explain why the set provides all-things-considered reason.

It's important to see that this condition goes beyond what's required by Condition 1. Many philosophers would endorse the kind of mild rationalism that takes all normative cases to satisfy Condition 1; they would grant that oughts are always accompanied by supporting sets of facts. But Condition 2 requires that support to be structured in a very particular way: it has to flow *bottom-up* from local relations among those facts.³

³Compare (Dancy 2004, p. 16): "To talk of what there is overall reason to do (and note that 'reason' in this phrase is not a count noun) is to talk about where the contributory reasons come down—on this side or that."

Contrast this structure with Jonathan Way’s account of reasons. Way maintains that:

For the fact that p to be a reason for S to ϕ is for there to be a good pattern of reasoning from the belief that p , perhaps together with other correct attitudes which S has, to ϕ -ing. (Way 2017, p. 254)

Notice that on this view, for p to be a reason to ϕ *is for* it to factor in a good pattern of reasoning in a particular way. More generally, the facts in a set count as reasons for a conclusion C just in case—and *because*—those facts can act as premises in a piece of good reasoning to C . Way is explicit that his view “takes the correctness of attitudes and the goodness of reasoning to be prior to reasons.” (ibid.)⁴ Thus he has a top-down approach: the facts in the set are reasons for C because the set has some global property that can be independently understood. In a Franklin case, the determination works the other way; the set provides all-things-considered support for C *because* it contains reasons for C . This is clearly the way we’re meant to read Nair’s case; the movie and the restaurant provide reasons to cross the bridge that make it worth paying the toll.⁵

The final condition for being a Franklin case is:

- (3) In the case, the local relations among the facts and C that make the set provide all-things-considered reason to conclude C all belong to one of the following types:
- Qualitative, comparative, or quantitative reason for or against relations
 - Being a condition on a fact in the set
 - Being a modifier on a fact in the set

These relation-types require some explaining. First, the reason for and against relations: Here I’ve borrowed Carnap’s famous (1950) tripartite division of concepts. The qualitative relations in question are being a reason for C and being a reason against. (These are sometimes called *pro tanto* reason relations.) For example, that the restaurant is on the other side of the bridge is a reason for crossing the bridge. Comparative relations include one fact’s being a stronger reason for or against C than another. For example, Nair suggests that the amount of the toll is a stronger reason against crossing than the restaurant’s location is a reason for. Finally, quantitative relations say that a particular fact supports C to a particular degree; the fact has a specific weight as a reason for or against C . These weights may be combined, as when Nair suggests that the movie and restaurant locations together outweigh the reason against crossing provided by the toll.

⁴Kieran Setiya also affirms that “A reason to ϕ counts in favour of doing so by contributing a premise for sound reasoning to the desire or motivation to ϕ .” (2014, p. 228). But Setiya remains explicitly non-committal about whether the reason counts as a reason *because* it is a premise in a piece of sound reasoning, or the reasoning is sound *because* its premises count as reasons. (See (Setiya 2007, p. 13).)

⁵On either a top-down or a bottom-up view, the facts in a set together provide reason to draw the conclusion; they make it the case that the agent ought to conclude C . That’s what Condition 1 is about. Condition 2, and the top-down/bottom-up distinction, is about what makes it the case that the set has that global property—why it is that those facts make it the case that the agent ought to conclude C . In a Franklin case, the set provides reason to conclude C *because* the facts in that set stand in particular local relations to each other and to C .

What does it take for a fact to be a reason for or against some conclusion, for one fact to be a stronger reason than another, or for a fact to provide a reason of a particular weight? I am not going to provide any such analysis here (though I will mention one possibility later). I will assume only that these relations can be functionally characterized,⁶ and that they obey particular platitudes. For example:

- Whenever a set of facts provides all-things-considered reason to conclude C , at least one of the facts in the set is a reason for C .
- If a set of facts contains at least one reason for C and no reasons against C , then the set provides at least some all-things-considered reason to conclude C .⁷
- Combining two reasons for C produces a weight in favor of C no less than that of either reason taken alone.⁸

