

Knowledge In and Out of Contrast

Mikkel Gerken (Edinburgh) and James R. Beebe (Buffalo)

Work in Progress – Please do not cite – Comments welcome

1: INTRODUCTION. We report the results of a series of experiments that address a *contrast effect* on folk knowledge ascriptions: individuals' inclinations to agree with knowledge ascriptions to subjects in the same epistemic position differ depending upon what aspects of the subjects' epistemic positions are "in contrast." The contrast effect appears to be an important aspect of our folk epistemology, but there are several competing theoretical accounts of it. We aim to shed light the various accounts by providing novel empirical data and theoretical considerations.

In Sect. 2, we introduce the contrast effect in folk knowledge ascriptions and some interesting theoretical approaches to the existing data for it. We also provide a brief methodological prelude to our own experiments. In Sect. 3, we contribute to the accumulation of data for a contrast effect and by supplementing the existing data with two baseline conditions, as well as with some variations in the experimental manipulations.. In Sect. 4, we produce a contrast effect for *belief* ascriptions and consider the ramifications of this finding for the theoretical approaches. In Sect. 5, we present a series of experiments in which a key manipulation of the original experiment is altered and discuss how the results bear on the various theoretical approaches. In Sect. 6, we note some null-results of our experiments, and in Sect. 7, we conclude by discussing how the dataset constrains future theorizing about knowledge ascriptions.

We conclude that the dataset presents challenges for every theoretical approach to the contrast effect. However, we argue that the challenges it raises seem to be most serious for *contrastivist* accounts, according to the contrast effects straightforwardly reflect the truth-conditions of 'knows.' But we emphasize that the dataset remains limited and that both further experiments and further epistemological theorizing are called for.

2: THE CONTRAST EFFECT AND THEORETICAL APPROACHES TO EXPLAINING IT.

We introduce the contrast effect on knowledge ascriptions by outlining the existing evidence for it and sketch some interesting theoretical approaches to it.

2.1: The contrast effect. Schaffer and Knobe (2012: 689) first produced evidence for a contrast effect by three different manipulations, each of which was based upon the following vignette:

Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter forgot to wear gloves. He also forgot about the security camera. Today, Mary the detective has been called to the scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information.

Participants in the *thief contrast condition* were asked the extent to which they agreed with the following knowledge ascription:

K1: Mary now knows that Peter rather than anyone else stole the rubies. (rather than)¹

Participants in the *jewel contrast condition* were given the knowledge ascription:

K2: Mary now knows that Peter stole the rubies rather than anything else. (rather than)

Participants in the thief contrast condition were more inclined to agree than those in the jewel contrast condition. Mean ratings on a seven point Likert scale with '1' labeled as 'Disagree' and '7' labeled as 'Agree' were 4.6 for the thief contrast (K1), compared to 3.1 for the jewel contrast (K2). Schaffer and Knobe also produced the contrast using 'knowledge wh-' constructions in a pair of task probes:

K3: Mary knows who stole the rubies. (knows-wh)

K4: Mary knows what Peter stole. (knows-wh)

Participants were more inclined to agree with K3 (Mean = 4.91) than K4 (Mean = 2.62).

Finally, Schaffer and Knobe also generated the contrast by manipulating the reported conversational context for a knowledge ascription that appears in the vignette. Here the conditions were manipulated by an addition to the vignette above (henceforth 'the newspaper manipulation'). The jewel contrast condition was generated by the following addition (Schaffer and Knobe: 962):

Everyone is now asking the big question: Who stole the rubies? The news reporter is about to write a story about Mary. He is wondering if Mary now knows who stole the rubies. He writes: "Mary now knows that Peter stole the rubies."

K5: Please tell us whether you agree or disagree with the news reporter's claim, "Mary now knows that Peter stole the rubies." (newspaper manipulation)

The thief contrast condition for the newspaper manipulation was generated by the following addition to the vignette (Schaffer and Knobe: 963):

Everyone is now asking the big question: What did Peter steal? The news reporter is about to write a story about Mary. He is wondering if Mary now knows what Peter stole. He writes: "Mary now knows that Peter stole the rubies."

K6: Please tell us whether you agree or disagree with the news reporter's claim, "Mary now knows that Peter stole the rubies." (newspaper manipulation)

The mean rating was 5.24 for K5 versus 2.97 for K6. All of these differences were statistically significant.

2.2: Theoretical approaches to the contrast effect. The initial data suggests that the contrast effect is a robust feature of folk knowledge ascriptions. However, this assumption is compatible with a wide range of theoretical explanations. In this section, we briefly survey some interesting

¹ The names 'K1,' 'K2,' etc. were labels we used for our convenience and were not seen by participants. Likewise, the parenthetical characterizations are added here in order to make it easier for the reader to remember the conditions.

theoretical approaches. The approaches we consider do not constitute an exhaustive list, and several of them may be developed in different directions. Moreover, some of them may be combined in interesting ways. However, each of the approaches is independently interesting and has some traction with some of the novel data that we report below.

2.2.a: Semantic accounts. Schaffer and Knobe argue that the contrast effect provides evidence for a contrastivist theory of the term ‘knows.’ Roughly, whether a sentence of the form ‘S knows that p’ is *true* depends on certain contextual factors – specifically, the explicit or implicit contrast class. More specifically, ‘knows’ does not denote a two-place relation between a subject and a proposition. Rather, it denotes a three-place relation between a person, a proposition, and a contrast proposition (Schaffer 2005, Schaffer and Knobe 2012). Thus, all knowledge ascriptions are (tacitly or explicitly) of the form *S knows that p rather than q*. Since the contrast proposition (viz., *q*) is contextually determined, contrastivism is a brand of contextualism according to which the truth-conditions of ‘knows’ vary with contextual variances (DeRose 2009, Schaffer 2004). Such a view is a revisionist one. Traditionally it has been presupposed that whether a knowledge ascription is true is independent of such contextual factors.

A contrastivist account provides a fairly straightforward account of the contrast effect data. According to Schaffer and Knobe, the knowledge ascriptions in each thief contrast condition involves a different contrast proposition than the knowledge ascriptions in each jewel contrast condition—even when no contrast proposition explicitly appears in the knowledge ascription. Because Mary’s evidence points to Peter being the thief but does not indicate what items he may have stolen, contrastivism issues the verdict that the knowledge ascriptions are true in the thief contrast conditions (where the identity of the thief is the central issue) but not in the jewel contrast conditions (which focus on what items were stolen). This assessment seems to accord well with the data. Thus, according to contrastivism, the participants’ responses exhibit a semantic competence with the word ‘knows’ that should be reflected in the truth-conditions of the term.²

Opposition to contrastivism’s explanation of the contrast effect is motivated by further reflection on the case. Intuitively, the evidential requirements for knowing that Peter stole the rubies are higher than for knowing that Peter stole something. Yet since Mary has no evidence about what was stolen (and even lacks relevant beliefs about it), it seems that Mary does not know that Peter stole the rubies. Gerken offers the following line of reasoning (Gerken 2013):

- M1: In every case, Mary knows that Peter stole the rubies only if Mary is in a position to know that the rubies were stolen.
- M2: In the present case, Mary is not in a position to know that the rubies were stolen.
- M3: In the present case, Mary does not know that Peter stole the rubies.

Gerken takes the argument to support *strict invariantism* which is the traditional view that the *truth* of a knowledge ascription does not depend on contextual or conversational factors such as contrast class. However, DeRose (who defends a non-contrastivist brand of contextualism about ‘knows’) suggests that such an argument extends to the knowledge-wh case by highlighting the plausibility of the following conditional: “If Mary knows who stole the rubies, then Mary knows that Peter stole the rubies” (DeRose 2011: 107). A similar conditional has the ‘rather than’ knowledge ascription as the antecedent. There is thus a conflict between contrastivism’s account of

² A semantic approach need not be a brand of contextualism. For example, Aloni and Egré outline an invariantist semantics of ‘knows’ which may account for the contrast effect (Aloni and Egré 2010). According to such an account, the propositions expressed by the thief and jewel conditions differ due to differences that have nothing to do with ‘knows.’

the contrast data and epistemological theorizing about the relevant knowledge ascriptions. Anti-contrastivists may regard Gerken's and DeRose's arguments as examples of how overt epistemological theorizing may constrain the interpretation of the experimental data.

2.2.b: Psychological bias accounts. Schaffer and Knobe criticize a psychological account of the contrast effect that they label a 'shallow processing' account. According to the shallow processing account, the participants' response marks a performance mistake (Schaffer and Knobe, 2012: 700). Given Schaffer and Knobe's criticism, we will focus on psychological *bias* accounts such as Gerken's epistemic focal bias account. It differs from the shallow processing view that Schaffer and Knobe criticize by postulating a systematic *psychological bias*, rather than a mere *performance mistake*, in the participant's judgments.

The focal bias account presupposes strict invariantism about 'knows' (i.e., the view that the truth conditions of 'knows' do not vary with context.) The account regards the positive knowledge attributions as *false positives* on the basis of the arguments outlined above (Gerken 2012, 2013). According to the focal bias hypothesis, the participants in the thief contrast conditions fail to adequately process epistemically relevant facts that are not in focus. Instead, they form the intuitive judgment about the knowledge ascription in question on the basis of considering whether Mary's evidence allows her to rule out the alternatives that are in focus – i.e., the alternatives to Peter being the thief (Gerken 2012: 160ff, 2013). Hence, they do not adequately consider that Mary knows that Peter stole the rubies, her evidence must also indicate that it was rubies that were stolen.

The focal bias hypothesis can take the form of a *cognitive invisibility* account, according to which participants completely fail to register an epistemically relevant fact, or an *inadequate processing* account, according to which epistemically relevant facts are not adequately processed (Gerken 2013: 164). Gerken suggests that the focal bias only partly accounts for the contrast effect and that it should be integrated with other psychological considerations and with pragmatic accounts (Gerken 2012: 163ff, 2013: Sect. 6.2). However, he does not develop such integrations but considers the focal bias account in isolation. So, for the present purpose, we will also consider the focal bias hypothesis as a pure psychological bias account.

The focal bias account does not account for the initial contrast effect data as straightforwardly as the contrastivist account. This is because it regards participants as making a systematic mistake in their judgments about the knowledge ascriptions. In order for this claim to be vindicated, additional evidence must be produced in support of it. In Sections 4 and 5 below, we provide some of the requisite evidence.³

2.2.c: Pragmatic approaches. Schaffer and Knobe consider and dismiss a pragmatic account of the contrast effect that is based on Grice's maxims of Relevance and Quality (Grice 1989. See also Blome-Tillmann 2013). The account they consider has it that knowledge ascriptions K1 through K6 are all true although the jewel contrast "generates the *false implicature* that Mary can eliminate alternatives in which Peter stole something else" (Schaffer and Knobe 2012: 698).

