

CONFIRMATION, CAUSATION, AND SIMPSON'S PARADOX

BRANDEN FITELSON

branden@fitelson.org

ABSTRACT

In this paper, I review some recent treatments of Simpson's Paradox, and I propose a new rationalizing explanation of its (apparent) paradoxicality.

I. SIMPSON'S PARADOX: THE USUAL DIALECTIC

There was a period in the 1970s when the admissions data for the UC–Berkeley graduate school (hereafter, BGS) exhibited some (*prima facie*) peculiar statistical correlations.¹ Specifically, a strong negative correlation was observed between being female and being accepted into BGS. This negative correlation (in the overall population of BGS applicants) was (initially) a cause for some concern regarding the possibility of gender bias in the admissions process at BGS. However, closer scrutiny of the BGS admissions data from this period revealed that *no individual department's* admissions data exhibited a negative correlation between being female and being admitted. In fact, every department reported a *positive correlation* between being female and being accepted. In other words, a correlation that appears at the level of the general population of BGS applicants is *reversed in every single department* of BGS. This sort of *correlation reversal* is known as *Simpson's Paradox*. Because admissions decisions at BGS are made (autonomously) by each individual department, the lack of departmental correlations seems to rule-out the gender bias hypothesis as the best (causal) explanation of the observed correlations in the data. As it happens, there was a strong positive correlation between being female and applying to a department with a (relatively) high rejection rate. That is, women tended to apply to departments within BGS that have low(er) acceptance rates. And, this is what explains the observed correlations.

This informal presentation is typical of (most) discussions of Simpson's Paradox that one finds in the literature. Typically, one sees *causal explanations* of *how* some set of (“Simpson's-Paradoxical”) statistical correlations could have arisen in the first place (i.e., causal “how-possibly” explanations of the “paradoxical” correlational patterns in the statistical data). In this paper, I will be doing something different. I will be discussing what I will call *rationalizing explanations* of the paradoxicality of Simpson's Paradox. That is, I will be examining some explanations of *why* the cases may *seem* paradoxical

1 I am borrowing this example from Nancy Cartwright (1979). Cartwright was one of the first to offer (probabilistic) causal analyses of Simpson's Paradox. Since then, many statisticians and philosophers have studied the paradox and various causal explanations of how it can arise. See Malinas and Bigelow (2016) for a nice recent survey.

in the first place.² In this sense, my treatment of Simpson's Paradox will be more similar to the way philosophers have approached (many) other "paradoxes" — by trying to *explain (away)* the sense of paradoxicality itself.³ In the next section, I will provide some background on confirmation theory. In the remaining sections, I will (a) discuss some existing treatments (both causal and confirmation-theoretic) of Simpson's Paradox, and (b) offer a new (purely confirmation-theoretic) rationalizing explanation of the paradoxicality of Simpson's Paradox.

Here is a simplified, precise version of the BGS example, which will serve as our canonical example of Simpson's Paradox for the remainder of the paper. Let there be two departments: E and \bar{E} (think: *easy* and *not easy*) and two genders (of applicants): M and \bar{M} (think: *male* and *not male*).⁴ And, let A (\bar{A}) express the proposition that an applicant is *accepted* (*not accepted*). The following table gives all the relevant *acceptance rates* (i.e., statistical frequencies, which we will assume align with the objective probabilities/chances, in this situation) regarding this simplified characterization of BGS admissions (involving a total of 260 applicants, 130 of which applied to department E and 130 of which applied to \bar{E}).

These acceptance rates are to be understood as follows. For instance, $\Pr(A|E\&M) = 60/80$ means that the probability of acceptance for male applicants to department E is $60/80$ — i.e., $60/80$ of the male applicants to department E were accepted (and this frequency accurately reflects the corresponding objective probability of acceptance).

The crucial thing to notice here is that \bar{M} is *negatively* relevant to acceptance in the overall population, i.e., $\Pr(A|\bar{M}) < \Pr(A|M)$, while \bar{M} is *positively* relevant to acceptance *within each of the two departments*, i.e., $\Pr(A|\bar{M}\&E) > \Pr(A|M\&E)$ and $\Pr(A|\bar{M}\&\bar{E}) > \Pr(A|M\&\bar{E})$. This kind of *reversal of correlations* (between subpopulations and the overall population) is known as *Simpson's Paradox*.

In the next section, I will provide a bit of theoretical background on confirmation theory. Then, I'll examine two recent discussions of Simpson's Paradox. Finally, I will present a new rationalizing explanation of the (apparent) paradoxicality of Simpson's paradox.

