

**Two Types of Typicality:  
Rethinking the Role of Statistical Typicality in Ordinary Causal Attributions**  
Justin Sytsma, Jonathan Livengood, and David Rose

**Abstract:** Recent work on the role of norms in the use of causal language by ordinary people has led to a consensus among several researchers: The consensus position is that causal attributions are sensitive to both statistical norms and prescriptive norms. But what is a statistical norm? We argue that there are at least two types that should be distinguished—*agent-level statistical norms* and *population-level statistical norms*. We then suggest an alternative account of ordinary causal attributions about agents (the *responsibility view*), noting that this view motivates divergent predictions about the effect of information about each of the two types of statistical norms noted. Further, these predictions run counter to those made by the consensus position. With this set-up in place, we present the results of a series of new experimental studies testing our predictions. The results are in line with the responsibility view, while indicating that the consensus position is seriously mistaken.

Consider the following scenario based on a thought experiment by Joshua Knobe (2006):

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.

When this scenario is given to people without training in philosophy, they are significantly more likely to say that Jane caused the system to crash than Lauren (Livengood, Sytsma, and Rose,

---

To appear in *Studies in History and Philosophy of Science*.

The authors would like to thank Jill Cumby, David Danks, Ben Fraser, Greg Gandenberger, Joshua Knobe, Craig Roxborough, an anonymous referee for *Studies in History and Philosophy of Science*, and the audiences at the Causality in the Biomedical and Social Sciences, Southern Society for Philosophy and Psychology, and Society for Philosophy and Psychology conferences for their thoughtful comments on an earlier draft of this paper.

2011). And, they do so despite the fact that the actions of both agents were necessary for bringing about the outcome: If either Lauren or Jane had not logged-in, the system would not have crashed.

Results like this have struck many philosophers as quite surprising. They generally take the ordinary concept of causation to be purely descriptive in character; and, yet, ordinary causal judgments seem to be sensitive to broadly moral considerations. Thus, empirical work by Mark Alicke (1992) indicates that ordinary causal attributions are sensitive to whether or not an agent's motives are socially desirable and Knobe (2006) used intuitions about simple cases like that discussed above to argue that ordinary causal attributions are sensitive to prescriptive norms (to whether a behavior is permissible or impermissible).

Not surprisingly, some philosophers have sought to explain such results away, aiming to preserve the assumption that the ordinary concept of causation is purely descriptive in character. For example, Julia Driver (2008a) suggests that ordinary causal attributions might still be best explained in terms of a type of descriptive consideration—statistical norms. Thus, it might be that Jane is identified as the cause of the system crashing, in the above example, not because she violated a prescriptive norm *per se*, but because she did something that is statistically atypical in doing so—it presumably being unusual to violate company policy. Subsequent empirical work has suggested against this, however. In fact, a consensus position has emerged that focuses on norms more generally: Ordinary causal attributions are thought to be directly sensitive both to behaviors being out of the prescriptive norm (impermissible) and to behaviors being out of the statistical norm (atypical). As such, the consensus view takes the effect of prescriptive norms to be part of a wider phenomenon, with ordinary causal attributions being sensitive to whether or not a behavior is *out of the norm*, where this includes both purely descriptive and broadly moral considerations.

We are not convinced by this consensus view, however. Instead, we hold that ordinary causal attributions are much more intimately tied to broadly moral considerations than has been suggested. Our view is that the ordinary concept of causation, at least as applied to agents, is an inherently normative concept: Causal attributions are typically used to indicate something more akin to who is *responsible* for a given outcome than who *caused* the outcome in the descriptive sense of the term used by philosophers. As such, we hold that the current consensus position still goes too far in attempting to preserve the purely descriptive conception of causation. We will not argue directly for this *responsibility view* in this paper; rather, we call on it to motivate two predictions about the role of statistical norms in folk causal attributions that diverge from those given by the consensus view. Focusing on the empirical example that has led to the consensus position (Knobe and Fraser's *Pen Case*), we present the results of a series of studies testing these competing predictions. The results are in line with our predictions and indicate that the consensus view is badly mistaken.

