

Nothing at Stake in Knowledge

Original

Nothing at Stake in Knowledge / Rose, D., Machery, E., Stich, S., Alai, M., Angelucci, A., Berniunas, R., Buchtel, E.E., Chatterjee, A., Cheon, H., Cho, I.-R., Cohnitz, D., Cova, F., Dranseika, V., Lagos, A.E., Ghadakpour, L., Grinberg, M., Hannikainen, I., Hashimoto, T., Horowitz, A., Hristova, E., et al.. - In: NOÛS. - ISSN 0029-4624. - ELETTRONICO. - 53:1(2019), pp. 224-247. [10.1111/nous.12211]

Availability:

This version is available at: 11583/2960299 since: 2024-08-20T14:52:22Z

Publisher:

Blackwell Publishing Ltd

Published

DOI:10.1111/nous.12211

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Wiley postprint/Author's Accepted Manuscript

This is the peer reviewed version of the above quoted article, which has been published in final form at <http://dx.doi.org/10.1111/nous.12211>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions.

(Article begins on next page)

Nothing at Stake in Knowledge

Many philosophers hold that stakes affect ordinary knowledge ascriptions. Here's a version of a pair of cases aimed at supporting this: Bob and his wife are driving home on Friday and considering whether to stop at the bank to deposit a check. The lines at the bank are very long and so Bob considers coming back on Saturday. In the low stakes version, nothing of importance hinges on whether the check is deposited; in the high stakes version, it is very important that the check be deposited. Bob's wife asks whether the bank will be open on Saturday. Bob says he drove past the bank last Saturday, and it was open. However, his wife points out that banks sometimes change their hours. Bob says "I know the bank will be open tomorrow". In the low stakes case, many philosophers maintain that Bob does indeed know that the bank will be open; in the high stakes case, these philosophers maintain that Bob is ignorant – his statement that he knows the bank will be open tomorrow is false. These philosophers also maintain that this pattern of judgments is what we would expect from competent speakers confronted with this and similar cases (e.g., Cohen, 1999, 2013; DeRose, 1992, 2009; Fantl and McGrath, 2002; Nagel, 2008; Rysiew, 2001; Stanley, 2005).

Though many philosophers agree that stakes play a role in ordinary knowledge ascriptions, there is disagreement about what explains this. One view, *epistemic contextualism*, holds that "to know" is a context sensitive verb and that the truth conditions for knowledge ascriptions can vary across conversational contexts (e.g., DeRose, 2009). For instance, Bob's statement "I know the bank will be open tomorrow" can be true in low stakes contexts and false in high stakes contexts. Another view, *interest-relative invariantism*, denies that "to know" is a context sensitive verb and that the truth conditions for knowledge ascriptions vary according to conversational contexts. Instead, cases like the Bank cases show that practical factors—i.e., stakes—play a distinctive role in determining whether the knowledge relation obtains (e.g., Stanley, 2005). Yet another alternative, which we'll call *classical invariantism*, denies that "to know" is a context sensitive verb and that practical factors, such as stakes, play a direct role in determining whether the knowledge relation obtains. Instead, stakes affect knowledge ascriptions only by affecting our assessment of factors that have traditionally been taken to constitute or be necessary for knowledge, such as e.g., belief, quality of evidence, etc. (e.g., Bach, 2005; Weatherson, 2005; Ganson, 2007; Nagel, 2008). If this is right, then the role of stakes in knowledge ascriptions fails to motivate such surprising views as epistemic contextualism or interest-relative invariantism. Naturally, epistemic contextualists and interest-relative invariantists deny this, claiming that even when the factors that have traditionally been taken to constitute or be necessary for knowledge are held fixed, stakes continue to play a role in ordinary knowledge ascriptions (e.g., DeRose, 2009; Lawlor, 2013).