2. SOME BACKGROUND ON CONFIRMATION THEORY

From a deductive-logical point of view, "argument strength" does not come in degrees. An argument is either deductively valid or it is not. Confirmation theory aims to *generalize*

2 Such correlational structures are not really paradoxical at all — they occur in nature all the time. I will return to this point later, when I offer my *explanation-away* of the paradoxicality of Simpson's Paradox.

3 One analogous case that comes to mind here is Hempel's treatment of the Paradox of Confirmation (as explained in Fitelson and Hawthorne 2010). Hempel explains (away) the (apparent) paradoxicality of the paradox of the ravens by (a) isolating an ambiguous presupposition of the "paradox" (which has a false precisification and a true precisification), and then (b) arguing that the sense of paradoxicality arises from the conflation of these two conflicting precisifications of the ambiguous presupposition. My strategy for explaining away the paradoxicality of Simpson's Paradox will be very similar (at this level of abstraction) to Hempel's.

4 The explanations discussed below can all be generalized to cases of Simpson's Paradox that involve m departments and n genders, for any $m, n \geq 2$. But, for simplicity, I will speak only about the 2×2 case.

this purely qualitative deductive notion of “argument strength.” Ideally, it would be nice to have a *numerical measure of the degree to which* some premise P supports some conclusion C . At least, that was Carnap's (1962) dream – to construct a *confirmation function* $c(C, P)$, which (a) generalizes deductive entailment (in a sense to be made precise, below) and (b) is defined using probabilistic concepts, where (c) the probabilities in terms of which c is defined are “objective” (and, ideally, *logical*). For present purposes, we won't need anything like a Carnapian confirmation function. This is for two main reasons. First, we'll only be looking at cases in which there are (uncontroversial, well-defined and well-known) statistical probabilities (viz., *objective chances*). So, we won't need to worry about the “objectivity” of the probability assignments that will underlie our confirmation-theoretic relations (and we won't need to mess around with “logical” probabilities). Second, we'll only need *qualitative* confirmation-theoretic concepts in order to provide our analyses (i.e., numerical confirmation measures won't be required for our purposes). In this sense, we will be employing a more Bayesian (and less Carnapian) approach to confirmation. However, we will follow Carnap in making a key distinction between two kinds of probabilistic confirmation.

Confirmation as Firmness (confirms_f). P confirms $_f$ C (relative to background supposition S) iff $\Pr(C|P\&S) > t$, where t is some (possibly contextual) threshold.

Confirmation as Increase in Firmness (confirms_i). P confirms $_i$ C (relative to background supposition S) iff $\Pr(C|P\&S) > \Pr(C|S)$.

For present purposes, all of the conditional probabilities $\Pr(\cdot|\cdot)$ will be given by a specification of some (uncontroversial) statistical probabilities (viz., objective chances). So, there will be no controversy about where the probabilities in question come from. Our focus here will be on using these two confirmation-theoretic relations to characterize Simpson's Paradox and explain why it *seems* paradoxical (but, in fact, is *not*).

One important theoretical similarity between these two confirmation relations is that they both constitute *generalizations of deductive validity*, in the following sense. If P entails C , then P confirms C , relative to any background supposition S . To see why, note that if P entails C , then $\Pr(C|P\&S)$ will be *maximal* (i.e., equal to 1), and so we will (generally⁶) have *both* $\Pr(C|P\&S) > t$ and $\Pr(C|P\&S) > \Pr(C|S)$. That is, entailing premises will (generally) make their conclusions *highly probable and raise the probabilities* of their conclusions (*above* the values taken relative to whatever background suppositions are involved).

For our present purposes, the key theoretical *differences* between these two notions of confirmation can be seen by thinking of them as constituting (similar, but subtly different) explications of the concept of “argument strength.” If a premise P confers a high probability on conclusion C , then (at least in one sense) the argument can be said to be a “strong”

5 When we talk about “background suppositions,” we are always thinking of *indicative* suppositions (as opposed to *subjunctive* suppositions). See Joyce (1999) for more about the concept of “probability relative to an indicative supposition,” and how it differs from a *subjunctive*-suppositional conception of $\Pr(\cdot|\cdot)$, which may be more appropriate for causal or explanatory applications (e.g., causal analyses of Simpson's Paradox).