Schaffer and Knobe argue against such a pragmatic account. However, they do not argue against a version of the view according to which all the relevant knowledge ascriptions are *false*.

³ Furthermore, Participants responses may be claimed to be mistaken due to inadequate conceptual resources – sometimes called "mindware gaps" (Stanovich 2009, 2011, Lerner and Tetlock 1999). So, strict invariantism may be defended by arguing that participants lack the conceptual resources for responding adequately to the task. For example, the armchair arguments contra contrastivism are fairly complex (Kelp 2011, DeRose 2011, Gerken 2013). Such arguments may require conceptual resources that participants lack. Likewise, participants are insensitive to the pragmatics-semantics distinction (Rysiew 2007). As these candidate examples of mindware gaps indicate, this approach may be combined with another psychological or pragmatic approach.

Given that this is the conclusion of the a priori arguments noted above, we take this to be the strongest version of the pragmatic approach to the finding of a contrast effect (DeRose 2011: 107, Gerken: 2013, Sect. 3). A pragmatic defense of this type of invariantism must explain why participants are inclined to agree with a false knowledge ascription. Among the resources for such an approach is the independently motivated idea of presupposition accommodation (Stalnaker 1974, Karttunen 1974, Beaver 1997). Thus, it may be argued that the participants are, given the thief contexts, reasonable in accommodating the false presupposition that Mary knows what was stolen (Gerken 2013, Sect. 6.2, Steglich-Petersen *forthcoming*). Such an approach is not prone to Schaffer and Knobe's objections although it still faces the challenges of explaining why participants readily accommodate a *false* presupposition.

The pragmatic approach noted above is compatible with a psychological account. Indeed, there is an important difference between postulating presupposition accommodation and the idea that participants are more likely to unreflectively accept a presupposition of a sentence. Moreover, independently motivated frameworks of cognitive pragmatics, such as relevance theory, seek to integrate psychological and pragmatic elements (Sperber and Wilson 1986/1995, Carston 2002). According to such accounts, communication is constrained by our cognitive capacities. So, what is conveyed by a sentence in any given conversational context is partly determined by how ordinary speakers process information. On the other hand, pragmatic facts concerning conversational significance partly determine what information is processed or how it is processed.

Cognitive pragmatics offers resources for addressing the challenges associated with pure psychological and pure pragmatic accounts. The program of cognitive pragmatics is sometimes criticized for failing to yield clear predictions. But given that it aims to combine psychological and pragmatic approaches, it is worth considering in the present context.

The above list of approaches is not exhaustive but it marks some important distinctions that are important for theorizing about folk epistemology. For example, the distinction between contrastivism (Schaffer's brand of contextualism) according to which the truth conditions of knowledge ascriptions vary with conversationally implicit or explicit contrasts and the traditional view, strict invariantism, according to which whether a knowledge ascription is true is independent of such conversational factors. Among the defenses of strict invariantism, we find pragmatic accounts as well as psychological accounts. Among the latter we find accounts that postulate *performance error* (e.g., the shallow processing account criticized by Schaffer and Knobe) and those that postulate a *systematic bias* (e.g., Gerken's focal bias account.) Psychological bias accounts, in turn, may appeal to *cognitive invisibility*, *inadequate processing* or to *mindware gaps*. Finally, there are various combinations of these approaches that have yet to be developed. Such developments should be constrained by empirical data pertaining to knowledge ascriptions.

We will contribute to the debates by providing a range of new data and discussing their bearing on the various approaches. We will highlight cases in which the data provide a challenge for a given approach to the original contrast effect. However, we will also engage in constructive work by pointing out when the data support a given approach or suggests a revision of development of it. Although all of the approaches sketched above should be considered, we will focus primarily on the how the data bears on Schaffer's contrastivism and Gerken's focal bias version of the psychological bias approach. This is in part due to considerations of space and in part because these accounts are comparatively well developed. Moreover, by considering these accounts, we consider the issue as an instance of the broader invariantist-contextualist dispute, and it is reasonable to address this major dispute before turning, for example, to in-house disputes among invariantists.

2.3: Methodological considerations. All of the accounts sketched above are consistent with the presence of contrast effects in a wide variety of knowledge ascriptions. In consequence, further epistemological reflection and further empirical data are called for. We aim to provide novel empirical data that may serve as input for further reflection. We do not assume that the issue can be settled by empirical investigation alone. Both theorizing from the relative isolation of the armchair and reflection on the sort of empirical data that we will consider are important for the theory of knowledge as well as for an account of our folk epistemology.

Hence, we regard our experimental findings as contributing to an overlap between philosophy and cognitive psychology. In contrast, we eschew experimental philosophy's *negative program*, according to which empirical data is invoked to cast doubt on or replace traditional philosophical methodology (see, e.g., Nagel 2012b, 2013). Just as empirical investigations may shed light on traditionally philosophical subject matters, traditional philosophy may shed light on empirical matters. For example, traditional epistemology may (on a good day) provide the "gold standard" response to a task which is required for interpreting participant responses to it. Furthermore, the experimental data may underdetermine the proper philosophical account of the truth-conditions of 'knows.' Indeed, we will not claim that the data that we present provide decisive evidence for or against any one of the approaches sketched above. Theoretical virtues that go beyond consistency with experimental data must figure in the choice of epistemological theory.

Despite these qualifications, the experimental evidence that we will present remains important data for theorizing about knowledge ascriptions. In particular, some of the data we are about to present will provide explanatory challenges for some theoretical approaches and thus form a part of an abductively based theory choice.

3. STUDY 1: REPLICATIONS, SUPPLEMENTATIONS, AND BASELINES. We began our empirical investigation of contrast effects in folk knowledge attributions by attempting to replicate and supplement the results obtained by Schaffer and Knobe in the K1 through K6 conditions by a minor change in the stimuli and by obtaining some baseline measures.

3.1. Methodology. In a between-subjects design, participants in Study 1 read either Schaffer and Knobe's original vignette or one of the supplemented vignettes used in the newspaper manipulation described above. Each participant was then asked to indicate the extent to which they agreed or disagreed with the relevant knowledge ascription. While Schaffer and Knobe used the temporal qualifier 'now' in some of their prompts (K1, K2, K5, and K6) but not in others (K3 and K4), we omitted the term from each of the target knowledge ascriptions for the sake of uniformity.⁴

We also attempted to generate an additional contrast by constructing a pair of knowledge ascriptions that only mentioned either the jewel or the thief aspect of the vignette:

K7: Mary knows that Peter stole something from the jewelry store.

K8: Mary knows that rubies were stolen from the jewelry store.

Schaffer and Knobe's initial study did not involve a base-line condition that would provide evidence as to whether the contrast effect should be seen as increasing or decreasing a simple, non-contrasted knowledge ascription. So, in Study 1 we also provided two differently worded base-line prompt questions to be used with the original vignette:

⁴ The omission also helps to prevent the prompt question from being interpreted as asking whether Mary knew more after her investigation than before. Furthermore, it also sidesteps any confusion that may arise from the fact that 'now' figures in each of the newspaper manipulation vignettes but not in their accompanying prompt questions. Despite this change in the stimuli, we will continue to use 'K1-K6' as labels for the conditions in question.

K9: Mary knows that Peter stole the rubies. (baseline)

K10: Mary knows that Peter stole rubies from the jewelry store. (baseline)

400 participants (average age = 33, 43% female, 77% Anglo-American, 83% with at least some college education) were recruited through Amazon's Mechanical Turk (www.mturk.com) and were assigned to one of the experimental conditions K1 through K10.⁵ Each participant was asked two comprehension questions, and those who failed to answer both questions correctly were excluded from the analysis. Participants were asked to indicate the extent to which they agreed or disagreed with the target knowledge ascription on a seven-point scale marked with the labels 'Completely Disagree,' 'Mostly Disagree,' 'Slightly Disagree,' 'Neither Agree nor Disagree,' 'Slightly Agree,' 'Mostly Agree,' and 'Completely Agree.' Before providing demographic information about themselves, participants were asked if they had any comments they would like to share with the researchers. After indicating their gender, age, ethnicity, and highest educational level attained, participants completed the Cognitive Reflection Task (Frederick 2005). Toplak et al characterize the CRT as designed to test individuals' ability "to override a prepotent response alternative that is incorrect and to engage in further reflection that leads to the correct response" (Toplak, West and Stanovich 2011).

3.2. Results. The data depicted in Figure 1 reveal that we succeeded in generating a contrast effect similar to that found by Schaffer and Knobe in the three contexts represented by conditions K1 through K6. A statistically significant difference between mean knowledge attributions was found within the 'rather than' pair (K1 and K2, small effect size), the 'knowledge wh-' pair (K3 and K4, large effect size), and the newspaper manipulation pair (K5 and K6, medium effect size).⁶ When we reran some of our conditions with 'Master Workers'⁷ from Mechanical Turk in an effort to obtain further confirmation of the reliability of our most important results, a larger difference between the mean knowledge attributions in the K1 (5.93) and K2 (4.12) condition was found.⁸ Mean knowledge ratings fell significantly above the neutral midpoint in four of the six conditions.⁹ Notably, all of the mean knowledge ratings we obtained appear to be significantly higher than their counterparts in Schaffer and Knobe's study.

⁵ Workers were paid \$.25 for their participation and were not allowed to participate in more than one condition.

⁶ K1 vs. K2: $U = 599$, $z = -2.037$, $p < .05$, $r = -.23$. K3 vs. K4: $U = 261.5$, $z = -5.16$, $p < .001$, $r = -.58$. K5 vs. K6: $U = 438$, $z = -2.882$, $p < .01$, $r = -.33$.

⁷ According to the Mechanical Turk help pages (<https://requester.mturk.com/help/faq>), "Masters are elite groups of Workers who have demonstrated accuracy on specific types of HITs [i.e., Human Intelligence Tasks] on the Mechanical Turk marketplace. Workers achieve a Masters distinction by consistently completing HITs of a certain type with a high degree of accuracy across a variety of Requesters. Masters must continue to pass our statistical monitoring to remain Mechanical Turk Masters." Master workers were each paid \$1 for their participation.

⁸ $U = 225.5$, $z = -3.136$, $p < 0.01$, $r = -0.41$, medium effect size.