The “condition” and “modifier” terminologies come from (Bader 2016); other authors employ similar notions under different names.⁹ A condition is a fact that, while perhaps not itself a reason for or against C , affects whether some fact counts as a reason for or against. For instance, the fact that the agent enjoys movies is a condition making the location of the movie a reason to cross the bridge. A modifier affects the weight of a reason for or against C . For instance, if the agent has ten billion dollars in the bank, this may weaken the \$25 toll as a reason against crossing the bridge.¹⁰

I will continue to describe reasons for and against C , conditions on those reasons, and modifiers on them, as being among the facts that “bear on” C . It might seem more natural to describe all these kinds of facts as “reasons” relevant to C . But the word “reason” is used in two different ways in the literature. Some authors define “reason” so that a reason must always be for or against, as when Setiya writes, “‘Reasons,’ here, are considerations that count in favour.” (2007, p. 9) Other authors use “reason” more expansively, so that conditions and modifiers also count as reasons.¹¹ To avoid having to settle this issue, I will describe all the facts in a set that normatively supports concluding C as facts that “bear on” C .¹² Condition 3 says that in a Franklin case, all of the facts bearing on C are reasons for or against, conditions, or modifiers. Together, Conditions 1 through 3 entail

⁶Not necessarily *defined* or *reduced*—just characterized!

⁷Compare (Berker 2007, p. 125): “According to one common conception, a reason for action is a consideration that would decisively count in favor of a given action were no other reasons present.”

⁸Again, Berker (2007, p. 130): “The individual contribution made by a reason of positive valence always positively affects the total reason in favor of the action in question.”

⁹For more on conditions and modifiers, see (Lord and Maguire 2016). (Dancy 2004) draws similar distinctions using different terminology.

¹⁰Franklin’s description of his method doesn’t seem to allow for conditions and modifiers, so I guess the “Franklin case” moniker is a bit of an anachronism. I include cases with conditions and modifiers in the class under consideration both to bring our discussion in line with the contemporary reasons literature, and to make it at least somewhat plausible that the class might be fully general.

¹¹One example is (Scanlon 1998, p. 51): “Reasons can be related to one another in more complex ways. I may, for example, judge one consideration, C , to be a reason for taking another consideration, D , not to be relevant to my decision whether or not to pursue a certain line of action.”

¹²We could also describe these facts as “relevant to” C . However, I will try to avoid “relevance” talk because it has a technical usage, which I will discuss later.

that in any Franklin case in which an agent ought to conclude C , there exists at least one fact that is a reason for C (perhaps in the context of other facts acting as conditions or modifiers). To deny that all cases involving normative support are Franklin cases, I will provide examples in which an agent ought to conclude C even though no fact is a reason for C .¹³

1.2. Before I do so, however, a few more clarifications. First, we need to slightly complicate the functional structure described by Conditions 1 through 3. Suppose you're confronted with the following three facts:

Only left-handers wear gloves on their right hands.
Johnny wears a glove on his right hand.
The team's best pitcher is left-handed.

It seems to me that this collection of facts together supports (at least to some extent) concluding C , that Johnny is the team's best pitcher. But that's not because any of these three facts is a reason for C . Instead, the first fact is a condition that makes the second fact a reason for the proposition that Johnny is left-handed. The third fact then acts as a condition making this proposition a reason for the conclusion that Johnny is the team's best pitcher.

I would want to class this example as a Franklin case. Admittedly, we have a set of facts that together provides all-things-considered reason to conclude C , even though none of those facts is a reason for C . Yet the set comes to have this property in virtue of local relations among the facts it contains, and all of those relations are of the types listed in Condition 3. It's just that some of the facts, instead of being reasons for C , are reasons for an *intermediate conclusion* (that Johnny is left-handed), which—while not actually a member of the set in question—nevertheless acts as the crucial reason for C . A subset of the original set of facts supports the intermediate conclusion in exactly the manner described by Conditions 1 through 3, then adding that intermediate conclusion to the original set provides support for C in the manner typical of a Franklin case.

Since this type of case is obviously in the spirit of a Franklin case, I propose amending our official Franklin case definition to bring it into the fold. The amended definition would probably be easiest to write up with a series of recursive clauses. Hopefully the idea is clear enough that I don't have to go through the details here.