6 Strictly speaking, I am assuming here that $\Pr(C|S)$ and $\Pr(P\&S)$ do not take extreme values, and that t is *strictly between* $1/2$ and 1. But, in all the cases we'll be discussing, these (weak) assumptions will hold.

one.⁷ However, even if P confers a high probability on C , it may also be true that P *lowers* the probability of C (as compared with its initial/prior probability). In such cases, there also seems to be a sense in which P can *not* form the basis of a “strong argument in favor of” C . There are many examples of this sort in the literature. Some of the most interesting examples fall under the rubric “the base rate fallacy” (Kahneman and Tversky 1982; Levi 1983).⁸ For instance, let P be the claim that a positive test result has been obtained from an imperfect but reliable test for some rare disease (performed on a sample drawn from some individual a). And, let C be the claim that a does *not* have the rare disease in question. If the disease is sufficiently rare (i.e., if $\Pr(C)$ is sufficiently high), then $\Pr(C|P)$ will (still) be very high. Nonetheless, $\Pr(C|P)$ will be *less* than $\Pr(C)$, and so P will *lower* the probability of C . In such a case, it would be odd to say that P constitutes strong evidence in favor of C . The problem here is that “argument strength” has (at least) two probabilistic components: confirms_f and confirms_i , and these two components can *run in opposite directions*. This is what happens in “base rate fallacy” cases (and many other cases besides – see Fitelson (2015) for further discussion). Such cases reveal one important theoretical difference between confirms_f and confirms_i .

These two concepts also exhibit some more subtle theoretical divergences that will be important for our analysis of Simpson’s Paradox. In order to grasp these more subtle theoretical differences, it is useful to consider the following principle regarding “argument strength.”

The Suppositional Sure Thing Principle (SSTP).⁹ If P constitutes a strong argument for C on the supposition that S and P constitutes a strong argument for C on the supposition that \bar{S} , then P constitutes a strong argument for C (*unconditionally*).

Basically, the SSTP involves a sort of suppositional “reasoning by cases.” It asks us to consider whether (a) the argument $P \therefore C$ is strong, given some (indicative) supposition S and (b) the argument $P \therefore C$ is strong, given *the negation of* that (indicative) supposition (viz., \bar{S}). And, if both of these assessments come out positive, then the SSTP implies that $P \therefore C$ is strong *unconditionally* (i.e., in the *absence* of background suppositions).

This kind of suppositional “reasoning by cases” is correct for some understandings of “argument strength” but not for others. Specifically, SSTP is satisfied by confirmation as firmness (confirms_f), but *not* by confirmation as increase in firmness (confirms_i). That is to say, we have the following two key theoretical facts.¹⁰

7 The idea that “argument strength” is best explicated in terms of confirmation as firmness (confirms_f) is a very popular one (historically). For instance, Skyrms (2000) defends this approach. In general, I think it is a mistake to ignore considerations of increase in firmness (confirms_i) when assessing “argument strength.” See Fitelson (2015) for a detailed discussion regarding the history of probabilistic explications of “argument strength.”

8 The $\text{confirms}_f/\text{confirms}_i$ distinction is useful for understanding what is going on in *various* “reasoning fallacies” that crop up in the contemporary cognitive psychology literature. See Crupi *et al.* (2008, 2013) for explanations (away) of the so-called “conjunction fallacy” which make essential use of this distinction.

9 The “sure thing” moniker is borrowed from Savage (1972), who uses it to describe a similar “reasoning by cases” property of preference relations.

10 See the Appendix for elementary proofs of these two key Facts, as well as Fact 3 below.

Table 1 Probabilistic structure of canonical (BGS) example of Simpson's Paradox.

	M	\bar{M}	Overall
E	$60/80 = \Pr(A E \& M)$	$40/50 = \Pr(A E \& \bar{M})$	$100/130 = \Pr(A E)$
\bar{E}	$10/50 = \Pr(A \bar{E} \& M)$	$20/80 = \Pr(A \bar{E} \& \bar{M})$	$30/130 = \Pr(A \bar{E})$
Overall	$70/130 = \Pr(A M)$	$60/130 = \Pr(A \bar{M})$	$130/260 = \Pr(A)$

Fact 1. Confirmation_f satisfies SSTP. In other words, it is generally true that if (a) P confirms_f C , relative to S and (b) P confirms_f C , relative to \bar{S} , then (c) P confirms_f C , unconditionally.

Fact 2. Confirmation_i violates SSTP. In other words, there exist cases/contexts in which (a) P confirms_i C , relative to S and (b) P confirms_i C , relative to \bar{S} ; but, (c) P does not confirm_i C , unconditionally.

Fact 2 is the theoretical property of confirms_i which makes Simpson's Paradox possible in the first place. Indeed, Simpson's Paradox cases just are the examples which establish the truth of Fact 2 (e.g., the simplified BGS example depicted in Table 1 above).

While confirms_f and confirms_i disagree regarding the SSTP, they agree on a similar (but subtly different) theoretical principle regarding "argument strength." To wit:

The Conjunctive Sure Thing Principle (CSTP). If $P \& S$ constitutes a strong argument for C (unconditionally) and $P \& \bar{S}$ constitutes a strong argument for C (unconditionally), then P constitutes a strong argument for C (unconditionally).