Here is how we will proceed. In Section 1, we review the debate leading up to the consensus position. We then take a closer look at the notion of a statistical norm, distinguishing between two types—*agent-level* and *population-level*. In Section 3, we briefly describe our alternative to the consensus view, and on the basis of it make two competing predictions about the role of statistical norms in ordinary causal attributions. These predictions are empirically tested in Sections 4 through 6.

## **1. The Consensus Position**

What types of information do ordinary people call on in making causal attributions about agents? Many philosophers have assumed that ordinary causal judgments are by and large based on descriptive information. Calling on intuitions about some simple cases, however, Joshua Knobe

(2006) has argued that this is incorrect: Ordinary causal attributions are sensitive to prescriptive norms. In response, Julia Driver (2008a) has suggested that the role of prescriptive norms can be placed within a wider framework when we consider norms more generally. She writes that “we might ask whether or not it was that the agent acted wrongly or, rather, somehow ‘out of the norm’” (430).

Driver’s suggestion can be fleshed out in a couple of different ways. Ordinary causal attributions might be directly sensitive to behaviors that are either out of a statistical norm or out of a prescriptive norm. Alternatively, ordinary causal attributions might be directly sensitive only to behaviors that are out of a statistical norm. Knobe and Fraser (2008) understand Driver to be suggesting the second of these two hypotheses, namely “that it might be possible to explain all of the puzzling results by appealing to the concept of *atypicality*” (2008, 442). To test Driver’s hypothesis they gave people a story—*the Pen Case*—describing two characters (Professor Smith and an administrative assistant) who behave in ways that are typical of the populations to which they belong, but where one of their behaviors is permitted while the other is not.<sup>1</sup> Specifically, each character takes a pen from the receptionist’s desk and these behaviors jointly lead to a problem. Despite both behaviors being statistically typical (being in the statistical norm), Knobe and Fraser found that people were far more likely to judge that Professor Smith had caused the problem. These results appear to indicate against the hypothesis that ordinary causal attributions are only directly sensitive to behaviors being out of a statistical norm, since both characters in the Pen Case are described as acting in a way that is statistically typical.

Note, however, that Knobe and Fraser’s study does not tell against the *disjunctive hypothesis* noted above: It leaves open the possibility that ordinary causal attributions are directly sensitive to behaviors that are either out of a statistical norm or out of a prescriptive

---

<sup>1</sup> Each of the vignettes discussed in this paper, including Knobe and Fraser’s original Pen Case vignette, can be found in the appendices to this paper available on the PhilSci Archive.

norm. Craig Roxborough and Jill Cumby (2009) conducted a further study to directly test this disjunctive hypothesis. They suggest that there are three versions of the Pen Case worth considering—Knobe and Fraser’s original version, plus two variations on it. Knobe and Fraser tested the version of the Pen Case in which the behaviors of *both* characters are statistically typical of their respective populations, but one might also look at ordinary causal attributions for variations where the behavior of *only one* of the characters is statistically typical and the other is statistically atypical. Roxborough and Cumby tested the variation in which Professor Smith’s behavior is statistically typical, while the administrative assistant’s behavior is statistically atypical. They predicted that participants would be more likely to judge that the administrative assistant caused the problem when he acted in a manner that was statistically atypical than when he acted in a manner that was statistically typical. What they found, however, is that there was no significant difference between the mean response for the question about the administrative assistant in their study and the mean response for the same question in Knobe and Fraser’s study. Having found no effect where they were looking for one, Roxborough and Cumby went on to conduct a post hoc comparison of the mean response for the question about Professor Smith in their study and the mean response for the question about Professor Smith in Knobe and Fraser’s study. There is a significant difference.<sup>2</sup>

Roxborough and Cumby take these studies to provide evidence that statistical norms impact ordinary causal attributions, even if they do not do so in the way originally predicted. As such, they advocate the disjunctive hypothesis that ordinary causal attributions are directly sensitive to behaviors being out of either a prescriptive norm or a statistical norm. In fact, this hypothesis is now the consensus position amongst the primary participants in this debate: Driver

---

<sup>2</sup> To compare the means from Roxborough and Cumby’s study with the means from Knobe and Fraser’s study, the appropriate test to use is a Welch, two-sample t-test, which assumes neither equal sample sizes nor equal variances for the two means to be compared. For this test,  $t = -3.3532$ ,  $df = 27.093$ , and the  $p\text{-value} = 0.002370$ .