Next, I should clarify how I think normative support relates to reasoning. I believe that if a set of facts all-things-considered supports drawing a conclusion, and does so by virtue of local relations among those facts, then it ought to be possible for an agent to reason from those facts to that conclusion by recognizing those relations (and perhaps writing them down to keep track, as Franklin suggests). This need not be the *only* way to reason from those facts to that conclusion, and it need not even be the *best* way in some broad sense. We're all familiar with the Franklinian image of making a decision by "weighing up the pros and cons", but it's unclear how often we actually do that in our heads. (In fact, I started thinking about the topic of this essay when I noticed that it almost never seems to me like I'm making a decision that way.) Psychologists study mental heuristics that reach conclusions while skipping over some of the normative details, but confer

¹³Kearns and Star write, "The following is a very plausible principle to many people: *If one ought to F, then there is a reason for one to F.* It seems impossible for it to be the case that one ought to perform an action if there is no reason at all to do so." (2008, p. 51–2, emphasis in original) I intend to demonstrate the possibility of what seems to them impossible.

advantages in speed or reduced cognitive load. They also have repeatedly warned against reading the structure of one’s reasoning off of what seems to be happening when you’re doing it.¹⁴ Among other things, we sometimes reason so quickly and automatically that there’s nothing the process “seems like” at all.

So I am not interested in classifying cases according to how the agent’s reasoning in the case actually goes, or seems to go (even if we confine our attention to cases in which the agent is reasoning well). Instead, I am interested in how the reasoning in a case *could* go:¹⁵ In every Franklin 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 should be *possible* for that agent to slow down, work explicitly through a series of correct reasoning steps, and tell a bottom-up story about how that all-things-considered feature of the set arises from individual reasons for and against the conclusion within the set. As Jay Wallace puts it: “We resolve the question of what we are to think or to do by... reflecting on the considerations that count in favor of holding certain beliefs or acting in certain ways.” (2012, p. 19)

Finally, while certain kinds of holists about normativity will immediately reject the suggestion that all (or even any) cases of normative support are Franklin cases, I am not a holist. I believe there is a bottom-up, atomistic, compositional process by which facts combine to support drawing conclusions.¹⁶ That is, I believe that every case of normative support satisfies Conditions 1 and 2 of the definition of a Franklin case. My issue is with Condition 3. I think that the ideology of reasons, conditions, and modifiers misdescribes the functional structure of normative support in many cases. That’s what the counterexamples are meant to illustrate.

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. While

¹⁴See, for example, (Nisbett and Wilson 1977).

¹⁵Compare (Silverstein 2016, p. 2, n. 5).

¹⁶I also tend to think that this compositional process is systematic, and might be expressible in general principles. So I tend to reject particularism. But I will not address that issue in this essay.

decision theory offers a machinery for composing sets of facts into normative conclusions, reasons for are no part of that machinery.¹⁷

As I noted above, one need not accept my reasoning, nor my reports about it, at face value. First, 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 second, just because I don't explicitly think 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 could be a reason for rosé, and this example doesn't seem to be a Franklin case. But there are some maneuvers available that might bring this example back into the fold. First, one could simply insist that there must be a reason for rosé in there somewhere. Call this the Foot-Stomping rescue strategy. Perhaps 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 one could try to tell a 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 rescue strategy 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, there's the Summary Fact rescue 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 for the conclusion (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.¹⁸

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

¹⁷Dietrich 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.)

¹⁸Thanks to Eric Sampson, among others, for this suggestion.

reason to the summary fact. A summary fact is meant to operate as an intermediate conclusion, of the type I modified the Franklin case definition to accommodate. But it was crucial to that modification that the intermediate conclusion be supported by some subset of the original facts in the case, along the lines specified by the Franklin definition. In the present example, you should be able 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 pushes the problem back to the reasoning that gets you the conclusion that rosé maximizes expected utility. If this truly is a Franklin case, then the utility calculations that get you the intermediate conclusion should operate with reasons for and against. And again, I don't see which of the facts outlined in the original story is 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 also 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, 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 is 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.¹⁹ 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 example, the two other strategies might be able to recast that example as a Franklin case. 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.