So we see a dispute over what best explains the role of stakes in ordinary knowledge ascriptions. It is thus extremely surprising that a wide range of empirical evidence suggests that ordinary knowledge ascriptions fail to display any sensitivity to stakes (e.g., Buckwalter, 2010; Buckwalter and Schaffer, 2015; Feltz and Zarpentine, 2010; May, Sinnott-Armstrong, Hull, and Zimmerman, 2010; Turri, forthcoming; though see e.g., Pinillos, 2012; Pinillos and Simpson, 2014; Sripada and Stanley, 2012). If stakes really do not play any role in ordinary knowledge ascriptions, one of the main motivations for epistemic contextualism and interest relative invariantism would be undermined. Perhaps these different explanations of the role of stakes in ordinary knowledge ascription are born out of nothing more than a myth (Schaffer and Knobe, 2009). If so, classical invariantism about knowledge might be best

supported—not because it provides the best explanation of the role of stakes in ordinary knowledge ascriptions, but rather because the failure of stakes to play a role in ordinary knowledge ascription would undercut an important motivation for its two competitors, epistemic contextualism and interest-relative invariantism. These radical alternatives to classical invariantism, lacking evidence in support of one of their important motivations, should perhaps then fall. Classical invariantism would stand.

In the remainder of this article, we'll disarm an important motivation for epistemic contextualism and interest-relative invariantism. We'll accomplish this by presenting a stringent test of whether there is a stakes effect on ordinary knowledge ascription. Having shown that, even on a stringent way of testing, stakes fail to impact ordinary knowledge ascription, we will conclude that we should take another look at classical invariantism.

Here is how we will proceed. Section 1 lays out some limitations of previous research on stakes. Section 2 presents our study and concludes that there is little evidence for a substantial stakes effect. Section 3 responds to objections. The conclusion clears the way for classical invariantism.

1. Strengthening the Case for Stakes

The role of stakes in ordinary knowledge ascriptions is taken to be illustrated by patterns of judgments allegedly made by competent speakers. While a number of philosophers have taken for granted the sensitivity to stakes of knowledge ascription among competent speakers, empirical evidence has suggested otherwise. A wide range of empirical research has failed to uncover evidence that stakes play a role in ordinary knowledge ascriptions (see below for discussion of evidence seemingly supporting such role). In light of this evidence, it would be tempting to conclude that perhaps stakes do not, after all, play a role in ordinary knowledge ascription. As tempting as this may be, however, it seems to us that there are a number of issues that have yet to be addressed.

First, virtually every study exploring the role of stakes in knowledge ascription has failed to ask participants whether they view the situation as a high- or low-stakes situation. This is especially surprising given that the main finding is a null result i.e., that knowledge ascriptions do not differ between high and low stakes cases. It may well be that the manipulation—i.e., high vs. low stakes—was ineffective, perhaps because participants failed to pay sufficient attention to key details varying between the cases. For instance, in the Bank cases, it may be that participants fail to appreciate that in one case it is “very important” that a check be deposited, while in the other it is “not very important”. If so, then the fact that no stakes effect was found wouldn't show that competent speakers fail to display a sensitivity to stakes. So, in the study reported below, the first question examined whether participants have understood what was at stake.

Another, perhaps more serious issue is that some of the results suggesting that stakes fail to play a role in knowledge ascriptions might be due to protagonist projection. Protagonist projection occurs when a subject takes up a protagonist's perspective and imagines what seems true from the protagonist's point of view (Holton, 1997). Importantly, protagonist projection looks to be at least partly responsible in producing otherwise surprising findings. For instance, Weinberg, Nichols and Stich (2001) present evidence that Westerners and East Asians view Gettier cases differently: Westerners intuit that protagonists in Gettier cases do not know the relevant proposition while East Asians intuit that protagonists in Gettier cases do know the relevant proposition. But recent research suggests that the differences uncovered by Weinberg, Nichols and Stich disappear when taking into account participants' tendency to engage in protagonist projection. Indeed, recent work by Machery et al. (2015) found that, across four cultures—USA, India, Japan and Brazil—rates of knowledge denial when confronted with Gettier cases were both high and similar across each of the four cultures sampled when a question

targeting protagonist projection was introduced. While some participants attributed knowledge to a Gettierized protagonist when asked whether that protagonist “knows” or “does not know”, when given a question aimed at probing for protagonist projection—i.e., being asked whether the Gettierized protagonist “really knew” or “didn’t really know but only thought she knew”—rates of knowledge denial increased with the vast majority of participants indicating that the Gettierized subject “didn’t really know” but “only thought that she knew.”

Similarly, protagonist projection also looks to be behind apparently non-factive knowledge ascriptions such as “Everyone knew that stress caused ulcers, before two Australian doctors in the early 1980s proved that ulcers are actually caused by bacterial infection” (Buckwalter, 2014, p. 396).