(2008b) endorses it; Roxborough and Cumby (2009) endorse it; and Hitchcock and Knobe (2009) endorse it.

## 2. Atypicality and Statistical Norms

We have just seen that a consensus position has emerged in the debate about the effect of information about norms on ordinary causal attributions: It is held that all else being equal, people are more likely to say that an agent who behaves in a way that is out of the norm caused an outcome than they are to say that an agent who behaves in a way that is in the norm caused the same outcome. With regard to statistical norms specifically, the consensus position holds that all else being equal, people are more likely to say that an agent who behaves in a way that is statistically atypical caused an outcome than they are to say that an agent who behaves in a way that is statistically typical caused the same outcome.

What does it mean to say that an agent behaves in a way that is statistically typical or atypical? Here are two options. One might say that an agent's behavior is statistically typical/atypical relative to how *people* generally behave in a given type of situation. Or, one might say that an agent's behavior is statistically typical/atypical relative to how *the agent herself* generally behaves in that type of situation. Call the first kind of statistical norm a *population-level statistical norm* and the second kind an *agent-level statistical norm*.

To illustrate, consider the population consisting of philosophy professors at Fictitious State University (FSU). Suppose that 90% of philosophy professors at FSU almost never smoke during breaks in their seminars. In contrast, the remaining 10% of professors almost always smoke during breaks. Suppose that Professor Madeline Madeup is a smoker and that we observe her smoking during a break. Professor Madeup's smoking is *atypical for the population* of philosophy professors to which she belongs. The philosophy professors at FSU do not typically

smoke during breaks. Given that Professor Madeup is a smoker, however, her behavior is *typical for her*.

In what follows, we will say that ordinary causal attributions are *sensitive to behaviors being population-level atypical* if, other things being equal, untutored people tend to be more likely to say that an agent who behaves in a way that is population-level atypical caused an outcome than they are to say that an agent who behaves in a way that is population-level typical caused the outcome. We will follow the same convention for sensitivity to behaviors being population-level typical, agent-level atypical, agent-level typical, impermissible, or permissible.

Having now distinguished between these two types of statistical norms, how should we understand the consensus position? When it is claimed that ordinary causal attributions are directly sensitive to behaviors being out of either a prescriptive norm or a statistical norm, what type of statistical norm is at issue? The answer is not obvious, since the distinction between agent-level statistical norms and population-level statistical norms has not been drawn in the literature. That this distinction has not been drawn, however, suggests that it is statistical norms *generally* that are at issue, not one or the other type specifically. And there are other reasons to think that this is the case. For example, Driver both discusses cases that involve agent-level statistical norms (e.g., 2008a, 430), and takes Knobe and Fraser's Pen Case results to be relevant to her account—even though it concerns population-level statistical norms.<sup>3</sup> As such, we will treat the consensus position as making a claim about statistical norms generally. Nonetheless, it is worth noting that there is a problem with the consensus position however we understand “statistical norm”—as we will see in the following sections.

---

<sup>3</sup> Knobe and Fraser do not describe the behaviors of the characters in the Pen Case as being typical *of those agents*; rather, the vignette only gives information about the typicality of their behaviors relative to the *populations* to which they belong.

### 3. Two Competing Predictions

We do not believe that the consensus position is correct. Specifically, we are not convinced that information about statistical norms *directly* impacts ordinary causal attributions concerning agents. Instead, we conjecture that such attributions primarily reflect people's broadly moral judgments: Causal attributions concerning agents are typically used to express judgments related to normative responsibility. Defending this *responsibility view* is a large project, and we will not attempt to establish its correctness in this paper. Rather, we simply note that this view provides motivation for two predictions about the role of statistical norms in ordinary causal attributions that are directly at odds with the consensus position; we then limit our empirical aims to testing these predictions with respect to the example that has driven the debate—Knobe and Fraser's Pen Case.

