Supplemental Materials to "Following the FAD"
Jonathan Livengood, Justin Sytsma, and David Rose

1. Following the FAD

Having articulated the theories of Halpern and Pearl (2005), Hitchcock (2007a), and Halpern and Hitchcock (2015), we are left with a critical question: *Do any of these theories deliver the correct causal attributions?* To answer this, we need a way to independently determine the correct causal attributions for a given scenario. This is *the* problem of actual causation. And it is not one that we solve in this article. Rather, we assess the theories we're considering against folk causal attributions. We do so because these theories, and DTN accounts in general, are committed to the folk attribution desideratum (FAD), according to which the deliverances of a theory of actual causation are correct insofar as they agree with folk causal attributions.

In laying out the DTN terminology that we've adopted, Halpern and Hitchcock make it clear that they take such accounts to follow the FAD:

> One approach to solving these problems that has been gaining increasing popularity… is to incorporate 'defaults', 'typicality', and 'normality' into an account of actual causation. These approaches gain further support from empirical results showing that such considerations do in fact influence people's judgments about actual causation. (2015, 415)

In line with this, they suggest that one goal of giving an account of actual causation it that "doing so lets us explain the responses that people make to queries regarding actual causation" (436). And Halpern and Hitchcock conclude their paper with the hope that "future empirical work will examine more carefully how people use normality in judgments of actual causation" (453)—a project that we take this paper to be contributing to.

Further, it is equally clear that the other two theories of actual causation that we are examining also follow the FAD. For example, Halpern and Pearl (2005) write that "the best ways

to judge the adequacy of an approach are the intuitive appeal of the definitions and how well it deals with examples" (846). With regard to the examples they consider, they again appeal to intuition. For instance, with regard to the deliverance of their theory for their first example (4.1), they write that "this seems to us intuitively right" (860). And in discussing their second example (4.2) they assert that "common sense suggests" a specific causal attribution. We believe that in these examples Halpern and Pearl are appealing to (their assessment of) the deliverances of common sense or ordinary intuition and that failing to satisfy the FAD would be evidence against their theory.

Hitchcock (2007a) is also clear in his endorsement of the FAD. For instance, he notes that his theory of actual causation is not entirely objective owing to its reliance on the distinction between default and deviant values of a variable. He then considers the possibility of making his theory of actual causation entirely objective by appealing to fundamental laws and using those laws to determine the default and deviant values of the variables in the causal model. However, Hitchcock then rejects this move, writing, "perhaps a case could be made for allowing only genuine laws of nature to determine default values of variables, *but if we disallow folk theories, we are not likely to arrive at a theory that accords with folk intuitions*" (506, emphasis added). The consequent of the conditional in this quotation is supposed to provide reason to reject the suggestion in the antecedent. Hitchcock wants to allow folk theories to determine default and deviant values precisely because he wants to arrive at a theory that respects what people say about actual causation.

And a similar commitment to the FAD can be found in other work by Hitchcock, including his 2009 article with Knobe and the article with Halpern discussed above. Further, such a commitment is consistent with Hitchcock's (2007b) discussion of three distinct concepts

of causation—scientific, folk attributive, and metaphysical. Our argument is simple. First, Hitchcock's (2007a) theory of actual causation cannot be a theory of the scientific concept, which Hitchcock (2007b, 510) identifies with the causal structures characteristic of graphical causal modeling. But Hitchcock tells us that *token* causal relations (by which he means what we are calling *actual* causal relations) cannot be read off from a causal model directly (511). Since in his (2007a), Hitchcock produces a theory of token causation, we think it is plausible to read his (2007b) as asserting that *actual* causal relations cannot be read directly from causal models, and hence, a theory of actual causation is not a theory of the scientific concept. Second, Hitchcock's theory of actual causation cannot be a theory of the metaphysical concept, for the sake of pragmatic consistency. Since Hitchcock (2007b) claims that the metaphysical concept of causation is an unstable and dispensable hybrid of the scientific and folk-attributive concepts, it would be pragmatically inconsistent to offer a metaphysical account of causation while simultaneously claiming that no one should be offering metaphysical accounts of causation. Finally, Hitchcock's theory of actual causation is congruent with the folk-attributive concept. Both are retrospective. Both are (or could be) employed to facilitate expressions of praise and blame. Neither makes predictions about what happens under interventions; instead, both make predictions about what people will say. And both are discriminatory. Given that Hitchcock (2007b) is attempting to characterize the folk attributive concept of causation, requiring that his theory satisfy the FAD is a plausible demand.