⁹ K1: $t(39) = 7.154$, $p < .001$, $r = .75$ (large effect size). K2: $t(39) = 3.615$, $p < .01$, $r = .50$ (large effect size). K3: $t(38) = 13.346$, $p < .001$, $r = .91$ (large effect size). K4: $t(38) = -0.219$, $p > .05$, $r = .04$. K5: $t(36) = 6.602$, $p < .001$, $r = .74$ (large effect size). K6: $t(37) = 1.42$, $p > .05$, $r = .23$.




Figure 1. Mean knowledge attributions in the ‘rather than’ (K1 = 5.95, K2 = 5.14), knowledge-wh (K3 = 6.31, K4 = 3.92), and newspaper manipulation (K5 = 5.76, K6 = 4.49) pairs of conditions in Study 1. An ‘*,’ ‘**,’ or ‘***’ indicates that the mean differs significantly from the neutral midpoint at either the .05, the .01, or the .001 level. An ‘*,’ ‘**,’ or ‘***’ with a bracket indicates a statistically significant difference between pairs of conditions at either the .05, the .01, or the .001 level. Error bars represent 95% confidence intervals in all figures.

The results from K7 through K10 are represented in Figure 2. The means in three of the four conditions fell significantly above the midpoint.¹⁰ The manipulations in which only ‘Peter’ (K7) or only ‘rubies’ (K8) figured in the target knowledge ascription generated a statistically significant contrast effect with a medium effect size.¹¹

There was no statistically significant difference between the two baseline conditions (K9 and K10). Moreover, there was no statistically significant difference between either of the baseline conditions and any of the thief contrast conditions (i.e., K1, K3, K5, or K7). In other words, participants were no less likely to think that that Mary knows that Peter stole the *rubies* from the jewelry store (K10) as they were to think Mary knows that that Peter stole *something* from the jewelry store (K7). However, there was a statistically significant difference between each of the baseline conditions and each of jewel contrast conditions (i.e., K2, K4, K6, and K8), with medium to large effect sizes in every case.¹² That is, participants were more likely to think that Mary knows that Peter stole the rubies (K9) than they were to think that Mary knows what Peter stole (K4).

There were no effects of gender or level of education on knowledge attributions in Study 1. However, there was a medium-sized positive correlation between the number of items a participant answered correctly on the Cognitive Reflection Task and that participant’s inclination to attribute

¹⁰ K7: $t(36) = 6.356, p < .001, r = .73$ (large effect size). K8: $t(39) = -0.134, p > .05, r = .02$. K9: $t(36) = 7.613, p < .001, r = .79$ (large effect size). K10: $t(37) = 7.561, p < .001, r = .78$ (large effect size).

¹¹ $U = 448, z = -2.763, p < 0.01, r = -.32$.

¹² K9 vs. K2: $U = 520, z = -2.010, p < .05, r = -.23$ (small effect size). K9 vs. K4: $U = 285, z = -4.426, p < 0.001, r = -.51$ (large effect size). K9 vs. K6: $U = 362, z = -3.452, p < .01, r = -.40$ (medium effect size). K9 vs. K8: $U = 363.5, z = -3.452, p < .01, r = -.40$ (medium effect size). K10 vs. K2: $U = 552, z = -2.195, p < .05, r = -.25$ (small effect size). K10 vs. K4: $U = 299, z = -4.659, p < .001, r = -.53$ (large effect size). K10 vs. K6: $U = 379.5, z = -3.689, p < .001, r = -.42$ (medium effect size). K10 vs. K8: $U = 397, z = -3.512, p < .001, r = -.40$ (medium effect size).

knowledge in the thief contrast conditions K1, K3, and K5¹³. There was no significant correlation between CRT score and knowledge attributions in the jewel contrast conditions K2, K4, and K6. In other words, high CRT participants were more inclined than low CRT participants to think that Mary had knowledge when the focus of the prompt was on who committed the crime; but high CRT participants were not any more or less inclined than low CRT participants to attribute knowledge to Mary when the focus of the prompt was on what was stolen.




Figure 2. Mean knowledge attributions in the ‘Peter only’ (K7 = 5.62), ‘rubies only’ (K8 = 3.95), and baseline (K9 = 5.97, K10 = 6.03) conditions in Study 1.

Above we noted that all of the mean knowledge ratings that appear in Figure 1 were significantly higher than their counterparts in Schaffer and Knobe’s original study. Relatedly, the lowest mean knowledge attribution we obtained in Study 1 was 3.92, even though both Schaffer and Knobe’s contrastivism and Gerken’s focal bias account contend that the correct verdict in these conditions is that knowledge should be denied to Mary.

3.3. Discussion. Our findings support the robustness of the contrast effect in ordinary assessments of knowledge ascriptions and provide some interesting perspectives on the nature of this contrast effect.

Consider that knowledge ascriptions in the baseline conditions were significantly higher than in the jewel contrast conditions but *not* significantly higher than in the thief contrast conditions. This raises explanatory challenges to contrastivist accounts. Given that we found no difference between mean participant agreement with K1 (‘Mary knows that Peter rather than anyone else stole the rubies’) and K9 (‘Mary knows that Peter stole the rubies’) using the same vignette, the original data may indicate that a *general* inclination to agree with a knowledge ascription is diminished only given focus on a certain epistemic alternative. So, the contrastivist owes an account of why there is no difference between the thief contrast case and the baseline case. Contrastivists might try to claim that the responses in the baseline conditions are at the “ceiling” of agreement with any knowledge ascriptions, and so the inclination to agree could not be stronger in the thief contrast. A competing contrastivist account has it that the baseline condition involves an

¹³ $r = .31, p < .01$.

unarticulated contrast that is equivalent to the contrast in the thief condition.¹⁴ But in order for the latter response to not be *ad hoc*, it should be accompanied by an account of what generates the implicit contrast in the baseline cases.

Consider also the fact that we found no significant difference between mean participant responses to K7 ('Mary knows that Peter stole something from the jewelry store') and K9 ('Mary knows that Peter stole the rubies.') Even contrastivists should admit that the evidential requirements for being in a position to know that Peter stole the rubies are higher than for being in a position to know that Peter stole something. If contrastivists wish to maintain that participant responses to contrast effect cases exhibit semantic competence with the word 'knows' that should be reflected in its truth conditions, they must find a way to accommodate these unfavorable data.

According to the most promising version of strict invariantism, the 'rubies' component of the complement clause of the knowledge ascription is epistemically relevant. So, invariantists postulating a psychological bias may gain modest support from the fact that there were significant differences between participant responses in the baseline conditions and the jewel contrast conditions but not between the baseline and thief contrast conditions. For example, according to the focal bias account, the epistemically relevant 'rubies' component is not adequately processed unless it is in focus. On this account, participants who fail to adequately consider the epistemic relevance of the jewel component will be generally inclined to agree with the knowledge ascription.

Invariantist's pragmatic accounts are also compatible with this finding although they too face an explanatory challenge. The general challenge to pure pragmatic accounts is to explain why the false presupposition that Mary knows what was stolen is accommodated. Pure pragmatic accounts may appeal to syntactic features of the knowledge ascriptions – i.e., to the 'rather than' and 'wh-' phrases (Steglich-Petersen *forthcoming*). But such an account must then explain why the presupposition is also accommodated in the baseline knowledge ascriptions despite the apparent absence of such syntactic features. Consider, in addition, the difference between the two baseline knowledge ascriptions. Recall that one, K9, read "Mary knows that Peter stole the rubies" whereas the other, K10, read "Mary knows that Peter stole rubies from the jewelry store." The former might be thought to lend more easily to an accommodation of the presupposition that Mary knows that rubies were stolen insofar as the determinate phrase "the rubies" may be taken to indicate that the object of theft is common knowledge. However, there was no statistically difference between these two baseline conditions. This is a further data point that a pure pragmatic account must explain.

Let us turn to the lack of significant disagreement with the knowledge ascriptions in any of the conditions – including our jewel contrast that only asks about Mary's knowledge about the rubies. The "high floor" might be taken to suggest a problem with the stimuli that has been noted by, among others, DeRose (DeRose 2011: 99ff). DeRose argues that experimental participants must always presuppose some aspects of Mary's background information. For example, she must be presupposed to know where the theft took place. So, according to DeRose, it may, given the narrative, be quite natural to suppose that Mary is informed about what was stolen. We may add that this is particularly so in the vignette in which participants are told that "Everyone is now asking the big question: Who stole the rubies?" (Schaffer and Knobe: 692). It would be odd to suppose that everybody *but the detective* knew what was stolen. So, particularities of the vignette might partly explain that participants do not tend to disagree with the knowledge ascription in any conditions. This candidate explanation of lack of disagreement with the knowledge ascriptions presupposes an aspect of pragmatic accounts – namely, presupposition accommodation. It does not, however, provide direct evidence for such pragmatic views insofar as they do not predict a high "floor" any more than the competing theories. For example, on a focal bias account, it may be that it is not

¹⁴ Indeed, given a contrastivist account, it is misguided to speak of "baseline" conditions since every knowledge ascription involves a tacit contrast.

enough that there is an epistemically relevant alternative that is salient in the knowledge ascription—it may also need to be salient for the participants that it is an epistemically relevant alternative. If so, the fact that the epistemic relevance of the object of theft is not salient in Schaffer and Knobe’s *vignette* may contribute to an explanation of the high floor.¹⁵

Furthermore, Schaffer has, in a replication with a new version of the jewel thief vignette in which Mary’s evidence is more explicit found a lower floor. Given this manipulation, participants in the jewel contrast condition tended to disagree in a manipulation mirroring the newspaper addition (Thief: $M = 5.05$. Jewel: $M = 2.13$). A within-subject design with the modified vignette resulted in a comparable pattern for a range of knowledge ascriptions. (These results are cited from Schaffer and Szabo, *forthcoming*, which, in turn cites, Schaffer *forthcoming*). These findings may be taken to support the interpretation according to which the high floor of the original data is explained by insufficiently explicit stimuli.

In sum, the statistically significant differences between all of the thief and jewel conditions in the original data and the present replication provide empirical reason to assume that there is a fairly robust contrast effect in folk assessment of knowledge ascriptions. Our supplementations provide some explanatory challenges for some of the theoretical approaches to the data.

4. STUDY 2: A CONTRAST EFFECT FOR BELIEF ASCRIPTIONS. In order to shed light on whether the contrast effect is a distinctive feature of knowledge ascriptions or whether it is a more general feature of mental state ascriptions, we reran most of the conditions from Study 1 replacing ‘knows’ with ‘believes.’