¹⁹I take it this is how many contemporary ethical utilitarians interpret the utilities appearing in their theories.

The folks at RealClearPolitics began with a set of facts about the results of individual polls, and reasoned their way to a conclusion of what to 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 rescue 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,²⁰ 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 set of averages available.²¹ 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 reasons against misleading the American people with respect to basic mathematics). But now we have the same universal problem as the first time we considered the 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 rescue strategy is available in this case.²² 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%.

²⁰And 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%.

²¹In 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.

²²Thanks to Sophie Horowitz for this suggestion.

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.

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? This certainly doesn't seem to be a Franklin case.

The Foot-Stomper will grant that of course none of the facts above 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 Foot-Stomper provide any non-arbitrary grounds for designating one as the reason and the others as conditions or modifiers?

In case he thinks he can, consider 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 rescue 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? In each case we 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 we'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 when the agent determines that the conclusion is entailed; the Summary Fact response needs something to say about that reasoning.

But then there's also the problem that the Summary Fact response misreads the significance of 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 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)

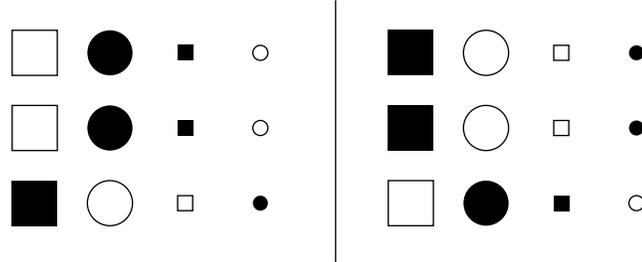
Which brings us back to the Default-Comparison Strategy. We can't apply that strategy to these examples yet, because we don't know what the default confidences/probabilities are supposed to be going into our deductions. So let's make some reasonable ones up. In the more abstract deduction case, each 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.²³ 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.

²³Inductive 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."

FIGURE 1



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.)

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 this example is a Franklin case, at least one of the three facts must be a reason for the left-of-line conclusion, either considered in isolation or in light of the other two. So which one(s) is it?²⁴

Let's work through our rescue 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.

To sum up: In this 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

²⁴I 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.

together those facts provide all-things-considered reason to draw that conclusion. So the example, while clearly an instance of normative support, is not a Franklin case.

2.5. I'd like to pause for a moment to relate the shapes example to some points I made earlier, and to some notions common in discussions of inductive reasoning.

Earlier I listed some platitudes that come along with functional picture of reasons for and against. The first platitude was that whenever a set of facts provides all-things-considered reason to conclude C , at least one of the facts in the set is a reason for C . I've just argued that the shapes example violates that platitude.²⁵

Since the shapes case involves inductive reasoning, and in particular probabilistic inductive reasoning, we can recast that platitude in probabilistic terms. Bayesian analyses of induction typically start with the assertion that evidence E inductively 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 also take it to be necessary and sufficient for E 's constituting *evidence for* H . If that's right, then in the shapes 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.²⁶ This is no strike against Bayesian confirmation theory; it simply shows that the Bayesian approach to evidence doesn't satisfy a platitude central to the structure of reasons for.

What about the other reasons-for platitudes? Well, consider this example: I randomly select a positive integer between 1 and 10, inclusive, such that each integer has an equal chance of being selected. You are interested in the conclusion C that my integer is prime. Consider these facts:

My integer is odd.
My integer is less than 3.

Each of these facts, considered in isolation, raises the probability of C , and so on a probability-raising account supplies a reason for C . Yet combined they provide conclusive reason for $\sim C$. (Recall that the number 1 is not prime.) Thus this example violates the platitude already mentioned. It also violates the other two platitudes I mentioned earlier:

- If a set of facts contains at least one reason for C and no reasons against C , then the set provides at least some all-things-considered reason to conclude C .
- Combining two reasons for C produces a weight in favor of C no less than that of either reason taken alone.