2. Vignettes

Studies 1-4, 7, and 8: Original Lauren and Jane Vignette

> Lauren and Jane both work for a company that uses a mainframe that can be accessed from terminals on different floors of its building. The mainframe has recently become

unstable, so that if more than one person is logged in at the same time, the system crashes. Therefore, the company has instituted a temporary policy restricting the use of terminals so that two terminals are not used at the same time until the mainframe is repaired. The policy prohibits logging into the mainframe from the terminal on any floor except the ground floor.

One day, Lauren logged into the mainframe on the authorized terminal on the ground floor at the exact same time that Jane logged into the mainframe on the unauthorized terminal on the second floor. Lauren and Jane were both unaware that the other was logging in. Sure enough, the system crashed.

Studies 5 and 6: Feature not a bug

Lauren and Jane both work for a company that uses a mainframe that can be accessed from terminals on different floors of its building. The operating system for the mainframe was designed to only allow a single user to log into the mainframe at a given time, however. As such, if more than one person logs into the mainframe at the same time, the system crashes. Because of this, the company has a policy restricting the use of terminals so that two terminals are not used at the same time. The policy prohibits logging in to the mainframe from the terminal on any floor except the ground floor during business hour.

One day, Lauren logged into the mainframe on the authorized terminal on the ground floor during business hours; at the exact same time, Jane logged into the mainframe on the unauthorized terminal on the second floor. Lauren and Jane were both unaware that the other was logging in. Sure enough, the system crashed.

Studies 9-14: Without permissibility information

Lauren and Jane both work for a company that uses a mainframe that can be accessed from terminals on different floors of its building. Though the company does not know it, the mainframe has recently become unstable, so that if more than one person is logged in at the same time, the system crashes.

One day, Lauren logged into the mainframe on the ground floor at the exact same time that Jane logged into the mainframe on the second floor. Lauren and Jane were both unaware that the other was logging in. Sure enough, the system crashed.

Studies 15 and 16: Feature not a bug, without permissibility information

Lauren and Jane both work for a company that uses a mainframe that can be accessed from terminals on different floors of its building. The operating system for the mainframe was designed to only allow a single user to log into the mainframe at a given time,

however. As such, if more than one person logs into the mainframe at the same time, the mainframe crashes.

One day, Lauren logged into the mainframe on the ground floor at the exact same time that Jane logged in to the mainframe on the second floor. Lauren and Jane were both unaware that the other was logging in. Sure enough, the mainframe crashed.

Study 17: Lauren alone (second page same on each)

Without Typicality Information:

First page: Lauren works for a company that uses a mainframe that can be accessed from terminals in its building. Though the company does not know it, the mainframe has recently become unstable, so that if anyone logs into the system, the system crashes.

One day, Lauren logged into the mainframe. Sure enough, the system crashed.

Second page: When the company realized that the mainframe was unstable, they implemented a temporary policy prohibiting employees from logging into the mainframe until it is repaired. A memo was sent to each employee informing them of the new policy. Each employee was required to return a signed form indicating that they had read and understood the policy. As an employee of the company, Lauren received the memo and returned the signed form.

Despite having done so, later that week Lauren logged into the mainframe. Sure enough, the system crashed.

Agent-level Typical: Lauren works for a company that uses a mainframe that can be accessed from terminals in its building. Though the company does not know it, the mainframe has recently become unstable, so that if anyone logs into the system, the system crashes.

Lauren often logs into the mainframe at work, since her job requires the use of a computer. One day, Lauren logged into the mainframe to do some work. Sure enough, the system crashed.

Agent-level Atypical: Lauren works for a company that uses a mainframe that can be accessed from terminals in its building. Though the company does not know it, the mainframe has recently become unstable, so that if anyone logs into the system, the system crashes.

Lauren almost never logs into the mainframe at work, since her job rarely requires the use of a computer. However, one day, Lauren logged into the mainframe to do some work. Sure enough, the system crashed.

Population-level Typical: Lauren works for a company that uses a mainframe that can be accessed from terminals in its building, which company employees often use. Though the company does not know it, the mainframe has recently become unstable, so that if anyone logs into the system, the system crashes.

One day, Lauren logged into the mainframe. Sure enough, the system crashed.

Population-level Atypical: Lauren works for a company that uses a mainframe that can be accessed from terminals in its building, but company employees very rarely use it. Though the company does not know it, the mainframe has recently become unstable, so that if anyone logs into the system, the system crashes.

One day, Lauren logged into the mainframe. Sure enough, the system crashed.