Like the SSTP, the CSTP also involves "reasoning by cases." However, unlike the SSTP (which involves suppositional reasoning), the CSTP involves conjunctive (and non-suppositional) "reasoning by cases." That is, the CSTP does not ask us to assess the strength of a single argument $P : C$, relative to a pair of suppositions (viz., S and \bar{S}). Rather, the CSTP asks us to consider the unconditional strength of a pair of arguments ($P \& S : C$ and $P \& \bar{S} : C$). Interestingly, both confirmation_f and confirmation_i satisfy the CSTP. That is, we have:

Fact 3. Confirmation_f and confirmation_i satisfy CSTP. In other words, it is generally true that if (a) $P \& S$ confirms C and (b) $P \& \bar{S}$ confirms C , then (c) P confirms C .¹¹

The disagreement between confirms_f and confirms_i with respect to SSTP vs. CSTP can be understood in terms of their ability to distinguish conjunctive vs. suppositional support/confirmation. Confirmation as firmness simply cannot distinguish between the two. If you want to know whether P confirms_f C , on the supposition that S , you check whether $\Pr(C|P \& S) > t$. And, this is exactly the same thing you must do if you want to check whether $P \& S$ confirms_f C , unconditionally. In other words, from the point of view of confirmation as firmness, there is no distinction between conjunctive and suppositional support (in general). In light of this, it is easy to see why Facts 1 and 3 both obtain (for confirms_f) – they are actually one and the same fact, from the perspective of

¹¹ Unsurprisingly, it's also true that deductive entailment satisfies CSTP. That is, if $P \& S = C$ and $P \& \bar{S} = C$, then $P = C$. Thus, on pretty much any way of explicating "argument strength," the CSTP will come out true.

confirmation_r. This is not so for confirmation_i, which distinguishes suppositional and conjunctive support in such a way that allows for Simpson's-Paradoxical cases (i.e., violations of the SSTP) to arise.

Facts 1–3 are the key theoretical facts that will undergird my (ultimate) rationalizing explanation of Simpson's Paradox. In the next section, I will examine two recent discussions of Simpson's Paradox. This will lead naturally into a new rationalizing explanation, based on Facts 1–3 above.

3. TWO RECENT DISCUSSIONS OF SIMPSON'S PARADOX

It is clear that there can be nothing truly *paradoxical* about "Simpson's Paradox." After all, examples like Cartwright's BGS example *actually occur in nature* (all the time). So, what needs to be explained (or perhaps explained *away*) is the *appearance* of paradoxicality in cases of Simpson's Paradox (i.e., their *prima facie* paradoxicality). In this section, I'll discuss two recent treatments of Simpson's Paradox. Then, I'll propose a new, purely confirmation theoretic rationalizing explanation.

3.1 Pearl's Causal Rationalizing Explanation

Pearl (2009: §6.1) provides a *causal* rationalizing explanation of Simpson's Paradox, according to which its apparent paradoxicality involves the conflation of an invalid argument involving confirmation_i and a valid argument involving *causation*. Let's use our BGS example to illustrate Pearl's approach. In that example, there is a correlation between being female and being rejected from graduate school. But, this correlation is *spurious*. It is explained by the fact that females tend to apply to departments with lower acceptance rates. Furthermore, this "tendency" is *not* itself *causal*. How can we tell? Pearl would explain this by postulating that if we were to *intervene* on an applicant's gender, this would not have an effect on the probability that they apply to one department rather than another. In Pearl's notation, this amounts to the following assumption:

$$(\dagger) \Pr(E|do(M)) = \Pr(E|do(\bar{M})) = \Pr(E).^{12}$$

The main idea behind Pearl's rationalizing explanation is that people (to the extent they see the BGS case as *paradoxical*) are (falsely) assuming that the correlations involved all reflect *causal* relationships. In Pearl's notation, we can express the three key correlations in the BGS case as *causal* relations, in the following way:

- (1) $\Pr(A|do(M)\&E) < \Pr(A|do(\bar{M})\&E)$
- (2) $\Pr(A|do(M)\&\bar{E}) < \Pr(A|do(\bar{M})\&\bar{E})$
- (3) $\Pr(A|do(M)) > \Pr(A|do(\bar{M}))$

12 Here, "*do*(X)" means "the value of X is set by an *intervention* on X." Where, Pearl-style interventions are ways of setting the value of X, which cancel any other causal influences that may be upstream from X. Thus, interventions are ways of revealing whether there really is a direct causal connection between X and some potential (downstream) effect of X. Here, the truth of (†) would imply that the "tendency" of females to apply to more selective departments is *not* a *causal* relationship.