Examples like this one violate everything an account of normative support based on reasons for would lead us to expect.²⁷ This is what I mean when I say that,

²⁵Some reasons theorists are suspicious of comparative or especially quantitative reason relations. So suppose we pared down the definition of a Franklin case to include only qualitative relations. The new definition would still invoke the platitude I've just discussed, and so still would not apply to the shapes example (and many of the others I've presented).

²⁶This is why I opted earlier for "bears on" rather than "relevance" talk—see note 12.

²⁷Sorensen (2016, pp. 8–9) offers an example in which "reasons that separately raise hope can jointly dash that hope." His moral is "that what is individually confirming can be collectively disconfirming." Meanwhile, Salmon (1975) finds in Carnap's (1950) work on probability examples with all of the following structures:

while I'm comfortable with the idea that normative support has a bottom-up, compositional structure, I don't think that structure always has the functional shape suggested by an account based on reasons for.

2.6. There is one last, Hail Mary strategy that has the potential to cast all of these examples as Franklin cases. In each of the examples I've offered, one 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 Conjunctive 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.²⁸ So if we treat the conjunction of the original facts as an intermediate conclusion, we don't have the problem (suffered by the Summary Fact Strategy) of explaining how the original facts count as reasons for the intermediate conclusion. Moreover, it's clear at least in the deductive and inductive reasoning cases I offered above that this intermediate conclusion (the conjunction of the pieces of evidence described) does in fact raise the probability of the ultimate conclusion C .²⁹

So does this strategy allow us to read every example of normative support as a Franklin case? The idea would be that in every such example, the individual facts bearing on C are reasons for their conjunction, and that conjunction is then a reason for C . This seems to satisfy the recursive definition of a Franklin case that accommodates intermediate conclusions, and so make every example of normative support a Franklin case.

Yet I don't think the Conjunctive Strategy really gives us what we want from a Franklin case. At best, it recovers the mild rationalism of Condition 1 and an approach to reasons like Jonathan Way's. For our purposes, it's not enough to

-
- Two facts each raise the probability of a conclusion, but their conjunction lowers that probability
 - Two facts each raise the probability of a conclusion, but their disjunction lowers that probability
 - A fact raises the probability of each of two conclusions, but lowers the probability of their conjunction
 - A fact raises the probability of each of two conclusions, but lowers the probability of their disjunction

In personal communication, Branden Fitelson notes that if two facts each raise the probability of a conclusion, then either their conjunction or their disjunction must raise its probability. And if a fact raises the probability of each of two conclusions, then it must also raise the probability of either their conjunction or their disjunction. While these are fascinating theorems, they underscore for me the ways in which the functional structure of probability-raising is very different from that of additive reason weights.

²⁸The "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.

²⁹In 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 $2/3$.

establish that wherever there's normative support, it's possible to find something reasonish in the vicinity. Given Conditions 2 and 3, whenever a set of facts all-things-considered supports drawing a conclusion, we want an *explanation* in terms of reasons for and against of why it does so. It's not enough to say just that the facts are reasons for their conjunction, and then the conjunction is a reason for the conclusion. Saying that the conjunction of the facts in a set is reason for a conclusion is another way of saying that the set supports that conclusion. You don't explain why the set has a particular property by stating that it has that property in a slightly different way. Notice that on the Conjunctive Strategy there's no account of why certain sets of facts wind up supporting and others don't—no subtle interplay of reasons, conditions, modifiers, or other local relations to generate a global result.

In each of the examples we've considered, a localized explanation for why the set supports drawing the conclusion (or why the conjunction of elements in the set provides a reason for that conclusion) is readily available. In Nair's example, that explanation operates in the Franklin style: the set contains reasons for and against the conclusion, the reasons for together outweigh the reasons against, so the set as a whole provides all-things-considered reason to conclude *C*. In the other examples, an explanation is available, but that explanation doesn't work in terms of reasons for and against. In the wine example, the explanation is decision-theoretic; in the polling example, it's mathematical; in the deduction examples, it's logical; and in the inductive examples, it's probabilistic. I have argued that none of those explanations has the right functional structure to operate in terms of reasons, conditions, and modifiers.