4.1. Methodology. It was unproblematic to substitute ‘believes’ for ‘knows’ in K1, K2, K5, K6, K9, and K10 to yield the following prompts:

- B1: Mary believes that Peter rather than anyone else stole the rubies. (rather than)
- B2: Mary believes that Peter stole the rubies rather than anything else. (rather than)
- B5: Mary believes that Peter stole the rubies. (newspaper manipulation)
- B6: Mary believes that Peter stole the rubies. (newspaper manipulation)
- B9: Mary believes that Peter stole the rubies. (baseline)
- B10: Mary believes that Peter stole rubies from the jewelry store. (baseline)

However, the same substitution in K3 and K4 would have resulted in ungrammatical belief ascriptions, such as “Mary believes who stole the rubies.” So, instead we constructed the following pair of ‘belief-about-wh’ prompts:

- B3: After considering the evidence, Mary has a belief about who stole the rubies. (belief-about-wh)
- B4: After considering the evidence, Mary has a belief about what Peter stole. (belief-about-wh)

In addition, data for the following two conditions were collected.

- B7: After considering the evidence, Mary has a belief about who stole what.(who stole what)

¹⁵ Thanks to...

B8: After considering the evidence, Mary has a belief about what was stolen by whom.
(what was stolen by whom)

Apart from these variations in the target belief ascriptions, the methodology did not differ from the methodology described in Sect. 3.1. The syntactic and semantic differences between the ‘belief-about-wh’ and the ‘knows-wh’ ascriptions will be considered below.

In a between-subjects design, 520 participants (average age = 34, 54% female, 75% Anglo-American, 85% with at least some college education) were recruited through Amazon’s Mechanical Turk and were assigned to one of the experimental conditions B1 through B10.

4.2. Results. The results of Study 2 reveal a contrast effect for ‘believes’ that is analogous to the one Schaffer and Knobe obtained for ‘knows.’ There was a statistically significant difference between the mean belief attributions in each of the three pairs of conditions represented in Figure 3.¹⁶ When we reran the B1 and B2 conditions with Master Workers from Mechanical Turk, a somewhat larger difference was found between the mean belief attribution ratings (B1 = 6.33, B2 = 5.30).




Figure 3. Mean belief attributions in the ‘rather than’ (B1 = 6.42, B2 = 5.91), ‘belief-about-wh’ (B3 = 6.61, B4 = 3.96), and newspaper manipulation (B5 = 6.24, B6 = 4.62) pairs of conditions in Study 2.

We even obtained a medium-sized contrast effect for ‘believes’ in the ‘who stole what’ (B7) and ‘what was stolen by whom’ (B8) conditions (cf. Figure 4).¹⁷ As in Study 1, there was no statistically

¹⁶ B1 vs. B2: $U = 5256$, $z = -3.456$, $p < .01$, $r = -.23$ (small effect size). B3 vs. B4: $U = 116$, $z = -4.750$, $p < .001$, $r = -.64$ (large effect size). B5 vs. B6: $U = 438.5$, $z = -2.489$, $p < .05$, $r = -.29$ (small to medium effect size). Mean belief attributions also fell significantly above the midpoint in four of the six conditions: B1: $t(116) = 18.229$, $p < .001$, $r = .86$ (large effect size). B2: $t(115) = 12.165$, $p < .001$, $r = .75$ (large effect size). B3: $t(27) = 11.526$, $p < .001$, $r = .91$ (large effect size). B4: $t(26) = -.083$, $p > .05$, $r = .02$. B5: $t(37) = 11.122$, $p < .001$, $r = .88$ (large effect size). B6: $t(33) = 1.485$, $p > .05$, $r = .25$.

¹⁷ B7 vs. B8: $U = 296.5$, $z = -2.365$, $p < .05$, $r = -.31$.

significant difference between the means in the two baseline conditions in Study 2.¹⁸ There were no effects of gender or level of education on belief ascriptions in Study 2, and there was no significant correlation between participants' CRT score and their inclination to ascribe belief.




Figure 4. Mean belief attributions in the ‘who stole what’ (B7 = 6.45), ‘what was stolen by whom’ (B8 = 5.57), and baseline (B9 = 6.05, B10 = 5.98) conditions in Study 2.

4.3. Discussion. The data set provides evidence for a contrast effect for belief ascriptions. We emphasize that the ‘belief-about-wh’ conditions do *not* amount to analogs of the ‘knowledge-wh’ conditions. This mode of ascription differs both syntactically and semantically from a simple ‘-wh’ ascription. These differences are substantive and may well impact participant responses. For example, one can have a belief about who stole the rubies without having any belief about Peter. E.g., a belief that a male stole the rubies will satisfy the ‘belief-about-wh’ ascription. So, we do *not* conclude that the ‘belief-about-wh’ finding provides evidence for a contrast effect for ‘belief-wh’ ascriptions. Our claim is rather that the finding suggests independent but related and important contrast effects for belief. Along with the data that duplicates the ‘rather than’ and ‘newspaper’ contrast for ‘knows,’ the ‘belief-about-wh’ contrast effect provides converging but defeasible evidence for a pervasive contrast effect for belief ascriptions. The question, then, is what such an effect shows about the original contrast effect for knowledge ascriptions.

It has been suggested that a contrast effect for belief would raise a challenge for epistemic contrastivism and provide evidence for the focal bias account given some auxiliary premises (Gerken 2013: fn. 18). More generally, the finding could be taken to provide evidence for a psychological account according to which knowledge ascriptions share the biases that are *generally* found in mental state ascription (see Nagel 2010, 2012a although she does not address contrast effects). To assess these suggestions, the positive and negative arguments should be spelled out.

Let us first consider the negative argument against contrastivism. In general, if some effect that is first found for some mental verb X is also found for a broader range of mental verbs, this is

¹⁸ All four means differed significantly from the midpoint. B7: $t(28) = 10.158, p < .001, r = .89$ (large effect size). B8: $t(29) = 4.230, p < .001, r = .62$ (large effect size). B9: $t(38) = 7.981, p < .001, r = .79$ (large effect size). B10: $t(39) = 7.339, p < .001, r = .76$ (large effect size).

generally taken to suggest that the effect is not particular to X. For example, the finding of a Knobe effect for belief and knowledge is taken to suggest that the Knobe effect does not reflect particular properties of intention (Beebe and Buckwalter 2010, Beebe 2013). Likewise, finding a stakes effect for belief has led to the suggestion that such an effect does not tell us much about knowledge (Buckwalter and Schaffer *forthcoming*). A more specific challenge may be set forth as a dilemma: the epistemic contrastivist either rejects or adopts doxastic contrastivism – i.e., contrastivism about ‘believes.’ Let us consider the two horns of the dilemma in turn.

First horn: The trouble with rejecting doxastic contrastivism is that the epistemic contrastivist then owes an account of why the experimental data is evidence for contrastivism in the case of ‘knows’ but not in the case of ‘believes.’ The experimental set-up is the same in the two cases. So, it would seem *ad hoc* to postulate a pragmatic or psychological explanation of the belief case and retain a semantic explanation of the knowledge case.

Second horn: Part of the trouble with accepting doxastic contrastivism concerns its motivation. Schaffer has done much to motivate contrastivism about ‘knows’ on epistemic grounds (Schaffer 2005, Knobe and Schaffer 2012). But much of this motivation appears to be distinctively epistemic. For example, it appeals to a contextualist version of a widely accepted relevant alternatives program in epistemology (Dretske 1970, 1971). So, in the case of knowledge ascriptions, there is an independently motivated idea of what the contrast proposition is – i.e., a relevant alternative that the subject must be able to rule out. But in the case of belief, it is typically not clear what the contrast proposition is supposed to be. Hence, it is unclear that the motivation of epistemic contrastivism is applicable to belief. Epistemic contrastivism is interesting because it may be given independent motivation from epistemological theory. In the absence of a similar motivation, contrastivism for belief appears to be troublesome.

Schaffer and Szabo (*forthcoming*) address the issue within a general discussion of whether their comparativist semantics for ‘knows’ generalizes to other attitude verbs.¹⁹ They argue that it generalizes to *specific epistemics* such as ‘see’ and ‘remember’ as well as to *emotive factives* such as ‘regret’ and ‘care.’ However, with regards to ‘belief’ and other attitude verbs, they “leave the prospect of generalizing comparativism to other attitude verbs unsettled” (Schaffer and Szabo, *forthcoming*).

The present findings does not settle the issue pertaining to ‘belief’ since there are, as Schaffer and Szabo note, a range of relevant behaviors of the verb to be considered. Nevertheless, our findings suggest that the contrast effect generalizes at least to ‘belief’ despite other asymmetries between ‘believes’ and ‘knows.’ So, the contrastivist faces the challenge of either motivating doxastic contrastivism or explaining the contrast effect as merely pragmatic quasi-association (as Schaffer and Szabo mentions (following Beaver and Clark (2008))). Whether the finding is damaging for contrastivism about ‘knows’ depends on the plausibility of a developed contrastivist view. So, we only conclude that the data raises a challenge for contrastivism that should be addressed.

Let us consider whether a contrast effect for belief may be taken to provide evidence for a psychological account. As mentioned, the broad rationale for psychological approaches to knowledge ascriptions is that they should be expected to exhibit biases that apply to other mental states as well (Spicer 2005, Nagel 2010, 2012, Gerken 2012). So, evidence of contrast effects for ‘believes’ may be taken as evidence for this broad rationale for a psychological approach to the knowledge ascription data. Moreover, the finding of an *word order* effect between the ‘who-what’ and ‘what-who’ conditions (B7 and B8) is important because order effects may be explained by considerations pertaining to the agents’ processing patterns and limitations.

¹⁹ Comparativism is a genus under which contrastivism is a species. We will focus on Schaffarian contrastivism.

These considerations motivate the pursuit of a psychological account. But they are too broad to provide evidence for any particular version of it. Since the focal bias account is a specific psychological approach that has been suggested as an account of the contrast effect, we will briefly consider how the present data bear on it. The focal bias account does not by itself predict a contrast effect for belief. A couple of auxiliary assumptions are required. One such auxiliary is the assumption that participants take the subject of the knowledge ascription to be epistemically rational. In the present case, this involves the more specific presumption that Mary is responsive to her evidence. Yet more specifically, that she forms beliefs in only what her total set of evidence suggests. The second, associated, assumption is that folk ascription of belief is constrained by principles similar to the ones that characterize epistemic focal bias for knowledge ascriptions. Roughly, the analog of Gerken’s *Principle of Contextual Salience* claims that for an agent, q is normally a salient alternative to S ’s belief that p just in case the agent processes q as such (Gerken 2012: 155, 2013: 50). Likewise, a rough belief-analog of Gerken’s *Principle of Epistemic Satisficing* is that agents normally form epistemic judgments based on processing a limited part of the available evidence (2012: 155, 2013: 51). Given analogs of these principles, a focal account of the contrast effect for belief ascriptions ensues. Since the auxiliary assumptions seem fairly reasonable, the present finding is congenial to the focal bias account.