Claim (1) asserts that being male *causally decreases* the chance of acceptance – within department *E*. Claim (2) asserts that being male *causally decreases* the chance of acceptance – within department \bar{E} . And, claim (3) asserts that being male *causally increases* the chance of acceptance – in the overall population.

It can be shown (1)–(3) *cannot all be true*, if gender has no causal influence on which department one applies to (i.e., if claim (†) is true). That is, assuming (†), it is impossible for the three key (Simpson-reversal) correlations in the BGS case to reflect *causal* relationships. And, according to Pearl, this explains the sense of paradoxicality one may feel when one first hears about (e.g.) the BGS case. For *if* all of the correlations involved in the reversal *were* causal, *then* – assuming (†) – this *would* be paradoxical (as it would involve an *inconsistent* set of probabilistic claims). Thus, Pearl's account predicts that once one learns the underlying causal structure of the case (e.g., that condition (†) holds in the case, etc.), one's sense of paradoxicality should disappear.¹³

Pearl's causal rationalizing explanation is interesting and illuminating. But, I think it has a couple of potential shortcomings. First, it is not a purely confirmation-theoretic explanation. Simpson's Paradox does not essentially involve any causal concepts. It only (essentially) involves relations of confirmation._r As such, it would be nice to have a purely confirmation-theoretic rationalizing explanation. Second, while Pearl's account does involve the conflation of a valid argument (his causal argument) and an invalid one (the purely confirmation_r-theoretic one), the valid argument is *enthymematic*. That is, Pearl's causal paradox does not just involve the correlations that are given in the statement of the paradox [i.e., (1)–(3)] it also involves an *unstated causal independence* premise/assumption [i.e., (†)] which is not essential to the statement of Simpson's Paradox. In these two senses, Pearl's rationalizing explanation brings in considerations that are not essential to the Paradox (*per se*). It would be preferable (*ceteris paribus*) to have a rationalizing explanation that appeals only to information that is given (essentially) in the statement of the paradoxical cases.

Next, I will discuss Kotzen's (2013) recent take on Simpson's Paradox. That will lead us toward our purely confirmation-theoretic rationalizing explanation.

3.2 Kotzen on Simpson's Paradox and The Miners Puzzle

Kotzen (2013) draws an analogy between Simpson's Paradox and the Miners Puzzle (Kolodny and MacFarlane 2010). For our purposes, the key feature of Kotzen's approach is the reconstruction of Simpson's paradoxical cases as violations of the following classically valid form of deductive inference:

13 Miriam Schoenfield (personal communication) has conjectured a similar, causal rationalizing explanation. Hers is expressed in terms of subjunctive conditional probabilities (Joyce 1999) rather than Pearl's (2009) *do*-calculus. But, like Pearl, Schoenfield will need to assume that the causal independence expressed by (†) holds in her examples. Schoenfield has done some preliminary experiments which indicate a decrease in the sense of paradoxicality of Simpson's Paradox cases, when one learns about the causal structure of the case. It would be interesting to do some experiments which also involve my suppositional/conjunctive distinction [and which involve possible violations of (†)], so as to tease apart the roles of causal and non-causal influences on our sense of paradoxicality in cases of Simpson's Paradox. This is fertile ground for work in experimental philosophy (and the psychology of inductive/causal inference). See, also, fn. 17.

If S , then P confirms _{i} Q .
 If \bar{S} , then P confirms _{i} Q .
 Therefore, P confirms _{i} Q .

While this form of argument is valid for the *material* conditional (i.e., if \lceil if p , then q \rceil is interpreted as $\lceil p \supset q \rceil$), it may not be valid for other kinds of conditionals (e.g., the English indicative conditional). Some authors have taken the Miners Puzzle to provide reasons to think that the English indicative does not obey various classically valid forms of deductive inference (e.g., the form of inference above, or even *modus ponens*), when the consequent(s) of the conditionals contain various sorts of *modal* expressions (Kolodny and MacFarlane 2010). Similarly, Kotzen presents Simpson's Paradox as having an analogous structure/upshot – but with respect to consequent(s) containing *confirmation _{i} -theoretic* expressions.

While Kotzen's aim is not to provide a rationalizing explanation of why Simpson's Paradox seems paradoxical in the first place, I think his analogy with the Miners Puzzle puts us on the right track toward formulating such an explanation.