The amendment of the Franklin case definition to allow for intermediate conclusions was meant to open another route for local reasons relations among the facts in a set to explain why that set supports drawing an ultimate conclusion. The Conjunctive Strategy abuses that amendment to create an artificial reason that doesn't explain anything.³⁰ We know what the relevant explanations are in all of our examples; some of those explanations have the Franklin structure and some of them do not. Rather than obscure the true explanations with an artificial conjunctive substitute, we should admit that some examples are Franklin cases and some are not.

3

We're philosophers. There's always another move available. Perhaps we could go epistemicist: in each example 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 two of the examples (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 generalization that doesn't want saving. As I said a moment ago, in each of my examples we know what the correct explanation is of the relevant normative support; that explanation just doesn't have the structure

³⁰For a different set of arguments against rolling up all of the facts bearing on a conclusion into a conjunction that's reason for the conclusion, see (Bader 2016, §5).

Franklin envisioned. So let's accept this, and move on to what follows. What are the consequences of the fact that not every example of all-things-considered reason involves a reason for?

Well, one question will be whether an alternative, fully general compositional account of normative support is available. I think this question remains wide open. In the meantime, I don't want to deny that at least some examples of normative support are genuine Franklin cases; in certain situations it does seem appropriate to make a decision by tallying up reasons pro and con. And in other situations, the reasons for approach may provide a useful heuristic even if it doesn't track the true normative structure of the case.

But once we recognize that Franklin's approach isn't fully general, we should stop blithely assuming that every normative situation can be analyzed in terms of reasons for and against. I can't tell you how many times I've read a philosophical paper that otherwise doesn't even mention reasons, but at a crucial juncture the author defends a normative claim by suddenly invoking the weight of reasons on each side.³¹ This is particularly egregious when the topic is probabilistic relations; we've seen that probabilistic confirmation doesn't pattern with reasons for and against. But it also can be problematic when one discusses philosophical methodology.³²

Often, a philosopher will defend a thesis about a particular area with an argument structured as follows: There are a variety of popular positions in this area; the thesis is plausible on this position, the thesis is plausible on this other position, etc.; so in general my thesis is well-supported. To the extent I can make sense of this argumentative strategy, the idea seems to be that the plausibility of the thesis on each position is a reason in favor of that thesis; those reasons stack up; so the thesis is well-supported in general. Yet this assumes exactly the sort of additive normative structure that I've been arguing doesn't universally apply. Are the enumerated positions mutually exclusive? Could the plausibility of the thesis on one position undermine its plausibility on another? Is the plausibility in question meant to be probabilistically construed?³³ David Enoch (2011, pp. 14ff.) likes to assess philosophical theories by tallying up "plausibility points"; the metaphor is apt only if the relevant form of plausibility can be added and subtracted like scores in a game.³⁴

But the lesson of this essay shouldn't only affect how we defend philosophical theories; it should also effect which theories we defend. Over the past few decades, more and more philosophers have come to subscribe to Joseph Raz's position that "[t]he normativity of all that is normative consists in the way it is, or provides, or is otherwise related to reasons." (Raz 1999, p. 67) Along the way, a particular approach to metanormativity has gained in popularity. Lord and Maguire write,

It is . . . increasingly common to think that the weighted [normative] notions will explain the non-weighted, and that within the weighted

³¹Once you're attuned to the phenomenon, I bet you'll start to notice it too!

³²Thanks here to Luca Ferrero, who after hearing a talk on this material talk asked me, "That's all fine, but now how am I supposed to do philosophy?"

³³For what it's worth, the only formal theory of plausibility I'm aware of (a component of Demster-Shafer theory) is intimately related to probability theory and bears nothing like the functional structure of reasons for.