By subsuming the impact of information about statistical and prescriptive norms on ordinary causal attributions as different instances of a general phenomenon, the consensus position treats the two kinds of norms as having equal standing. Information about both kinds of norms is thought to have a direct, independent impact on ordinary causal attributions. The responsibility view, by contrast, maintains that in examples like the Pen Case information about statistical norms impacts ordinary causal attributions only indirectly: Information about norms sometimes plays a role in people's broadly moral judgments about normative responsibility, which are expressed in the causal attributions that they make. And these diverging views lead to competing predictions. Specifically, on the basis of the responsibility view we make two predictions about the (indirect) role of statistical norms in ordinary causal attributions that diverge from the (direct) role specified by the consensus position: First, ordinary causal attributions for the Pen Case are insensitive to population-level statistical norms; second, ordinary causal attributions for the Pen Case are sensitive to agent-level typicality, not atypicality.



The reasoning behind our first prediction is that we expect that information about population-level statistical norms will have little impact on judgments about whether the characters are normatively responsible for the problem. More generally, we expect that how other people typically act in a given type of situation will largely be treated as irrelevant to whether or not a specific person is taken to be normatively responsible for an outcome. Thus, while excuses of the form “everybody was doing it” might help to explain an agent’s action, we suspect that people generally do not take such excuses to actually mitigate normative responsibility. As such, we expect that changing the information given about the population-level typicality of a character’s behavior in the Pen Case will not affect how likely people are to say that the character caused the problem.

The reasoning behind our second prediction is that we expect patterns of behavior to be relevant to people’s normative judgments. In cases where the agent could reasonably be expected to know that a bad outcome might result from her behavior, we expect that people will be more likely to judge that the agent is normatively responsible for the outcome when she typically acts in a reckless way. The agent will be judged to be more responsible because her pattern of behavior increased the chance that the bad outcome would eventually occur. For example, compare the case of a person who habitually jaywalks with that of a person who jaywalks only very occasionally. Suppose that each person jaywalks on otherwise identical occasions and that each one’s behavior leads to an accident. We (the authors) have the intuition that the habitual jaywalker is more blameworthy than the occasional jaywalker because her pattern of behavior increased the chance that an accident would eventually occur (in comparison to the occasional jaywalker). Expecting that ordinary causal attributions are used for expressing normative judgments, we therefore conjecture that with regard to scenarios like the Pen Case people will be more likely to say that an agent who behaves in a way that is in the agent-level statistical norm

caused a bad outcome than to say that an agent who behaves in a way that is out of the agent-level statistical norm caused the same bad outcome.

If the two predictions made on the basis of the responsibility view are correct, then the consensus position is seriously mistaken: (1) ordinary causal attributions are *not* sensitive to information about population-level statistical norms at all, and (2) ordinary causal attributions are sensitive to information about agent-level *typicality* (rather than *atypicality*). While the current empirical evidence for the Pen Case suggests that the first prediction is incorrect, we are not convinced that the current empirical data is reliable. We are concerned that the finding that ordinary causal attributions are sensitive to information about population-level atypicality involves a post hoc comparison of the results from two different studies, with the samples drawn from undergraduates in two different countries, and where one of the studies had a relatively small number of participants (N=18). As such, we feel that the results of these studies should not be accepted until they have been replicated. Moreover, the role of information about agent-level statistical norms on ordinary causal attributions has not been tested at all.

#### **4. New Studies on Population-Level Statistical Norms**

In this section, we present two studies that we conducted to test the effect of information about population-level statistical norms on ordinary causal attributions about the Pen Case. The first study includes permissibility information, while the second does not.