Recent work by Turri (forthcoming) suggests that something like protagonist projection—what he calls “deferral”—may play a crucial role in some of the cases aimed at showing stakes play a role in knowledge ascriptions. Instead of people’s knowledge ascriptions shifting along with variations in stakes, people may simply defer to others’ mental-state reports. For instance, in one version of the Bank cases (but not in the version we used), the protagonist in the high-stakes case says, “I don’t know it will be open tomorrow,” while the protagonist in the low stakes case says, “I know the bank will be open tomorrow” (see, e.g., DeRose, 2009, 2011). In this version of the Bank cases, agreement with both statements might have nothing to do with a shift in stakes. Instead, people may naturally defer to the mental state reports of others.

In addition to producing the misleading appearance of a stakes effect, protagonist projection can also *mask* a genuine stakes effect. This can happen when the protagonist says in both conditions, “I know the bank will be open tomorrow,” as is the case in the vignettes we used. To ensure that protagonist projection does not mask any genuine effect of stakes, we introduced a probe aimed at capturing whether participants are making genuine knowledge ascriptions or merely projecting: It contrasts “knows” and “thinks he knows, but doesn’t actually know.” By comparing participants’ answer to this probe to their answers to a “know/does not know” probe, we will also be able to examine the role of protagonist projection in knowledge ascription in a context where stakes are manipulated.

Finally, the failure of stakes to play a role in knowledge ascription may only reflect something peculiar about the practice of knowledge ascription within a narrow linguistic community. Indeed, all of the empirical work done thus far has been conducted with participants drawn from the USA. Perhaps a wide range of other linguistic communities display a sensitivity to stakes. So we would like to know whether the effect of stakes (or lack thereof) in knowledge ascription is cross culturally robust. Putting all of this together, our questions are:

- Is the difference in stakes appreciated by participants?
- Is the lack of a stakes effect due to protagonist projection?
- Is the effect of stakes (or lack thereof) in knowledge ascriptions cross-culturally robust?

In taking these up, our strategy was to undertake a cross cultural study, introducing a number of measures aimed at addressing the questions under consideration in order to determine whether stakes sensitivity (or the lack thereof) reflects a core aspect of folk epistemology.

2. Nothing at Stake in Knowledge: A Study in Folk Epistemology

2.1. Method

We collected data from 4504 people across nineteen sites, spanning fifteen countries. Each participant was randomly assigned to one of two conditions, a Low- or High-stakes version of a Bank case. Here is the Low-Stakes version:

Bob and his wife are driving home on a Friday afternoon. They both received some money earlier in the day, and so they plan to stop at the bank on the way home to deposit it. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Although they generally like to deposit any money they receive at the bank as soon as possible, it is not especially important in this case that it be deposited right away, and so Bob suggests that they drive straight home and deposit their money on Saturday morning. His wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." Bob replies, "No, I know the bank will be open. I was just there two weeks ago on Saturday. It was open until noon." As a matter of fact, the bank will be open on Saturday morning.

And here is the High Stakes version:

Bob and his wife are driving home on a Friday afternoon. They both received some money earlier in the day and so they plan to stop at the bank on the way home to deposit it. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. They have recently written a very large and very important check. If the money is not deposited into their bank account before Monday morning, the important check they wrote will not be accepted by the bank, leaving them in a very bad situation. Bob suggests that they drive straight home and deposit their money on Saturday morning. His wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." Bob replies, "No, I know it'll be open. I was just there two weeks ago on Saturday. It was open until noon." As a matter of fact, the bank will be open on Saturday morning.

To see whether stakes are appreciated, participants were first asked:

Comprehension: According to the story, which of the following statements is correct? [It is not very important that Bob and his wife deposit their money/It is very important that Bob and his wife deposit their money.]

They were then asked:

Knowledge Attribution: In your personal opinion, when Bob says "I know the bank will be open" is his statement true? [Yes, Bob's statement is true./No, Bob's statement is not true.]

And finally to see if responses to Knowledge Attribution are due to protagonist projection, we asked:

Strict Knowledge Attribution In your personal opinion, which of the following sentences better describes Bob's situation? [Bob knows the bank will be open on Saturday./Bob thinks he knows the bank will be open on Saturday, but he doesn't actually know it will be open.]

The cases were translated into fourteen languages by competent native speakers and presented in the respective native language for each group.