³⁴Enoch takes his inspiration from Lewis and Lewis (1970), who prefer to speak of plausibility "prices".

notions, reasons will explain the rest. (Lord and Maguire 2016, p. 4)

In a similar vein, Selim Berker discusses

a widely held framework according to which reasons for action are the fundamental normative units whose interactions determine, through a metaphorical balancing of the weight of reason, all other normative properties of actions. (Berker 2007, p. 112)

These are claims about the fundamental structure of normativity, and they have consequences. Suppose that interactions among reasons for and against explain (or determine) all other normative properties. Then in any case in which a normative property is present, reasons for or against must be as well, and the former must obtain because of the latter. This implies that all examples of normative support must be Franklin cases. So my argument in this essay rules out a certain kind of reasons fundamentalism.³⁵

Finally, I don't think it's a coincidence that reasons fundamentalism and the tacit assumption that all cases are Franklin cases have gained popularity in tandem with metaethical realism. On a rival metaethics like, say, constructivism, normative properties are grounded in features of the agent. Those features give rise to a limited collection of principles or standards, which in turn generate all else that is normative (including reasons). Reasons fundamentalism, on which a plethora of facts out in the universe stand in reasons relations without those relations' deriving from anything else, is much more at home in a realist metaethics. On the other hand, explanations of normative support that do not boil down to reasons may undermine realism. In my first and last examples above, the relevant norms 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.³⁶ 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 reasons-based approach. In the end, this may be bad news for metanormative realism.

REFERENCES

- Bader, R. (2016). Conditions, modifiers, and holism. In E. Lord and B. Maguire (Eds.), *Weighing Reasons*, pp. 27–55. Oxford: Oxford University Press.
- Berker, S. (2007). Particular reasons. *Ethics* 118, 109–39.
- Carnap, R. (1950). *Logical Foundations of Probability*. Chicago: University of Chicago Press.

³⁵In a series of articles (such as his (2016)), Daniel Fogal has attacked the reasons-first program and defended a top-down account of reasons. Many of his arguments start from the way we use “reasons” talk in ordinary language. While I agree with many of Fogal’s conclusions, I hope it’s clear how those arguments differ from mine. I don’t much care about the use of the word “reasons”; I’m more interested in the assertion that in every example of normative support, that support arises from facts (call them what you will) that are related to the conclusion via functional relations with a particular structure.

³⁶See, 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.

- Carroll, L. (1895). What the Tortoise said to Achilles. *Mind* 4, 278–280.
- Dancy, J. (2004). *Ethics without Principles*. Oxford: Oxford University Press.
- Dietrich, F. and C. List (2013). A reason-based theory of rational choice. *Noûs* 47, 104–34.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Finlay, S. (2006). The reasons that matter. *Australasian Journal of Philosophy* 84, 1–20.
- Fogal, D. (2016). Reasons, reason, and context. In E. Lord and B. Maguire (Eds.), *Weighing Reasons*, pp. 74–103. Oxford: Oxford University Press.
- Franklin, B. (1772/1975). Letter to Joseph Priestley. In W. B. Willcox (Ed.), *The Papers of Benjamin Franklin*, Volume 19, pp. 299–300. New Haven and London: Yale University Press.
- Jeffrey, R. C. (1983). *The Logic of Decision* (2nd ed.). Chicago: University of Chicago Press.
- Kearns, S. and D. Star (2008). Reasons: Explanations of evidence? *Ethics* 119, 31–56.
- 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.
- Lewis, D. and S. Lewis (1970). Holes. *Australasian Journal of Philosophy* 48, 206–12.
- 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.
- Salmon, W. C. (1975). Confirmation and relevance. In G. Maxwell and R. M. Anderson Jr. (Eds.), *Minnesota Studies in the Philosophy of Science*, Volume VI, pp. 3–36. Minneapolis: University of Minnesota Press.
- Scanlon, T. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Setiya, K. (2007). *Reasons without Rationalism*. Princeton and Oxford: Princeton University Press.
- Setiya, K. (2014). What is a reason to act? *Philosophical Studies* 167, 221–35.
- Silverstein, M. (2016). Reducing reasons. *Journal of Ethics & Social Philosophy* 10, 1–22.
- Snedegar, J. (2015). Contrastivism about reasons and ought. *Philosophy Compass* 10, 379–88.
- Sorensen, R. (2016). *A Cabinet of Philosophical Curiosities: A Collection of Puzzles, Oddities, Riddles and Dilemmas*. Oxford: Oxford University Press.

- 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.
- Way, J. (2017). Reasons as premises of good reasoning. *Pacific Philosophical Quarterly* 98, 251–70.