##### *4.1 Study 1: The Pen Case Revisited*

We began by collecting responses to the three versions of the Pen Case noted in Section 1.<sup>4</sup>

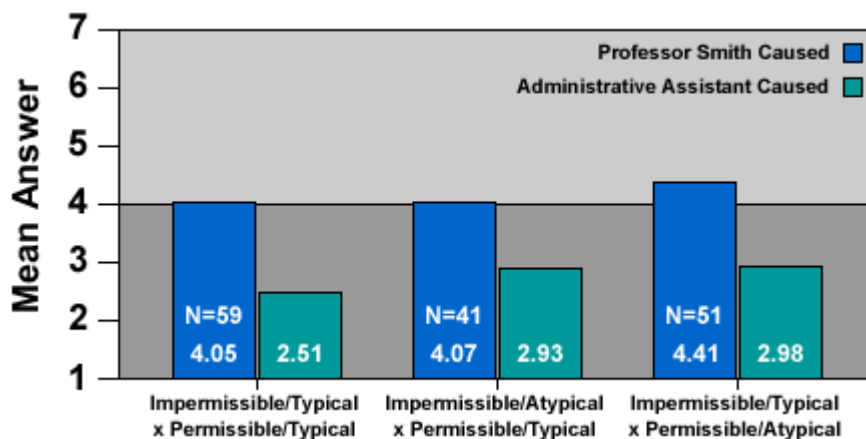
---

<sup>4</sup> The vignettes used for each of the studies described in this paper can be found in the appendices available on the PhilSci Archive. In addition, for ease of readability, further details about the samples for these studies, as well as the statistical analyses performed, are included in those appendices.

This includes the two versions compared by Roxborough and Cumby (2009), as well as the case in which the administrative assistant’s behavior is population-level typical, while the professor’s behavior is population-level atypical. In each version, both the administrative assistant and Professor Smith take pens; further, in each version it is permissible for the administrative assistant to do so, while it is impermissible for Professor Smith to do so.

Responses for these three probes were collected online from 151 native English speakers, 18 years of age or older, with at most minimal training in philosophy. Each participant was randomly assigned to one of the three versions of the Pen Case described above. After reading the vignette, participants were asked to indicate whether they agreed or disagreed with each of two causal claims—“Professor Smith caused the problem” and “the Administrative Assistant caused the problem”—on a 7-point scale anchored at 1 with “strongly disagree,” at 4 with “neutral,” and at 7 with “strongly agree.” The results are shown in Figure 1.

What we found is that information about population-level statistical norms had no significant impact on judgments about who caused the problem. Thus, contrary to the consensus position (but in line with our first prediction from Section 3), participants’ responses were not sensitive to behaviors being population-level atypical.



**Figure 1: Results for Study 1.**

#### 4.2 Study 2: The Pen Case without Permissibility Information

Permissibility information was included in each of the versions of the Pen Case that we ran in Study 1 (as well as in the studies conducted by Knobe and Fraser and by Roxborough and Cumby). Thus, population-level statistical norms might still play a role in ordinary causal attributions about the Pen Case: Perhaps the information about prescriptive norms overrides the information about population-level statistical norms. To test this we ran a second study, removing the information about prescriptive norms from the Pen Case vignettes used in our first study. In addition, we included a fourth variation in which both characters act atypically. Responses for these four vignettes were collected online from 183 native English speakers, 18 years of age or older, with at most minimal training in philosophy. The results are shown in Figure 2.

Once again we found that population-level statistical norms have no discernable effect on ordinary causal attributions for the Pen Case. Thus, in the absence of permissibility information, there was no significant difference between the mean responses for the two statements in any of the conditions. In fact, if we average the responses for each of the questions in which the character behaved population-level typically, we find that it is virtually identical to the average for the responses for the questions in which the character behaved population-level atypically.

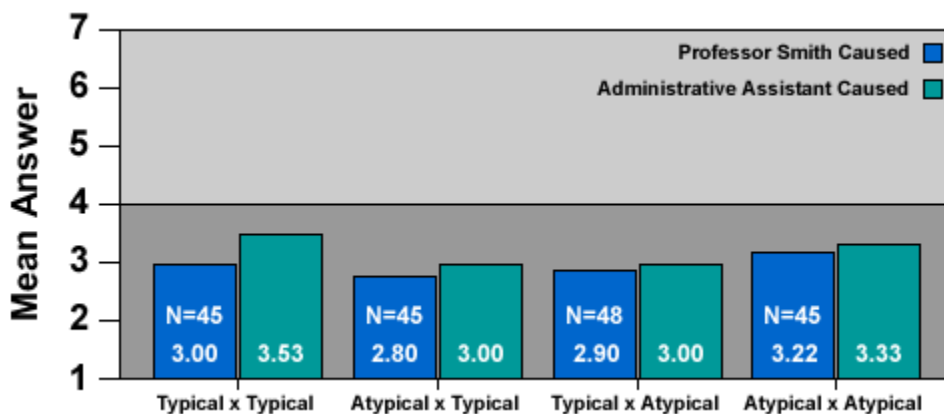


Figure 2: Study 2 results.