The foregoing argument is not intended to be conclusive. As emphasized above, arguments from empirical data should be thought of as aspects of an overall abductive argument rather than as self-standing conclusive arguments. The considerations based on the present data appear to do just that. They raise a challenge for contrastivist account and provide limited but notable evidence for psychological approaches such as the focal bias account.²⁰

5. STUDY 3: CONTRAST EFFECTS FOR ALTERNATIVE OBJECTS OF THEFT. In a third study, we replaced the original object of theft in Schaffer and Knobe’s vignettes (i.e., rubies) with alternative objects of theft.

5.1. Methodology. The methodology was identical with the methodology described in Sect. 3.1 above. The experimental manipulations differed only in that ‘rubies’ in some of the prompt questions from Study 1 was replaced with either ‘sapphires’ or ‘the 4.76 carat diamond set on an amethyst encrusted platinum necklace from the 17th century’ (henceforth abbreviated ‘diamond’). The prompt questions from the ‘rather than’ condition (K1 and K2), the first ‘knowledge-wh’ condition (K3), the two newspaper manipulation conditions (K5 and K6), and the two baseline conditions (K9 and K10) from Study 1 were transformed into the following two sets of prompts for Study 3:

- S1: Mary knows that Peter rather than anyone else stole the sapphires. (rather than)
- S2: Mary knows that Peter stole the sapphires rather than anything else. (rather than)
- S3: Mary knows who stole the sapphires. (knowledge-wh)
- S4: Mary knows that Peter stole the sapphires. (newspaper)
- S5: Mary knows that Peter stole the sapphires. (newspaper)
- S6: Mary knows that Peter stole the sapphires. (baseline)

²⁰ We focus on the ramifications of the finding for contrastivism and the focal bias version of the psychological bias account because the proponents of these approaches have explicitly addressed the issue (Gerken 2012, fn. 18, Schaffer and Szabo *forthcoming*). However, it should be noted that the finding is consistent with both pragmatic accounts and other psychological approaches. But rather than predicting what theorists who have not yet spoken to the issue would and could say, we set forth the present data set as an invitation to address it.

S7: Mary knows that Peter stole sapphires from the store. (baseline)

D1: Mary knows that Peter rather than anyone else stole the 4.76 carat diamond set on an amethyst encrusted platinum necklace from the 17th century. (rather than)

D2: Mary knows that Peter stole the 4.76 carat diamond set on an amethyst encrusted platinum necklace from the 17th century rather than anything else. (rather than)

D3: Mary knows who stole the 4.76 carat diamond set on an amethyst encrusted platinum necklace from the 17th century. (knowledge-wh)

D4: Mary knows that Peter stole the 4.76 carat diamond set on an amethyst encrusted platinum necklace from the 17th century. (newspaper)

D5: Mary knows that Peter stole the 4.76 carat diamond set on an amethyst encrusted platinum necklace from the 17th century. (newspaper)

D6: Mary knows that Peter stole the 4.76 carat diamond set on an amethyst encrusted platinum necklace from the 17th century. (baseline)

D7: Mary knows that Peter stole the 4.76 carat diamond set on an amethyst encrusted platinum necklace from the 17th century from the store. (baseline)

Because the second ‘knowledge-wh’ prompt (K4) from Study 1 did not contain ‘rubies,’ there was no substitution to be made. We also eliminated the following sentence from the original newspaper manipulation vignettes “He [i.e., the reporter] writes: ‘Mary now knows that Peter stole the rubies’” Importantly, the occurrences of ‘rubies’ in the vignettes that accompanied each of these prompts was left unchanged. In other words, the vignettes in Study 3 were about rubies, but the prompts that involved the target knowledge ascriptions were about sapphires or 17th century necklaces.

In a between-subjects design, 760 participants (average age = 31, 48% female, 56% Anglo-American, 84% with at least some college education) were recruited through Amazon’s Mechanical Turk and were assigned to exactly one of the experimental conditions represented by S1 through S7 and D1 through D7.

5.2. Results. The main results of Study 3 are represented in Figures 5 and 6. Among both the sapphire and the diamond prompts, there was a significant difference between the means in the contrasting pairs of ‘rather than’ conditions (S1 vs. S2 and D1 vs. D2) and between the means in the newspaper manipulation conditions (S4 vs. S5 and D4 vs. D5).²¹ When the S1 and S2 were rerun with Master Workers from Mechanical Turk, the difference between the means obtained was even larger.²² Using K4 (‘Mary knows what Peter stole’) as the knowledge-wh foil to S3 and D3, no significant differences were observed between K4 and S3 or between K4 and D3.

There was a small effect for gender in the sapphires cases, with females ($M = 3.54$) being less inclined than males ($M = 4.52$) to attribute knowledge, but there was no effect for gender in the diamond cases.²³ There was an effect for level of education in both the sapphires and the diamond conditions, with education level being negatively correlated with an inclination to attribute knowledge.²⁴ In other words, the higher the level of education attained, the less likely participants were to attribute knowledge in these cases. There was also a negative correlation between

²¹ S1 vs. S2: $U = 3645.5$, $z = -2.793$, $p < .01$, $r = -.20$ (small effect size). S4 vs. S5: $U = 226$, $z = -3.157$, $p < .01$, $r = -.41$ (medium effect size). D1 vs. D2: $U = 3385$, $z = -3.241$, $p < .01$, $r = -.23$ (small effect size). D4 vs. D5: $U = 261.5$, $z = -2.204$, $p > .05$, $r = -.29$ (small to medium effect size).

²² $U = 196.5$, $z = -3.767$, $p < .001$, $r = -.49$ (medium to large effect size).

²³ Sapphires: $U = 12,534$, $z = -3.828$, $p < .001$, $r = -.20$. Diamond: $U = 15,034$, $z = -1.391$, $p > .05$, $r = -.07$.

²⁴ Sapphires: $r = -.11$ (small effect size), $p < .05$. Diamond: $r = -.13$ (small effect size), $p < .05$

participants' CRT scores and their inclination to attribute knowledge in the diamond cases.²⁵ However, there was no significant correlation between CRT scores and knowledge attributions in the sapphires cases.





Figure 5. Mean knowledge attributions in the ‘rather than’ (S1 = 4.42, S2 = 3.46), ‘knowledge-wh’ (S3 = 4.38), newspaper manipulation (S4 = 5.18, S5 = 2.87), and baseline (S6 = 4.41, S7 = 4.14) pairs of sapphire conditions in Study 3.²⁶



²⁵ $r = -.18$ (small effect size), $p < .05$.

²⁶ The mean knowledge attributions differed significantly from the midpoint in three out of the seven conditions. S1: $t(95) = 1.615$, $p > .05$, $r = .16$. S2: $t(97) = -2.204$, $p < .05$, $r = .22$ (small effect size). S3: $t(36) = .898$, $p > .05$, $r = .15$. S4: $t(27) = 2.819$, $p < .01$, $r = .48$ (medium effect size). S5: $t(29) = -2.429$, $p < .05$, $r = .41$ (medium effect size). S6: $t(36) = .945$, $p > .05$, $r = .16$. S7: $t(35) = .340$, $p > .05$, $r = .06$.

Figure 6. Mean knowledge attributions in the ‘rather than’ (D1 = 4.20, D2 = 3.09), ‘knowledge-wh’ (D3 = 3.95), newspaper manipulation (D4 = 3.59, D5 = 2.45), and baseline (D6 = 2.71, D7 = 2.70) pairs of diamond conditions in Study 3.²⁷

Table 1 and Figure 7 combine the most important data from Studies 1 and 3 and reveal that there is a clear overall trend of decreasing knowledge attributions across the rubies, sapphires, and diamond conditions. Comparing the mean in each rubies cell with its sapphires counterpart and the latter with its diamond counterpart reveals that in every case the mean decreases. In nine out of fourteen such cases, the differences between the means were statistically significant.²⁸ Collapsing the data in each of the three columns in Table 1 and comparing them to one another reveals that the differences in mean knowledge attributions between the three condition categories are indeed statistically significant (Figure 7).²⁹

		Rubies	Sapphires	Diamond
Rather than	Thief contrast	5.95	4.42	4.20
	Jewel contrast	5.14	3.46	3.09
Knowledge-wh	Thief contrast	6.31	4.38	3.95
	Jewel contrast	3.92	(3.92)	(3.92)
Newspaper	Thief contrast	5.76	5.18	3.59
	Jewel contrast	4.49	2.87	2.45
Baseline	Baseline1	5.97	4.41	2.71
	Baseline2	6.03	4.14	2.70

Table 1. Mean knowledge attributions from the most important rubies (K1, K2, K3, K4, K5, K6, K9, K10), sapphires (S1, S2, S3, S4, S5, S6, S7), and diamond (D1, D2, D3, D4, D5, D6, D6) conditions in Studies 1 and 3.

²⁷ The mean knowledge attributions differed significantly from the midpoint in four out of the seven conditions. D1: $t(96) = .852, p > .05, r = .09$. D2: $t(94) = -4.106, p < .001, r = .39$ (medium effect size). D3: $t(36) = -.130, p > .05, r = .02$. D4: $t(26) = -1.056, p > .05, r = .20$. D5: $t(28) = -3.984, p < .001, r = .60$ (large effect size). D6: $t(37) = -3.873, p < .001, r = .54$ (large effect size). D7: $t(39) = -3.721, p < .01, r = .51$ (large effect size).

²⁸ K1 vs. S1: $U = 1283.5, z = -3.164, p < .01, r = -.27$ (small effect size). K2 vs. S2: $U = 1161.5, z = -3.834, p < .001, r = -.33$ (medium effect size). K3 vs. S3: $U = 435.5, z = -3.139, p < .01, r = -.36$ (medium effect size). K5 vs. S4: $U = 451.5, z = -.929, p > .05, r = -.12$. K6 vs. S5: $U = 354.5, z = -2.733, p < .01, r = -.33$ (medium effect size). K9 vs. S6: $U = 431.5, z = -2.559, p < .05, r = -.30$ (medium effect size). K10 vs. S7: $U = 367, z = -3.580, p < .001, r = -.42$ (medium effect size). S1 vs. D1: $U = 4239.5, z = -1.098, p > .05, r = -.08$. S2 vs. D2: $U = 4344, z = -.830, p > .05, r = -.06$. S3 vs. D3: $U = 614.5, z = -.780, p > .05, r = -.09$. S4 vs. D4: $U = 200, z = -3.059, p < .01, r = -.41$ (medium effect size). S5 vs. D5: $U = 433.5, z = -.025, p > .05, r = -.00$. S6 vs. D6: $U = 442.5, z = -2.868, p < .01, r = -.33$ (medium effect size). S7 vs. D7: $U = 480.5, z = -2.594, p < .01, r = -.30$ (medium effect size).