2.2. Results

First, do participants appreciate the difference in stakes? It turns out that they do. Overall, 78% of participants passed Comprehension (see Table 1 for Demographics). So we take our first concern—that participants may not be appreciating a difference in stakes—to be resolved.

Table 1. Demographic information about the study’s participants who passed Comprehension including countries in which data were collected, nature of the sample (students vs. non-students), mode of survey administrations (paper-pencil vs web-based, volunteers vs. in exchange for compensation, language of the survey).

Sample	Students	Method	Payment	Language	<i>N</i>
Europe					
Bulgaria	N	Web-based	Volunteers	Bulgarian	327
France	N	Web-based	Compensation & volunteers	French	367
Germany	N	Web-based	Compensation	German	153
Italy	Y	Paper-pencil	Volunteers	Italian	139
Portugal	Y	Paper-pencil	Volunteers	Portuguese	139
Spain	N	Web-based	Compensation	Spanish	239
Switzerland	N	Paper-pencil & web-based	Volunteers	French	54
Switzerland	Y	Paper-pencil & web-based	Compensation & volunteers	French	30
UK	N	Web-based	Compensation	English	255
Middle East					
Iran	N	Paper-pencil	Volunteers	Persian	164
Central & North America					
Mexico	N	Paper-pencil	Volunteers	Spanish	133
USA	N	Web-based	Compensation	English	225
South America					
Brazil	Y	Paper-pencil	Volunteers	Portuguese	135
East Asia					
China	Y	Paper-	NA	NA	128

Guangzhou China Mainland China	Y	pencil Paper- pencil	Volunteers	Chinese, Simplified	109
Hong Kong	N	Web- based	Compensation	Chinese, Simplified	180
Japan	Y	Web- based	Compensation	Chinese, Traditional	146
Japan	N	Web- based	Compensation	Japanese	151
Japan	Y	Paper- pencil	Volunteers	Japanese	165
Mongolia	N	Paper- pencil	Volunteers	Mongolian	115
South Asia India	Y	Paper- pencil	Volunteers		162

Next, do stakes affect Knowledge Attribution? Removing participants who failed Comprehension, we analyzed responses from the remaining 3530 participants. Overall, we found a significant, but negligible effect of stakes on Knowledge Attribution ($\chi^2(1, 3530)=9.040, p<.01$; see Figure 1). Ascription of knowledge in the low stake condition (85%) is only 3% larger than in the high stake condition (82%), and its conventional effect size is very small (Cramer's $V=.051$): On a standard interpretation of conventional effect sizes (see Ellis, 2010) it does not even count as small.¹

¹ We follow Ellis (2010) in interpreting the magnitude of the effect sizes. For Cramer's V we interpret values greater than or equal to .5 as large, greater than or equal to .3 but less than .5 as medium, and greater than or equal to .1 but less than .3 as small.




Figure 1: Overall Effect of Stakes on Knowledge Attribution

More importantly, across sites we find virtually no evidence that stakes affect knowledge attribution (see Table 2 and Figure 2). The exceptions are in italics.

Table 2: Effect of Stakes on Knowledge Attribution for Each Site

Sample	N	χ^2	p-value	Cramer's V
Central and North America				
Mexico	133	.000	.990	.001
USA	225	3.661	.056	.128
South America				
Brazil	135	.318	.573	.049
Europe				
Bulgaria	327	.084	.773	.016
France	367	2.990	.084	.090
Germany	153	.555	.456	.060
Italy	139	.117	.732	.029
Portugal	139	1.229	.268	.094
<i>Spain</i>	239	<i>6.219</i>	<i>.013</i>	<i>.161</i>
Switzerland	84	.841	.359	.100
<i>UK</i>	255	<i>4.470</i>	<i>.030</i>	<i>.132</i>
Middle East				
Iran	164	.020	.889	.011
East Asia				
China	128	.522	.470	.064
Hong Kong	146	.272	.602	.043

Guangzhou China	109	.690	.406	.080
Mainland China	180	1.345	.246	.086
Mongolia	115	.003	.959	.005
Japan	316	5.728	.017	.135
South Asia				
India	162	1.747	.186	.104




Figure 2: Rates of Knowledge Attribution for High- and Low-Stakes Cases for Each Site

Out of the nineteen sites sampled, only three (16%)—Spain, UK and Japan—displayed a significant, small-sized effect of stakes on knowledge ascriptions (the data from the USA is also near significant; see Figures 3 and 4). But the difference that stakes produced in these three sites is not significantly different from France, which displayed no difference of stakes on Knowledge Attribution (see Table 3).