## 5. Studies on the Effect of Agent-Level Statistical Norms

In line with the first prediction we made in Section 3, and against the consensus view, the two studies described in the previous section undermine the claim that population-level atypical behaviors are more likely to be judged causes than are population-level typical behaviors.


Recall, however, that we made two predictions: Not only did we predict that ordinary causal attributions for the Pen Case are insensitive to population-level statistical norms, we also predicted that ordinary causal attributions would be sensitive to agent-level statistical norms for examples like the Pen Case. Specifically, we conjectured that people would be more likely to say that a character who behaved *agent-level typically* was the cause of a bad outcome than a character who behaved *agent-level atypically*. To test this predication, we ran a series of further studies in which we rewrote the Pen Case to change the population-level statistical norms to agent-level statistical norms.

### 5.1 Study 3: The Agent-Level Pen Case

In our third study we gave participants four further variations on the Pen Case in which we varied the agent-level typicality of the behaviors while specifying that the administrative assistant (now given the name John) acts permissibly and Professor Smith acts impermissibly. Responses for these four vignettes were collected online from 216 native English speakers, 18 years of age or older, with at most minimal training in philosophy. The results are shown in Figure 3.

We found that, in contrast with our results in Study 1, assignment to a specific vignette affected participant responses. In fact, we found that when information about permissibility and agent-level typicality push in different directions, agent-level typicality information wins out. For example, looking specifically at the case in which Professor Smith acts both agent-level

atypically and impermissibly, while John acts agent-level typically and permissibly, the mean response for the former is actually significantly *lower* than the mean response for the latter. Thus, we find that in line with our second prediction, and counter to the consensus position, people are more likely to say that a character caused the problem when he behaves agent-level typically than when he behaves agent-level atypically. It appears that ordinary causal attributions are *not* sensitive to behaviors being out of the statistical norm for cases like this, but are in fact sensitive to behaviors being in the statistical norm.



**Figure 3: Study 3 results.**

### 5.2 Study 4: The Agent-Level Pen Case without Permissibility Information

To further test the impact of information about agent-level statistical norms on ordinary causal attributions, in our fourth study we removed the permissibility information from the versions of the Pen Case used in Study 3. Responses for these four vignettes were collected online from 167 native English speakers, 18 years of age or older, with at most minimal training in philosophy. The results are shown in Figure 4.

What we find is that when agent-level statistical norms are used instead of population-level statistical norms, assignment to a specific vignette affects participant responses even when

no permissibility information is included. Specifically, participants were more likely to say that a character caused the problem when he behaved agent-level typically than when he behaved agent-level atypically. These results provide further support for the claim that ordinary causal attributions are sensitive to agent-level typicality (not atypicality), at least for examples like the Pen Case.

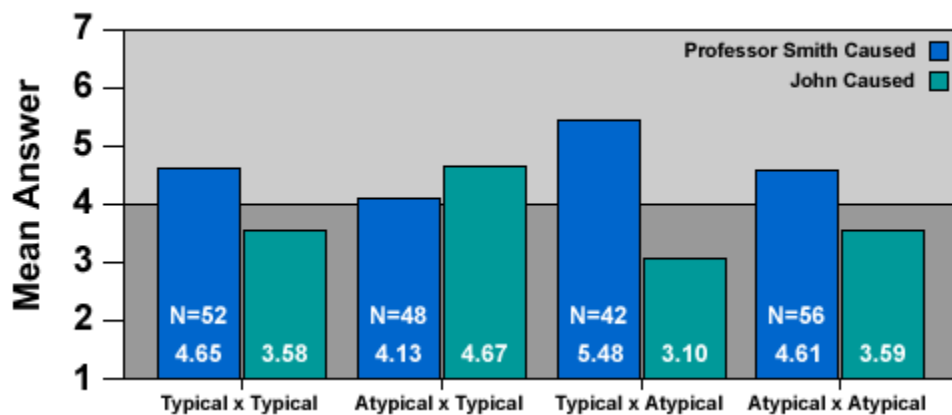


Figure 4: Study 4 results.