²⁹ $F(2, 989) = 76.639, p < .001, \text{partial } \eta^2 = .139$ (large effect size). Post-hoc Tukey’s HSD tests reveal significant differences (at the .001 level) between the overall means in the collapsed rubies condition the collapsed sapphires condition and between the overall means in the sapphires and diamond conditions.

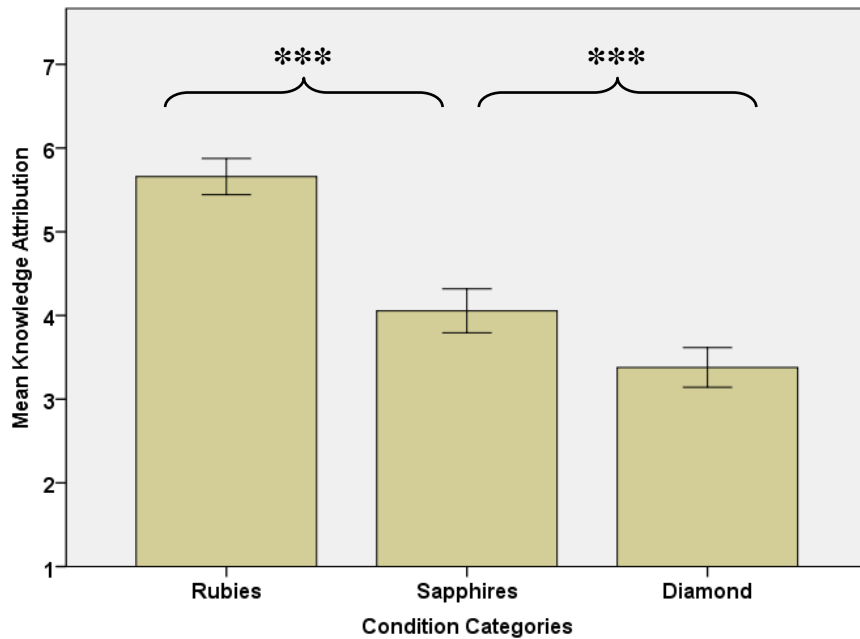


Figure 7. Mean knowledge attributions across the collapsed rubies (5.66), sapphires (4.06), and diamond (3.38) conditions represented in Table 1.³⁰

5.3. Discussion. The sapphire conditions involve a simple substitution—one that insufficiently attentive individuals might miss—whereas the diamond condition is designed to make the substitution easily detectable due to added complexity and specificity.

An overall trend suggested by the present dataset is that inclination to agree with the knowledge ascription decreases when the jewelry in the complement clause is not the object of theft described in the vignette. Moreover, the dataset suggests, agreement with the knowledge ascription decreases *more* in the diamond case in which the substitution is easier to detect. For brevity we let ‘rubies>sapphires>diamond’ denote the trend according to which the means are higher in the rubies conditions than in the sapphire conditions which, in turn, are higher than in the diamond conditions.

Let’s consider how the rubies>sapphires>diamond trend and these contrast effects bear on semantic, psychological and pragmatic accounts, respectively.

5.3.a. Semantic approaches. The rubies>sapphire>diamond trend raises a challenge to contrastivist accounts of the original data. Here we only consider contrastivist views that assume that the knowledge ascriptions in the sapphire and diamond conditions are *not true* in both thief and jewel conditions—i.e., we only consider views that assume the factivity of ‘knows.’³¹

Given factivity, it is not clear how a contrastivist may account for the rubies>sapphires>diamond trend. The motivation for contrastivism presupposes that the participant responses reflect a semantic competence that, in turn, reflects the truth-conditions of ‘knows.’ But the factivity-respecting contrastivist assumes the knowledge ascriptions to be false (not true) in *all* the thief contrast conditions—regardless of the object of theft. However, if the participants’ responses reflect the truth-conditions of ‘knows,’ the gradual decrease in inclination to agree with the thief contrast knowledge ascriptions is surprising. Contrastivism appears to lack the resources to

³⁰ The knowledge-wh rubies jewel contrast condition was omitted from this figure because it did not have sapphires or diamond counterparts.

³¹ We say ‘not true’ rather than ‘false’ in order to leave room for views according to which the knowledge ascriptions in question involve a type of presupposition failure that render them truth-valueless (Strawson 1964).

explain this trend. Of course, a contrastivist could appeal to psychological or pragmatic resources. But if such an appeal is *required*, the question for the contrastivist is this: Why are such explanations are good in some cases (when the explanandum is the trend under discussion) and bad in others (when the explanandum is the original contrast effects)?

Related challenges to contrastivism about ‘knows’ arise from the contrast effects in both the sapphire and diamond versions of the ‘rather than’ and ‘newspaper’ conditions. Given factivity, these knowledge ascriptions are false (or not true) in *both* the thief and jewel conditions. So, if participant responses reliably track the truth of contrasting knowledge ascriptions, no contrast effect should be expected. If factivity is invariably epistemically relevant, whereas it is contrast-sensitive whether the subject must have evidence bearing on the jewel-component, then contrast effects should only be expected in the latter cases. That is, contrast effects should *only* be expected in cases like Schaffer and Knobe’s original ones. So, unless the contrastivist accepts that the truth-value of the jewel-component is *also* contrast-sensitive (and thereby rejects factivity of ‘knows’), the account lacks the resources to explain the contrast effects reported above.³²

There is a methodological aspect to these challenges. The motivation for contrastivism presupposes that the contrast effects reflect semantic competence. But the present findings suggest that this is not always so. Contrasts can make participants agree with knowledge ascriptions that are false (or not true). Moreover, participants’ ability to respond correctly – i.e., to disagree with false (not true) knowledge ascriptions – is affected by the various manipulations. For example, participants were most inclined to disagree with the ‘diamond’ in the baseline and newspaper 2 conditions.

Thus, the ‘sapphire’ and ‘diamond’ data raise serious challenges to epistemic contrastivism and the methodology involved in motivating the view by the original contrast effects.

5.3.b. Psychological bias approaches. How do the present findings bear on invariantist psychological bias approaches to the original data? Bias approaches are motivated by such findings, since they claim that participants’ responses are *generally* affected by the contrast conditions. Because the ‘sapphire’ and ‘diamond’ conditions are false (or not true) in both thief and jewel conditions, the contrast effects cannot be explained by appeal to the participants’ semantic competence. So, given that there is some *principled* explanation of the contrast effects in these conditions, a psychological approach, according to which a bias is a contributing factor, deserves consideration. Indeed, the data reveals the participants’ mistake to be systematically influenced by the alleged object of theft and the way it is presented. So, it provides evidence for a *bias* as opposed to a mere *performance error* (as postulated by the shallow processing account that Schaffer and Knobe cogently criticize.)

However, the above findings are interesting not only because they provide general evidence for a psychological bias approach but also because they challenge particular brands of it. Consider, for example, a *cognitive invisibility* account, according to which participants are “blind” to the jewel alternative faces a challenge.³³ If participants were completely blind to the jewel component in the

³² In correspondence, Jonathan Schaffer has hypothesized that participants who noted the sapphire-substitution were more inclined to interpret it to be a typo than participants who noted the diamond-substitution. To consider this hypothesis, we looked at the comments section and found nothing to support it. For example, participants who indicated that they had noted the rubies-substitution were not more inclined to agree with the knowledge ascription. On the contrary, they appeared to be more inclined to strongly disagree and some participants even cited the substitution as their reason for doing so. However, a proper analysis of this data would require a coding of the comments and it would methodologically irresponsible to provide one post hoc. So, we remain content with noting that we have found no evidence for Schaffer’s hypothesis.

³³ It is not trivial how to articulate the idea that a participant is “blind” to the jewel alternative when it is out of contrast. One suggestion is that the participant is blind to the jewel alternative just in case she cannot *post hoc* correctly articulate

thief conditions, the responses should be on a par in those conditions. Since they are not, the present finding give proponents of a psychological account a reason to favor an approach postulating *inadequate processing* or a mindware gap over one that postulates *cognitive invisibility*.

Let us consider the rubies>sapphires>diamond trend. According to the focal bias account, the epistemic relevance of the jewel alternative is less likely to be adequately processed if it is out of focus – even if it is not unnoticed altogether.³⁴ Such an approach may be motivated by the present finding given two auxiliary assumptions. First, the epistemic relevance of the sapphires-component is easier to notice and process than the epistemic relevance of the rubies-component. The latter requires reflection on relevant alternatives. The former only requires appreciation of factivity.³⁵ Moreover, the mismatch between the actual object of theft and what is mentioned in the ‘sapphires’ case gives the participant a clue to critically assess the knowledge ascription that the participants in the ‘rubies’ conditions do not have available. Second, the epistemic relevance of the sapphire-component is not as easy to notice and process as the epistemic relevance of the diamond-component. After all the ‘diamond’- definite description is longer and more specific than the simple ‘sapphire.’

Given these two assumptions, the rubies>sapphire>diamond trend supports an *inadequate processing* version of the focal bias account. In the ‘rubies’ cases, it is very hard to adequately appreciate the epistemic relevance of the rubies, and it requires fairly sophisticated epistemological reflection about relevant alternatives to realize that the knowledge ascription is false in the thief conditions. But in the ‘sapphire’ cases, the participants are aided by the switch in jewelry and the fact that it does *not* require sophisticated reflection to disagree with a knowledge ascription on the basis of factivity violation. In the ‘diamond’ conditions, the switch is—because of the long definite description—even easier to notice than in the ‘sapphires’ conditions. This increases the likelihood that the switch is not missed by the participants. Thus, participants’ overall performance on the contrasted knowledge ascription task should, according to the focal bias account in question, be expected to improve in the manner of the rubies>sapphire>diamond trend. So, if the trend is robust, the inadequate processing version of the focal bias account gains evidential support.³⁶

The focal bias account explains the contrast effects in the ‘sapphires’ and ‘diamond’ conditions in the same manner that it explains the original ‘rubies’ contrast effects. According to the focal bias account, the aspects of the knowledge ascription that are not in focus are less likely to be adequately processed. While the violation of factivity may be easier to notice and process than the violation of the evidential requirement that Mary knows that rubies were stolen, it nevertheless requires processing. So, if aspects of the knowledge ascriptions that are out of focus are less likely to be adequately processed, the ‘sapphire’ and ‘diamond’ components may not be adequately processed in the thief contrasts. This explanation is consistent with the above account of the rubies>sapphires>diamond trend. To say that the ‘sapphire’ and ‘diamond’ components are more likely to be adequately processed than the ‘rubies’ component is consistent with the suggestion that they are less likely to be adequately processed in the thief than in the jewel contrast condition. In

the kind of jewelry figuring in the knowledge ascription. While the suggestion is rough, it has the advantage of being empirically testable.