Table 3: Difference in Stakes on Knowledge Attribution with France as Contrast Class²

Sample	B(SE)	Wald χ^2	p-value	Exp(B)
Central and North America				

² A logistic regression model with stakes, site and an interaction between stakes and site was run. Stakes did not predict Knowledge Attribution, Wald $\chi^2=2.897$, $p=.089$; site significantly predicted Knowledge Attribution, Wald $\chi^2=65.376$, $p=.000$; and there was no significant interaction between stakes and site on Knowledge Attribution, Wald $\chi^2=22.314$, $p=.218$.

USA	.231(.641)	.130	.718	1.260
Mexico	-.686(.579)	1.405	.236	.503
South America				
Brazil	-.942(.602)	2.445	.118	.390
Europe				
Bulgaria	-.614(.486)	1,601	.206	.541
France (Contrast)	---	---	---	---
Germany	-4.12(.553)	.556	.456	.662
Italy	-.872(.666)	1.716	.190	.418
Portugal	-1.237(.641)	3.723	.054	.290
Spain	.162(.535)	.091	.762	1.176
Switzerland	-1.314(.798)	2.713	.100	.269
UK	.142(.571)	.061	.804	1.152
Middle East				
Iran	-.740(.535)	1.914	.166	.477
East Asia				
China	-1.041(.633)	2.706	.100	.353
Hong Kong	-.360(.756)	.227	.633	.697
Guangzhou China	-1.110(.649)	2.924	.087	.329
Mainland China	.564(.331)	.000	.996	.000
Mongolia	-.669(.589)	1.289	.256	.512
Japan	.027(.507)	.003	.957	1.027
South Asia				
India	-.077(.621)	.015	.901	.926




Figure 3: Difference Between Low and High Stakes Cases for Those Who Attributed Knowledge on Knowledge Attribution




Figure 4: Effect Size (Using Cramer’s V) of Stakes on Knowledge Attribution for Each Site

Moreover, these three linguistic communities, along with every other linguistic community sampled, displayed overall high rates of knowledge attribution regardless of whether the case was low or high stakes (Figure 2). These results fit with a range of similar findings³ and extend the finding that stakes fail to play a role in knowledge ascription to a range of linguistic communities across the globe.

Given that we find virtually no evidence of a stakes effect on Knowledge Attribution, we want to know whether this may be due in part to protagonist projection. We look at this in two ways, first, using our second measure (“Strict Knowledge Attribution”). Overall, we again find a significant, but negligible effect of stakes on Strict Knowledge Attribution $\chi^2(1, 3522)=10.451, p<.01$, (see Figure 5). Ascription of knowledge in the low stake condition (63%) is only 5% larger than in the high stake condition (58%), and its conventional effect size is tiny (Cramer’s $V=.051$): Again, on a standard interpretation of conventional effect sizes (see Ellis, 2010) it does not even count as small.

³ See e.g., Buckwalter, 2010; Buckwalter and Schaffer, 2015; Feltz and Zarpentine, 2010; May, Sinnott-Armstrong, Hull and Zimmerman, 2010; Turri, forthcoming




Figure 5: Overall Effect of Stakes on Strict Knowledge Attribution

Moreover, across sites, we continue to fail to find evidence of a stakes effect on Strict Knowledge Attribution (see Table 4 and Figure 4). The exceptions are in italics.

Table 4: Effect of Stakes on Strict Knowledge Attribution for Each Site

Sample	N	χ^2	p-value	Cramer's V
Central and North America				
Mexico	131	1.740	.187	.115
<i>USA</i>	225	<i>5.453</i>	<i>.020</i>	<i>.156</i>
South America				
Brazil	135	2.603	.107	.139
Europe				
Bulgaria	327	.248	.618	.028
France	365	.505	.477	.037
<i>Germany</i>	<i>151</i>	<i>3.860</i>	<i>.049</i>	<i>.160</i>
Italy	139	.024	.878	.013
Portugal	139	.033	.856	.015
Spain	239	1.930	.165	.090
Switzerland	84	.350	.554	.065
UK	252	2.640	.104	.102
Middle East				
Iran	164	2.159	.142	.115
East Asia				
China	128	2.092	.148	.128