4. A NEW RATIONALIZING EXPLANATION OF SIMPSON'S PARADOX

I think Kotzen's discussion leads us in the direction of a *purely confirmation-theoretic* rationalizing explanation of Simpson's Paradox. I propose that the *prima facie* paradoxicality of Simpson's Paradox can be explained (at least in part) by appealing to the (subtle) distinction between *conjunctive* confirmation (or support) and *suppositional* confirmation (or support) that was implicit in our discussion of SSTP *vs.* CSTP above. To wit, contrast the following:

Suppositional Reading. If S , then P confirms C .
Conjunctive Reading. P & S confirms C .¹⁴

I propose that there is a tendency for hearers of descriptions of Simpson's Paradox cases to *conflate* these two readings. Let's apply this to our BGS case. To make things more concrete, let us now suppose that department \bar{E} is History. Then, we may describe the correlation between being female and being accepted to History, in the following two ways:

Suppositional. If one applies to History, then being female improves one's chances of acceptance.
Conjunctive. Being a female History applicant improves one's chances of acceptance.

14 Readers familiar with controversies regarding the *import-export* principle for the indicative conditional (Gillies 2012; Fitelson 2015) will notice that there are some structural similarities between these two controversies. In fact, if one adopts a semantics for the indicative conditional that involves a probabilistic relevance/confirmation _{i} component (Douven 2009), then the two controversies become (almost) *one and the same*. This is important, since many people think that import-export is *true* (for the English indicative), which suggests that people may well be susceptible to the kind of conflation I exploit in my rationalizing explanation, below.

In our BGS example, it is *only* the **Suppositional** reading that is true.¹⁵ Moreover, as we have seen, these two readings – while similar-sounding to the (untrained) ear – make all the difference, in terms of *paradoxicality*. If we slide from **Suppositional** to **Conjunctive**, then we are sliding from an interpretation on which the examples described are *possible* (because of Fact 2), to one on which they are *impossible* (because of Fact 3). And, this will be true for *any* example of Simpson's Paradox – even ones in which Pearl's causal independence assumption (†) *fails*. So, if my proposal is right, then we can exploit the fact that the *conjunctive* STP is a *truism* about confirmation (no matter how it is explicated) to furnish an explanation-*away* of the *apparent* paradoxicality of Simpson's Paradox.

Here's how that explanation would go. Examples of Simpson's Paradox involve violations of the SSTP. To the extent that hearers tend to conflate suppositional *vs.* conjunctive readings of the key locutions involved in descriptions of such cases, there will be a tendency to hear such cases as violations of the CSTP. But, *there can be no violations of the CSTP* (since, as Fact 3 reveals, it is a general truism about confirmation – no matter how it is explicated). Therefore, to the extent that hearers tend to conflate suppositional *vs.* conjunctive readings of the key locutions involved in descriptions of such cases, there will be a tendency for hearers to interpret the cases as implying something *paradoxical* (i.e., as somehow *contradicting a theoretical truism*).¹⁶ But, because there can be no such cases, the appearance of paradoxicality is a *mere* appearance (arising, at least in part, from said conflation).

Our rationalizing explanation has a couple of virtues, in comparison with Pearl's causal rationalizing explanation. First, ours is *purely confirmation-theoretic* – it makes no appeal to causal relations, which are *inessential* to the description of Simpson's Paradoxical cases. Moreover, Pearl's explanation relies on an *implicit causal independence assumption* [(†)] which is not stated in the descriptions of the Paradox. Our explanation applies to *all* instances of Simpson's Paradox – even ones in which Pearl's assumption (†) *fails*.¹⁷ In any case, even if causal facts/presuppositions are playing a role in the (*prima facie*) paradoxicality of Simpson's Paradox, my proposal is that the suppositional/conjunctive distinction/conflation also has (independent, rationalizing) explanatory relevance/force.

15 In our BGS example, being a female history applicant actually *lowers* one's probability of acceptance, since $\Pr(A|\bar{E}\&\bar{M}) = 1/4 < 1/2 = \Pr(A)$. So, the **Conjunctive** reading is *false* in the BGS example.

16 One reason for this conflation might be people's (rather general) tendency to conflate confirms_j and confirms_i (recall, *only* confirms_i is *capable of distinguishing* conjunctive and suppositional confirmation). This sort of conflation has been seen in other contexts as well. See Fitelson (2005) and Crupi *et al.* (2008, 2013) for further discussion. Another possible reason for the conflation is the general acceptance of the truth of import-export for the indicative conditional (Gillies 2012). See, also, fn. 14.