³⁴ Such an account may still postulate a *cognitive illusion*. Although it does not postulate that the jewel-component is “invisible” to participants, it still claims that their judgments are illusory. There is a discrepancy between what the knowledge ascriptions truth-value *appears* to be and what it is in fact *is*. Such a discrepancy is characteristic of the phenomena typically referred to by the umbrella term ‘cognitive illusions.’

³⁵ It is remarkable that the “floor” remains fairly high in the cases in which factivity is violated. One explanation may be that some participants do overlook the replacement of ‘rubies’ with ‘sapphires.’

³⁶ There are reasons to be cautious about concluding anything stronger. The present data set provides only limited defeasible evidence for the trend in question and should be replicated with other stimuli.

fact, these two suggestions are uniform insofar as both appeal the cognitive demands on adequately processing the jewel alternative.

Thus, the focal bias provide a uniform account of the contrast effects for the original contrast effects and the contrast effects for ‘sapphire’ and ‘diamond’ Hence it gains evidential support from the present findings.

5.3.c. Pragmatic approaches. Let us briefly turn to pragmatic approaches. As with the psychological bias approaches, pragmatic approaches do not require that the participant responses align with the truth-conditions of ‘knows.’ As mentioned, this is compatible with the fact that the data suggest a discrepancy between the response patterns and the truth-conditions. Modulo factivity denial, the ‘sapphire’ and ‘diamond’ are both false (not true) in both the thief and jewel conditions. But the rubies>sapphire>diamond trend and the contrast effects show a difference in response patterns. Of course, this motivation is mainly negative. It suggest that something other than the truth-conditions of ‘knows’ explains the response patterns. So, the pragmatic approaches’ plausibility depends on their positive account of the data. In this regard, pragmatic approaches have some explanatory potential, but they also face some challenges.

A central challenge to a pure pragmatic account of the rubies>sapphire>diamond trend is to explain why participants are *more* willing to mistakenly accommodate one false presupposition (i.e., that Mary knows that the rubies were stolen) than another (i.e., that she knew that sapphires were stolen). Worse yet, why are participants more inclined to mistakenly presuppose that sapphires were stolen than that the diamond was stolen? A pure pragmatic account may be unable to explain these findings. But, as noted, pragmatic approaches may *combine* with psychological approaches. In the present case, a pragmatic account might explain how the contrasting formulations generate mistaken presupposition accommodation. In turn, a psychological bias approach might explain why some mistaken presuppositions are more likely to be accommodated than others.

It is a major task to develop such a division of explanatory labor in a manner that is not ad hoc. As mentioned, independently motivated frameworks of *cognitive pragmatics* may be of interest. But such frameworks require considerable development themselves before they generate predictions about the contrast phenomena. So, we only conclude that pragmatic approaches may gain broad motivation from the present data but that *pure* pragmatic accounts are faced with explanatory challenges suggesting that they should be integrated with psychological approaches.

6. Some null-results. In the present section, we briefly report some null results we obtained in an effort to shed light on the nature of the candidate psychological accounts.

As we noted above, every participant in Studies 1 through 3 was directed to complete the Cognitive Reflection Task (henceforth: CRT). If a focal bias account fully explained the contrast effect, it might be hypothesized that more reflective participants would be less likely to agree with knowledge ascriptions that (according to the account) were false (Toplak, West and Stanovich 2011, Stanovich 2011). We found tidbits of data that accorded with this prediction. For example, the tendency to ascribe knowledge was negatively correlated with the CRT in the diamond cases and negatively correlated with higher education level in both sapphire and diamond cases (cf. Sec. 5.2). However, there was no general pattern within the data that supported it.

We also performed various manipulations in a controlled laboratory setting that failed to generate significant results. 306 undergraduate students from a large, public university in the northeastern United States were randomly assigned to either the K1 or the K2 conditions, each of which was coupled with one of the following four manipulations:

Manipulation 1: Participants read Schaffer and Knobe’s original vignette and responded to either K1 or K2 *before* completing the cognitive reflection task, along with a need for cognition test (Cacioppo & Petty 1982; Cacioppo, Petty & Kao 1984).

Manipulation 2: Participants read Schaffer and Knobe’s original vignette and responded to either K1 or K2 *after* completing the cognitive reflection task, and then completed a need for cognition test.

Manipulation 3: Participants read Schaffer and Knobe’s original vignette and responded to either K1 or K2 before completing the cognitive reflection task, along with a need for cognition test; the research materials were printed in a very faint, light grey font that was difficult to read.

Manipulation 4: Participants read Schaffer and Knobe’s original vignette and responded to either K1 or K2 before completing the cognitive reflection task, along with a need for cognition test; participants were told beforehand that they would need to justify the answers they gave and were directed to sign their names and initial each page.

We hypothesized that having participants complete the CRT before considering contrast effect cases might put them in a more reflective frame of mind than those who completed the CRT after responding to the vignettes. We wanted to see if this reflective priming might induce participants to think more critically about the cases and thus give fewer false positives. We also hypothesized that having the research materials printed in a faint font might induce a similarly reflective frame of mind (following Alter et al 2007, Oppenheimer and Frank 2007, Oppenheimer 2008).

In the fourth manipulation—an accountability manipulation—participants received the following instruction prior to the task (Simonson and Nye 1992, Lerner and Tetlock 1999):

Please read the following story and answer the questions below, providing a short written answer for the answer you give to Question 3. Please be advised that you may be invited by the researchers conducting the study to explain and justify your choices at a later time. Furthermore, you should be aware that the researchers may include the choices and justifications you make in the present study in a booklet that will serve as the basis for class discussion in a research methods course.

Please be sure to enter your name on the first page of the study and initial each subsequent page.

On the second page, the instruction for the post hoc task read: “In the space below, please explain why you answered Question 3 on the other side as you did.” Again, the hypothesis was that the predecisional accountability condition would induce more reflective thinking which, in turn, has been shown to debias certain cognitive illusions (following Simonson and Nye 1992, Lerner and Tetlock 1999).

No significant differences were found between participants to K1 or K2 in any of the four manipulations described above. ‘Need for cognition’ refers to an individual’s tendency to engage in and enjoy effortful cognitive endeavors and is correlated with general intelligence (Cacioppo and Petty 1982; Cacioppo, Petty & Kao 1984). There were no significant correlations between participants’ need for cognition scores and their agreement or disagreement with K1 or K2.

6.3. Discussion. It is generally problematic to draw strong conclusions from null results because they represent failures to find evidence. The present case is no exception from the rule that *absence of evidence is not evidence of absence*. For example, the disfluency and predecisional accountability results does not provide evidence against a general psychological bias account of the contrast

effects. Some biases are debiased by such manipulations and some are not (Simonsen and Nye 1992, Lerner and Tetlock 1999, Oppenheimer 2008). Likewise, the lack of a correlation between performance and the CRT or the need-for-cognition does not provide evidence against a bias account. Not all biases correlate with the need-for-cognition test, the CRT or general intelligence (Stanovich and West 2008, Stanovich 2011). Furthermore, the effects of the above manipulations can be somewhat fragile and fail to manifest themselves in every experimental setting. However, Toplak et al argue that the CRT constitutes “a unique predictor of performance on heuristics-and-biases tasks” (Toplak, West and Stanovich 2011). More specifically, they claim that “...the CRT is a measure of the tendency toward the class of reasoning error that derives from miserly processing” (Toplak, West and Stanovich 2011).

In the context of discovery, then, the converging null-results may be suggestive for the *types* of biases that proponents of a bias approach should consider. For example, accounts which postulate a heuristic process of a type that is debiased by an accountability manipulation should explain why the present studies failed to indicate any debiasing effect. In particular, *cognitive invisibility* accounts according to which participants simply overlook the rubies component in the thief conditions face such an explanatory challenge. Likewise, the failure to find a correlation between knowledge ascriptions and CRT scores should be explained, given the latter’s correlation with a wide array of heuristics-and-biases tasks. In the context of discovery, the converging null-results may be suggestive with regard to the types of biases that should be considered. So, we will conclude with some speculative considerations.

Consider bias approaches according to which the contrast effect has similarities with, for example, the Levesque task (as considered in Gerken 2012: 146ff. cf. Levesque 1986). Here is a Levesque task: “*Jack is looking at Ann but Ann is looking at George. Jack is married but George is not. Is a married person looking at an unmarried person?*” The answer possibilities are: *A: Yes, B: No, C: Cannot be determined.*

The Levesque task is demanding in part because the participants, in order to give the correct answer (*A: Yes*), must provide information not that is not provided – namely the disjunction that Ann is either married or not married (for this reason the task is also called the disjunctive insight task). 86% of participants fail to do so giving the wrong answer *C: cannot be determined* (Toplak and Stanovich 2002: 203). Importantly, the Levesque task does not correlate with high intelligence (Toplak and Stanovich 2002, Stanovich 2011).

In the contrast effect cases, the participants must also contribute something in order to respond correctly (i.e., negatively, according to strict invariantism). However, the Levesque task only requires that the participant contributes the assumption that someone is either married or not. This is a fairly simple assumption that most participants are aware of. But according to anti-contrastivism, the contrast effects require that the participant contributes several fairly complex conditionals that it is not clear that the participants can contribute. Consider DeRose’s argument that the knowledge ascription “Mary knows who stole the rubies” is false. This involves the conditionals “Mary knows who stole the rubies only if Mary knows that Peter stole the rubies” and “Mary knows that Peter stole the rubies only if Mary knows that the rubies were stolen.” (DeRose 2011). Gerken provides two similar arguments. As mentioned in Sect. 2.2a, one involves the conditional “In every case, Mary knows that Peter stole the rubies only if Mary is in a position to know that the rubies were stolen” (Gerken 2013: 45). The other involves the conditionals “If Mary does not believe anything about rubies (as far as her beliefs about the theft go), Mary does not believe that Peter stole the rubies” (Gerken 2013: 47).