## 6. Problems with the Pen Case

The studies reported in the previous two sections are in line with the predictions we made in Section 3 and provide evidence against the consensus position detailed in Section 1. Further, these four studies investigated the empirical example that has driven the consensus position—Knobe and Fraser’s Pen Case. Despite this, it might nonetheless be argued that the Pen Case is a bad example. Thus, it could be objected that there are problems with the vignettes used that undermine the conclusions drawn about the impact of statistical norms on ordinary causal attributions. Given that the empirical evidence for the consensus position comes from studies using the Pen Case, however, this objection applies to that position as much as it applies to the alternative that we have proposed.

Be that as it may, the objection that the Pen Case is a bad example also applies to the positive conclusions we drew in Sections 4 and 5. As such, in this section we discuss the most serious problem for the vignettes tested. We then respond to the objection by presenting the results of two final studies. What we find is that even after rewriting the vignettes to remove the problem, our predictions are still supported.

The most serious problem with the Pen Case is that there is a non-trivial time delay between the agents taking pens and the problem occurring. Thus, the administrative assistant and Professor Smith take pens in the morning, while the problem occurs later that day. As such, a participant might reasonably infer that other people would have taken pens in the meantime and this could plausibly mitigate their causal judgments about the example. Further, with regard to the Pen Case vignettes used in Studies 3 and 4, in particular, it might be argued that this problem is exacerbated by our having been heavy-handed in describing the agents' patterns of behavior: In the version in which both agents act typically, for example, they are described as each taking a pen almost every time they pass the receptionist's desk. Thus, participants might infer that they each took multiple pens between Monday morning and the time the problem occurred.

### *6.1 Study 5: The Modified Agent-Level Pen Case*

To test the objection, we rewrote the vignettes to further flesh-out the example and to remove the problem concerning the time delay noted above. As with the previous studies, the text of the rewritten vignettes is included in the online appendices. We began by testing the impact of agent-level statistical norms. In Study 5, participants were given rewritten versions of the vignettes used in Study 3. Responses for these four vignettes were collected online from 198 native English speakers, 18 years of age or older, with at most minimal training in philosophy. The results are shown in Figure 5.



What we find is that even after rewriting the vignettes, assignment to a specific probe still affected participant responses. Specifically, participants remained more likely to say that a character caused the problem when he behaved agent-level typically than when he behaved agent-level atypically. These results provide further support for the claim that ordinary causal attributions about agents are sensitive to agent-level typicality (not atypicality); further, they undermine the claim that our results can be explained away by calling on the problem with the Pen Case noted above.

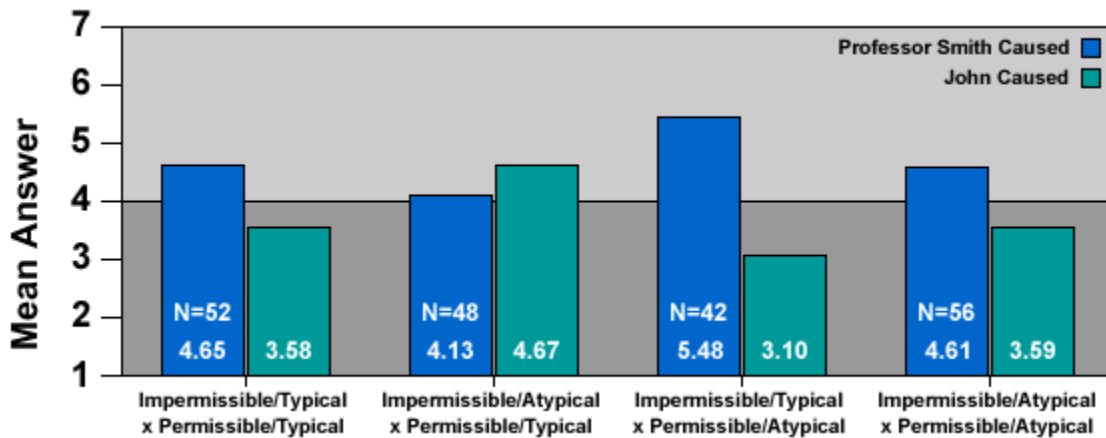


Figure 5: Study 5 results.