Hong Kong	146	2.776	.096	.138
Guangzhou China	109	.002	.969	.004
Mainland China	180	.479	.489	.052
Mongolia	116	.556	.456	.069
Japan	316	2.988	.084	.097
South Asia				
India	162	.144	.704	.030




Figure 6: Rates of Strict Knowledge Attribution for High and Low Stakes Cases for Each Site

Indeed, only two linguistic communities (11%)—Germany and USA—out of the nineteen sampled displayed evidence of a significant, small-sized effect of stakes on Strict Knowledge Attribution (Figures 7 and 8). But the difference that stakes produced in these two sites is not significantly different from France, which displayed no difference of stakes on Strict Knowledge Attribution (see Table 5). So, on our first way of looking at whether a stakes effect might be masked by projection, we find virtually no evidence that stakes affect knowledge attribution.

Table 5: Difference in Stakes on Strict Knowledge Attribution with France as Contrast Class⁴

Sample	B(SE)	Wald χ^2	p-value	Exp(B)
Central and North America				
USA	.477(.375)	1.621	.203	1.611
Mexico	.330(.458)	.519	.471	1.390
South America				
Brazil	-.769(.445)	2.990	.084	.464
Europe				
Bulgaria	-.061(.339)	.032	.857	.941
France (Contrast)	---	---	---	---
Germany	.539(.443)	1.483	.223	1.715
Italy	-.122(.429)	.081	.776	.885
Portugal	-.111(.437)	.064	.800	.895
Spain	.187(.360)	.269	.604	1.206
Switzerland	.090(.514)	.031	.861	1.094
UK	.248(.360)	.473	.492	1.281
Middle East				
Iran	-.640(.402)	2.535	.111	.528
East Asia				
China	.364(.449)	.656	.418	1.438
Hong Kong	.450(.452)	.991	.320	1.568
Guangzhou China	-.191(.459)	.173	.677	.826
Mainland China	-.504(.536)	.884	.347	.604
Mongolia	-.481(.480)	1.009	.315	.618
Japan	.334(.387)	.747	.387	1.397
South Asia				
India	-.050(.415)	.014	.905	.952

⁴ A logistic regression model with stakes, site and an interaction between stakes and site was run. Stakes did not predict Strict Knowledge Attribution, Wald χ^2 =.505, p=.478; site significantly predicted Strict Knowledge Attribution, Wald χ^2 =139.090, p=.000; and there was no significant interaction between stakes and site on Strict Knowledge Attribution, Wald χ^2 =22.071, p=.223.




Figure 7: Difference Between Low and High Stakes Cases for Those Who Selected “Really Knows” on Strict Knowledge Attribution for Each Site




Figure 8: Effect Size (Using Cramer’s V) of Stakes on Strict Knowledge Attribution for Each Site

On the second way of looking at whether the lack of a stakes effect might be due to protagonist projection, we looked at whether there was an effect of stakes on Strict Knowledge Attribution among those who attributed knowledge on Knowledge Attribution. Here we find a marginally significant effect $\chi^2(1, 2923)=3.567, p=.059$, Cramer’s $V=.035$: Of those who attributed knowledge on Knowledge Attribution in the high-stakes case, 68% selected “really knows” on Strict Knowledge Attribution; of those who attributed knowledge on Knowledge Attribution in the low-stakes case, 71% selected “really knows” on Strict Knowledge Attribution. This amounts to a mere 3% difference (see Figure 9).