In contrast to the simple disjunctive assumption that is required by the Levesque task, such conditionals are complex and unlikely to be accessible to lay participants. Indeed, they are overtly theoretical assumptions that epistemologists may not even agree about. The null-results are

consistent with assuming that a simple debiasing manipulation is insufficient to put the participants in a position to supply such conditionals and go through the required reasoning. Lerner and Tetlock puts the point as follows: “Predecisional accountability to an unknown audience will have no effect on bias if, even after increased attention to one’s decision process, no new ways of solving the problem come into awareness (Lerner and Tetlock 1999: 263).

As mentioned, most participants in the Levesque task may be attributed the relevant background assumption. In contrast, the relevant assumption is not clearly attributable to the participants of the present cases: The participants are not epistemologists. We conjecture that many participants would agree with the relevant conditionals if they were presented to them.³⁷ But it is unclear that participants can themselves contribute them without specific probing. If participants are generally unable to do so, the contrast effects should be expected to be robust against simple debiasing manipulations. This assumption may be articulated by claiming that the participants have “mindware gap” (Stanovich 2009, 2011. See also Boyd and Nagel *forthcoming*).

But, as mentioned, null-results are generally compatible with a wide array of psychological accounts – including ones that postulate a complex interaction between focal bias, mindware gaps and other factors such as pragmatic ones (Gerken 2012, Sect. 7.5.2, Gerken 2013, Sect. 6.2). So, since we are keen on respecting the difference between positive evidence and mere lack thereof, we remain content with presenting the null-results and the above speculations.

In short, the null-results do not provide positive evidence for or against any account. But perhaps they may be said to be compatible with *inadequate processing* or *mindware gap* approaches whereas simple *cognitive invisibility* approaches owe an explanation of them.

7. Conclusions. The present series of experiments contribute data to the empirical study of knowledge ascriptions and, more generally, to the developing field of *folk epistemology*. Our results raise some challenges for each of the approaches to the original contrast effect for knowledge ascriptions. However, the data set raises several serious challenges for a contrastivist account. For example, the contrastivist must explain the finding of a contrast effect for belief. Moreover, the contrastivist owes an explanation of the rubies>sapphire>diamond trend as well as the finding of contrast effects in most ‘sapphire’ and ‘diamond’ conditions. We noted at the outset that the contrastivist owes a response to the armchair arguments set forth by DeRose and Gerken (DeRose 2011, Gerken 2013). Likewise, the contrastivist owes a response to the present findings. The tenability of the view depends on the plausibility of the responses to this growing list of challenges.

However, we take the present dataset to provide some evidence for approaches according to which the contrast effects are explained by more general facts about mental state ascription. Both pragmatic and psychological bias approaches may provide resources for such an explanation, and we have noted that these approaches may combine. However, we have specifically considered how psychological bias approaches such as the focal bias account may account for the present data. Interestingly, the data provides some support for pursuing an *inadequate processing* version of this approach over a *cognitive invisibility* approach. So, apart from providing negative evidence and a broad motivation, the present findings may also contribute positively to the development of a theory about contrastive knowledge ascriptions. As we have emphasized, such a theory should account for the participant response patterns in a manner that is both empirically realistic and epistemologically sound. Both the empirical data and the theorizing provided here seek to contribute to this ambition.

³⁷ A problem for testing this conjecture experimentally is the possibility that the experiments might be confounded by experimenter bias. Even if this problem could be overcome, there would be methodological questions about what the lesson from such an experiment would be. Participant agreement with an instance of a methodological principle is clearly insufficient evidence for upholding the principle.

Future work in the emerging field of folk epistemology should be integrated with existing accounts of folk psychology. The present study indicates that research on knowledge ascriptions should be integrated with research on belief ascriptions and empirical accounts of the nature of cognitive processing. On the other hand, we also hope that the present investigation exemplifies that epistemological theorizing is required for fruitful empirical work on folk knowledge ascriptions.³⁸

LITERATURE

- Aloni, Maria & Égré, Paul (2010). Alternative questions and knowledge attributions. *Philosophical Quarterly* 60 (238):1-27.
- Alter, A. L., Oppenheimer, D. M., Epley, N., & Eyre, R.N. (2007). Overcoming intuition: Metacognitive difficulty activates analytic reasoning. *Journal of Experimental Psychology: General*, 136, 569-576.
- Beaver, David I. 1997. "Presupposition," in Johan van Benthem and Alice Ter Muelen (ed.), *Handbook of Logic and Language*: 939-1008. Elsevier.
- Beaver, David and Clark, Brady (2008). *Sense and Sensitivity: How Focus Determines Meaning*. Wiley-Blackwell.
- Beebe, James R. (2013). A Knobe Effect for Belief Ascriptions. *Review of Philosophy and Psychology* 4, 235-58.
- Beebe, James and Buckwalter, Wesley (2010). The Epistemic Side-Effect Effect. *Mind & Language* 25 (2010): 474-498
- Blome-Tillmann, Michael (2013). Conversational Implicatures (and How to Spot Them). *Philosophy Compass* 8 (2):170-185
- Boyd, Kenneth and Nagel, Jennifer (forthcoming). The Reliability of Epistemic Intuitions. In Edouard Machery (ed.), *Current Controversies in Experimental Philosophy*. Routledge.
- Buckwalter, Wesley and Schaffer, Jonathan (forthcoming). Knowledge, stakes and mistakes. *Noûs*.
- Cacioppo, John T., and Petty, Richard E. 1982. The Need for Cognition. *Journal of Personality and Social Psychology*, 42, 116-31.
- Cacioppo, John T., Petty, Richard E., and Kao, Chuan Feng. 1984. The Efficient Assessment of Need for Cognition. *Journal of Personality Assessment*, 48, 306-7.
- Carston, Robyn (2002). *Thoughts and Utterances. The Pragmatics of Explicit Communication*, Oxford: Blackwell.
- DeRose, Keith (2009). *The Case for Contextualism*. Oxford University Press.
- DeRose, Keith (2011). Contextualism, contrastivism, and X-Phi surveys. *Philosophical Studies* 156 (1): 81-110.
- Dretske, Fred. 1970: Epistemic Operators, *Journal of Philosophy*, 67: 1007-23.
- Dretske, Fred. 1971: Conclusive Reasons, *Australian Journal of Philosophy*, 49: 1-22.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19 (4): 25-42.
- Gerken (2012). On the Cognitive Bases of Knowledge Ascriptions. In Jessica Brown & Mikkel Gerken (eds.), *Knowledge Ascriptions*. Oxford University Press.
- Gerken, Mikkel (2013). Epistemic Focal Bias. *Australasian Journal of Philosophy* 91 (1):41-61.
- Levesque, H. J. (1986). Making believers out of computers. *Artificial Intelligence*, 30: 81-108.
- Lerner, J.S. and Tetlock, P.E. (1999). Accounting for the effects of accountability. *Psychological Bulletin*, Vol. 108, No. 3: 255-275.

³⁸ Acknowledgements...

- Karttunen, Lauri 1974. "Presupposition and Linguistic Context," *Theoretical Linguistics* 1: 181-94.
- Kelp, Christoph (2011). A problem for contrastivist accounts of knowledge. *Philosophical Studies* 152: 287-92.
- Nagel, Jennifer (2010). Knowledge ascriptions and the psychological consequences of thinking about error. *Philosophical Quarterly* 60 (239): 286-306.
- Nagel, Jennifer (2012). Mindreading in Gettier Cases and Skeptical Pressure Cases. In Brown and Gerken (eds.), *Knowledge Ascriptions*. Oxford University Press.
- Nagel, Jennifer (2012b). Intuitions and Experiments: A Defense of the Case Method in Epistemology. *Philosophy and Phenomenological Research* 85 (3): 495-527.
- Nagel, Jennifer (2013). Defending the Evidential Value of Epistemic Intuitions: A Reply to Stich. *Philosophy and Phenomenological Research* 86 (1): 179-199.
- Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Science*, 12(6), 237-241.
- Oppenheimer, D. M., & Frank, M. F. (2007). A rose in any other font would not smell as sweet: Effects of perceptual fluency on categorization. *Cognition*, 106(3), 1178-1194.
- Rysiew, Patrick (2007). Speaking of knowing. *Noûs* 41 (4):627–662.
- Schaffer, Jonathan (2004). From contextualism to contrastivism. *Philosophical Studies* 119 (1-2): 73-104.
- Schaffer, Jonathan (2005). Contrastive knowledge. In Tamar Szabo Gendler & John Hawthorne (eds.), *Oxford Studies in Epistemology I*. Oxford University Press.
- Schaffer, Jonathan, and Knobe, Joshua (2012). Contrastive Knowledge Surveyed. *Noûs* 46 (4): 675-708.
- Schaffer, Jonathan and Szabo, Zoltan Gendler (forthcoming). Epistemic comparativism: a contextualist semantics for knowledge ascriptions. *Philosophical Studies*:1-53.
- Simonson, I. and Nye, P. (1992). The effect of accountability on susceptibility to decision errors. *Organizational Behavior and Human Decision Processes*, Vol. 51, 3: 416-446.
- Sperber, Dan and Wilson, Deirdre (1986/1995). *Relevance: Communication and Cognition*, Oxford: Blackwell. (2nd revised edition Wiley-Blackwell, 1995.)
- Stalnaker, Robert C. 1974. "Pragmatic Presupposition," in Munitz et al (eds.), *Semantics and Philosophy: Essays*. New York University Press.
- Stanovich, K. (2009). Distinguishing the reflective, algorithmic and reflective minds: Time for a tripartite theory? In (Evans and Frankish 2009): 89-107.
- Stanovich, Keith (2011). *Rationality and the Reflective Mind*. Oxford University Press.
- Steglich-Petersen *forthcoming*. Knowing the answer to a loaded question.
- Strawson, P. F. (1964). Identifying reference and truth-values. *Theoria*, 30, 96–118 (Reprinted in P. F. Strawson *Logico-Linguistic Papers*, Ashgate, London 1971).
- Toplak M. E. and Stanovich, K. E. (2002). The domain specificity and generality of disjunctive reasoning: Searching for a generalizable critical thinking skill. *Journal of Educational Psychology*, 94: 197-209.
- Toplak, M. E., West, R. F., & Stanovich, K. E. (2011). The Cognitive Reflection Test as a predictor of performance on heuristics-and-biases tasks. *Memory and Cognition*. DOI:10.3758/s13421-011-0104-1.