17 It is important to note that Pearl's assumption (†) is *very strong*. It requires that $\Pr(E|do(M)) = \Pr(E|do(\bar{M})) = \Pr(E)$. I suspect there will be cases in which *either* $\Pr(E|do(M)) \neq \Pr(E|do(\bar{M}))$, or $\Pr(E|do(\bar{M})) \neq \Pr(E)$, or $\Pr(E|do(M)) \neq \Pr(E)$; but, some (residual) sense of paradoxicality is *nonetheless* felt in hearers of the (non-causal description of) the case. Pearl's explanation (and other similar causal explanations) *goes silent* in *all* such cases, whereas ours still applies (equally well). These sorts of empirical questions would be interesting to explore, experimentally. See, also, fn. 13.

APPENDIX: PROOFS OF FACTS I–3

For our proofs, we will adopt the algebraic style/methods of Fitelson (2008). The following *stochastic truth table* (Fitelson 2008) represents all possible probability distributions over the $\{P, C, S\}$ language, via the 7 real variables $a, b, c, d, e, f, g \in [0, 1]$.

P	C	S	$\Pr(\cdot)$
T	T	T	a
T	T	F	b
T	F	T	c
T	F	F	d
F	T	T	e
F	T	F	f
F	F	T	g
F	F	F	$1 - (a + b + c + d + e + f + g)$

Fact 1. Confirmation _{f} satisfies SSTP. In other words, it is generally true that if (a) P confirms _{f} C , relative to S and (b) P confirms _{f} C , relative to \bar{S} , then (c) P confirms _{f} C , unconditionally.

Proof. Fact 1 asserts that the following argument in probability calculus is valid (where t is some positive threshold value on the unit interval).

$$\begin{array}{l}
 (\mathcal{A}_1) \quad (1) \Pr(C|P\&S) > t. \\
 \quad \quad (2) \Pr(C|P\&\bar{S}) > t. \\
 \quad \quad \therefore (3) \Pr(C|P) > t.
 \end{array}$$

Algebraically (using the above setup (Fitelson 2008)), \mathcal{A}_1 is equivalent to:

$$\begin{array}{l}
 (\mathcal{A}_1) \quad (1) \frac{a}{a+c} > t. \\
 \quad \quad (2) \frac{b}{b+d} > t. \\
 \quad \quad \therefore (3) \frac{a+b}{a+b+c+d} > t.
 \end{array}$$

Cross-multiplying (1) & (2) yields:

$$a > (a + c) \cdot t \tag{1}$$

$$b > (b + d) \cdot t \tag{2}$$

Adding the lhs & rhs of these and collecting t yields

$$a + b > (a + b + c + d) \cdot t$$

which is equivalent to (3). □

Fact 2. Confirmation _{f} violates SSTP. In other words, there exist cases/contexts in which (a) P confirms _{f} C , relative to S and (b) P confirms _{f} C , relative to \bar{S} ; but, (c) P does *not* confirm _{f} C , unconditionally.

Proof. The BGS example depicted in Table 1 is sufficient to establish this claim. For completeness, I include here the full stochastic truth table representation of the probability model depicted in Table 1 above (in this example, the relevant instantiation of the SSTP schema is: $P \stackrel{\text{def}}{=} \bar{M}$, $S \stackrel{\text{def}}{=} E$, and $C \stackrel{\text{def}}{=} A$).

A	E	M	Pr(·)
T	T	T	3/13
T	T	F	2/13
T	F	T	1/26
T	F	F	1/13
F	T	T	1/13
F	T	F	1/26
F	F	T	2/13
F	F	F	3/13

It is easily verified that, in this model, the following three constraints are satisfied:

(i) \bar{M} confirms_i A, relative to E.

$$\Pr(A|\bar{M}\&E) = 4/5 > 3/4 = \Pr(A|M\&E)$$

(ii) \bar{M} confirms_i A, relative to \bar{E} .

$$\Pr(A|\bar{M}\&\bar{E}) = 1/4 > 1/5 = \Pr(A|M\&\bar{E})$$

(iii) \bar{M} disconfirms_i A, unconditionally.

$$\Pr(A|\bar{M}) = 6/13 < 7/13 = \Pr(A|M)$$

Thus, the BGS example illustrates the failure of confirms_i to satisfy the SSTP. □

Fact 3. Confirmation_f and confirmation_i satisfy CSTP. In other words, it is generally true that if (a) P&S confirms C and (b) P& \bar{S} confirms C, then (c) P confirms C.

Proof. We have already shown that confirmation_f satisfies CSTP, since this is equivalent to showing that argument form \mathcal{A}_1 , above, is valid.

To show that confirmation_i satisfies CSTP, we just need to show that the following argument form is valid.