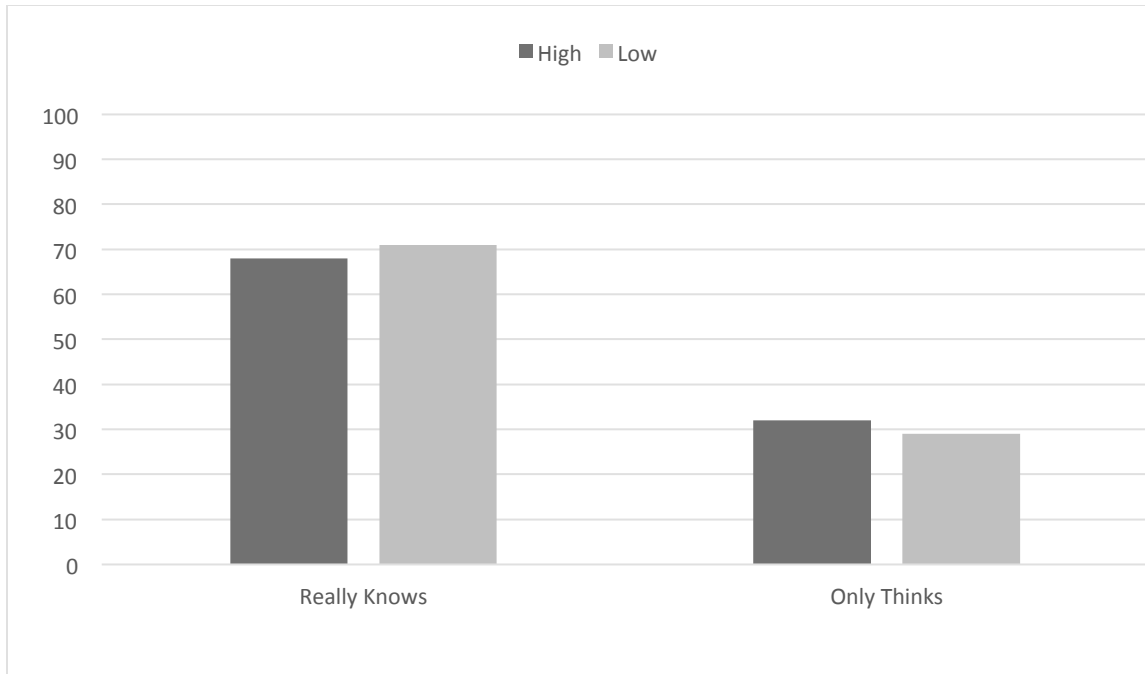


Figure 9: Overall Effect of Stakes on Strict Knowledge Attribution Among Those Who Attributed Knowledge on Knowledge Attribution

Yet again, across sites, we continue to find virtually no evidence of a stakes effect on Strict Knowledge Attribution among those who attributed knowledge on Knowledge Attribution (Table 6 and Figure 10).

Table 6: Effect of Stakes on Strict Knowledge Attribution Among Those Who Attributed Knowledge on Knowledge Attribution.

Sample	N	χ^2	p-value	Cramer's V
Central and North America				
Mexico	92	1.448	.229	.125
USA	202	3.182	.074	.126
South America				
Brazil	110	2.303	.129	.145
Europe				
Bulgaria	245	.056	.812	.015
France	333	.003	.995	.003
Germany	109	1.351	.245	.111
Italy	122	.095	.758	.028
Portugal	120	.650	.420	.074
Spain	190	.297	.586	.040
Switzerland	74	.724	.395	.099
UK	218	1.327	.249	.078

Middle East				
Iran	118	2.013	.156	.131
East Asia				
China	107	.613	.434	.076
Hong Kong	134	3.581	.058	.163
Guangzhou China	90	.356	.551	.063
Mainland China	179	.210	.617	.034
Mongolia	80	1.074	.300	.116
Japan	249	.034	.854	.012
South Asia				
India	139	.416	.519	.055




Figure 10: Rates of Strict Knowledge Attribution for High and Low Stakes Cases Among Those Who Attributed Knowledge on Knowledge Attribution for Each Site

On this way of looking at whether a stakes effect might be masked by projection, we find that stakes had no significant effect in any site, and approached significance in only two sites: Hong Kong and USA (Figures 11 and 12). Yet again, using France, which displayed no significant difference in stakes, as a contrast, we find that the difference in stakes produced in each site does not differ from the lack of difference in stakes in France (see Table 7).