### 6.2 Study 6: Modified Population-Level and Agent-Level Pen Case

In our final study, we randomly gave participants one of 16 variations on the rewritten versions of Pen Case used in the previous study. Each of these vignettes included information about both population-level statistical norms and agent-level statistical norms for each of the two characters (giving 16 combinations total for these four variables). Responses were collected online from 760 native English speakers, 18 years of age or older, with at most minimal training in philosophy.

As predicted, we found that agent-level statistical norms, but not population-level statistical norms, mattered for ordinary causal attributions. Specifically, participants were more likely to say that a character caused the problem when he behaved agent-level typically than when he behaved agent-level atypically, and participants were neither more nor less likely to say that a character caused the problem when he behaved population-level typically than when he behaved population-level atypically. These results once again support each of our two predictions, while indicating that the consensus position is seriously mistaken.

## **7. Conclusion**

In this paper, we have reconsidered the recent debate concerning the role of statistical norms in ordinary causal attributions. Participants in this debate have sought to explain the seemingly surprising finding that ordinary causal attributions are not just sensitive to purely descriptive considerations, but are also sensitive to broadly moral considerations: The consensus position that has emerged assimilates broadly moral considerations about the permissibility of behaviors to a wider phenomenon that includes purely descriptive considerations about the typicality of behaviors. This position holds that people are more likely to say that an agent caused an outcome when she acts in a way that is out of the norm than when she acts in a way that is in the norm, where these norms may be either prescriptive norms or statistical norms. Hence, according to the consensus position, both prescriptive norms and statistical norms directly and independently affect ordinary causal attributions.

By contrast, we conjectured that ordinary causal attributions concerning agents are related to attributions of normative responsibility. As such, we suggested that statistical norms might affect ordinary causal attributions only indirectly by affecting judgments about normative responsibility. On the basis of this suggestion we made two predictions concerning the role of

information about statistical norms in folk causal attributions for variants on Knobe and Fraser's Pen Case: (1) Ordinary causal attributions are insensitive to population-level statistical norms; (2) ordinary causal attributions are sensitive to agent-level typicality, not atypicality. Both of these predictions diverge from those suggested by the consensus position. We then ran a series of studies to test these predictions: Our predictions, but not those given by the consensus position, were consistent with the experimental results.

## References

- Alicke, M. (1992). "Culpable Causation." *Journal of Personality and Social Psychology*, 36: 368–378.
- Driver, J. (2008a). "Attributions of Causation and Moral Responsibility." In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Volume 2: The Cognitive Science of Morality*, pp. 423–439. Cambridge: MIT Press.
- Driver, J. (2008b). "Kinds of norms and legal causation: Reply to Knobe and Fraser and Deigh." In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Volume 2: The Cognitive Science of Morality*, pp. 459–461. Cambridge: MIT Press.
- Hitchcock, C., & Knobe, J. (2009). "Cause and Norm." *Journal of Philosophy*, 106(11): 587–612.
- Knobe, J. (2006). *Folk Psychology, Folk Morality*. Dissertation.
- Knobe, J., & Fraser, B. (2008). "Causal judgments and moral judgment: Two experiments." In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Volume 2: The Cognitive Science of Morality*, pp. 441–447. Cambridge: MIT Press.
- Livengood, J., Sytsma, J., and Rose, D. (2011). "Folk Intuitions and Theories of Actual Causation: A Hitch in Hitchcock's Account." Unpublished Manuscript.
- Roxborough, C., & Cumby, J. (2009). "Folk psychological concepts: Causation." *Philosophical Psychology*, 22(2): 205–213.