$ \begin{aligned} (\mathcal{A}_2) \quad & (1) \Pr(C P\&S) > \Pr(C). \\ & (2) \Pr(C P\&\bar{S}) > \Pr(C). \\ & \therefore (3) \Pr(C P) > \Pr(C). \end{aligned} $

Algebraically (using the above setup (Fitelson 2008)), \mathcal{A}_2 is equivalent to

$ \begin{aligned} (\mathcal{A}_2) \quad & (1) \frac{a}{a+c} > a+b+e+f. \\ & (2) \frac{b}{b+d} > a+b+e+f. \\ & \therefore (3) \frac{a+b}{a+b+c+d} > a+b+e+f. \end{aligned} $

Cross-multiplying (1) & (2) yields

$$a > (a + c) \cdot (a + b + e + f) \quad (1)$$

$$b > (b + d) \cdot (a + b + e + f) \quad (2)$$

Adding the lhs & rhs of these and collecting t yields

$$a + b > (a + b + c + d) \cdot (a + b + e + f)$$

which is equivalent to (3).¹⁸

REFERENCES

- Carnap, R. 1962. *Logical Foundations of Probability*, 2nd edition. Chicago, IL: University of Chicago Press.
- Cartwright, N. 1979. 'Causal Laws and Effective Strategies.' *Noûs*, 13: 419–37.
- Crupi, V., Tentori, K. and Fitelson, B. 2008. 'Probability, Confirmation and the Conjunction Fallacy.' *Thinking and Reasoning*, 14: 182–99.
- , — and — and Russo, S. 2013. 'On the Determinants of the Conjunction Fallacy: Probability vs. Inductive Confirmation.' *Journal of Experimental Psychology: General*, 142: 235–55.
- Douven, I. 2009. 'The Evidential Support Theory of Conditionals.' *Synthese*, 164: 19–44.
- Fitelson, B. 2005. 'Inductive Logic.' In J. Pfeifer and S. Sarkar (eds), *The Philosophy of Science: An Encyclopedia*, pp. 384–94. Oxford: Routledge.
- 2008. 'A Decision Procedure for Probability Calculus with Applications.' *Review of Symbolic Logic*, 1: 111–25.
- 2015. 'The Strongest Possible Lewisian Triviality Result.' *Thought*, 4: 69–74.
- Fitelson, B. and Hawthorne, J. 2010. 'The Wason Task(s) and the Paradox of Confirmation.' *Philosophical Perspectives*, 24: 207–41.
- Gillies, A. 2012. 'Indicative Conditionals.' In D. Graff Fara and G. Russell (eds), *The Routledge Companion to Philosophy of Language*, pp. 449–65. Oxford: Routledge.
- Joyce, J. 1999. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Kahneman, D. and Tversky, A. 1982. 'Evidential Impact of Base Rates.' In D. Kahneman, P. Slovic and A. Tversky (eds), *Judgment under Uncertainty: Heuristics and Biases*, pp. 153–60. Cambridge: Cambridge University Press.
- Kolodny, N. and MacFarlane, J. 2010. 'Ifs and Oughts.' *Journal of Philosophy*, 107: 115–43.
- Kotzen, M. 2013. 'Conditional Oughts and Simpson's Paradox.' Unpublished manuscript.
- Levi, I. 1983. 'Who Commits the Base Rate Fallacy?' *Behavioral and Brain Sciences*, 6: 502–6.
- Malinas, G. and Bigelow, J. 2016. 'Simpson's Paradox.' In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/sum2016/entries/paradox-simpson/>.
- Pearl, J. 2009. *Causality: Models, Reasoning, and Inference*, 2nd edition. Cambridge: Cambridge University Press.
- Savage, L. 1972. *The Foundations of Statistics*. New York, NY: Dover.
- Skyrms, B. 2000. *Choice and Chance: An Introduction to Inductive Logic*, 4th edition. Belmont, CA: Wadsworth.

¹⁸ I would like to thank Miriam Schoenfeld for very useful feedback/commentary on a previous draft of this paper. I would also like to thank the MIT philosophy department and the participants of the 2016 Episteme Conference for helpful discussions of this material.

BRANDEN FITELSON is Distinguished Professor of Philosophy at Northeastern University and visiting Professor at both the Munich Center for Mathematical Philosophy at LMU-Munich (MCMP @ LMU) and the Institute for Logic, Language and Computation at the University of Amsterdam (ILLC @ UvA). Before teaching at Northeastern, LMU, and UvA, Branden held teaching positions at Rutgers, UC-Berkeley, San José State, and Stanford. Branden got his MA and PhD in philosophy from UW-Madison. Before entering philosophy, Branden studied math and physics at Wisconsin, and he worked as a research scientist at Argonne National Laboratory and a NASA contractor.