Table 7: Difference in Stakes on Strict Knowledge Attribution Among Those Who Attributed Knowledge on Knowledge Attribution with France as Contrast Class⁵

Sample	B(SE)	Wald χ^2	p-value	Exp(B)
Central and North America				
USA	.523(.417)	1.572	.210	1.688
Mexico	.570(.568)	1.009	.315	1.769
South America				
Brazil	-.613(.488)	1.577	.209	.542
Europe				
Bulgaria	.048(.393)	.015	.903	1.049
France (Contrast)	---	---	---	---
Germany	.641(.640)	1.004	.316	1.899
Italy	.102(.477)	.046	.831	1.107
Portugal	.329(.516)	.407	.524	1.389
Spain	.147(.413)	.126	.723	1.158
Switzerland	.415(.583)	.506	.477	1.514
UK	.317(.407)	.606	.436	1.373
Middle East				
Iran	-.544(.470)	1.337	.248	.580
East Asia				
China	.309(.504)	.374	.541	1.361
Hong Kong	.776(.551)	2.301	.129	2.173
Guangzhou China	.248(.528)	.221	.638	1.282
Mainland China	-.239(.565)	.178	.673	.788
Mongolia	-.500(.549)	.829	.363	.607
Japan	.058(.493)	.014	.906	1.060
South Asia				
India	-.274(.492)	.310	.578	.760

⁵ A logistic regression model with stakes, site and an interaction between stakes and site was run. Among those who attributed knowledge on Knowledge Attribution, Stakes did not predict Strict Knowledge Attribution, Wald $\chi^2 = .003$, $p = .955$; site significantly predicted Strict Knowledge Attribution Wald $\chi^2 = 99.536$, $p = .000$; and there was no significant interaction between stakes and site on Strict Knowledge Attribution Wald $\chi^2 = 17.307$, $p = .502$.




Figure 11: Difference Between Low and High Stakes Cases for Those Who Selected “Really Knows” on Strict Knowledge Attribution Among Those Who Attributed Knowledge on Knowledge Attribution for Each Site




Figure 12: Effect Size (Using Cramer’s V) of Stakes on Strict Knowledge Attribution Among Those Who Attributed Knowledge on Knowledge Attribution for Each Site

Even on our second way of looking at whether a stakes effect might be masked by projection, we continue to find virtually no evidence that stakes affect knowledge attribution. We conclude that the lack of a stakes effect on knowledge ascriptions cannot be dismissed by appealing to protagonist projection.

2.3 Discussion

Given these results and the wide swath of research indicating that stakes do not play a role in ordinary knowledge ascription, the scales tilt against epistemic contextualism and interest-relative invariantism, at least to the extent that they attempt to account for everyday knowledge ascription. One of the important motivations for these views—that stakes play a role in ordinary knowledge attributions—is undermined. These views are “idle hypotheses” (Turri, forthcoming).

Interestingly, stakes had a marginally significant effect on both Knowledge Attribution and Strict Knowledge Attribution among those attributed knowledge on Knowledge Attribution for the USA sample. Stakes also had a significant, but small effect for the Strict Knowledge Attribution question. Perhaps this explains why some epistemologists, the most influential of whom are Americans, thought there was a stakes effect. In any case, the small effect of stakes is a far cry from what we would expect if stakes played an important role in knowledge ascription.

3. Objections and Responses

3.1 Evidence for Stakes Effect

Some experimental studies claim to have found some evidence for a stakes effect (Pinillos, 2012; Sripada & Stanley, 2012). These results clearly conflict with our findings as well as a range of other research. So perhaps there really is a genuine stakes effect on ordinary knowledge ascription.

We are skeptical that these studies provide genuine evidence that there is a stakes effect on ordinary knowledge ascription. Pinillos (2012) gave people a case about an individual, Peter, writing a paper for an English class. In the low-stakes version, it is not very important whether the paper has typos; in the high-stakes version, it is very important that the paper not have typos. What Pinillos found was that when participants were asked, “How many times do you think Peter has to proofread his paper before he knows that there are no typos?”, the median response was 2 in the low-stakes version while the median response was 5 in the high-stakes version. Pinillos interprets this as evidence that stakes do indeed affect ordinary knowledge ascriptions.

But Buckwalter (2014) and Buckwalter and Schaffer (2015) present compelling evidence that the effect Pinillos found has nothing to do with knowledge. The effect persists when “knows” in the probe Pinillos used is replaced with “believes,” “guesses,” and “hopes.” Rather than being an effect on knowledge ascription, the stakes effect Pinillos has uncovered is instead an effect on the modal expression “has to” (Buckwalter and Schaffer, 2015).

In Sripada and Stanley (2012), participants read about an individual Hannah who is allergic to Mongolian pine nuts. She is having dinner at a Mongolian restaurant, and the possibility is raised that there may be Mongolian pine nuts in her dish. They found an effect of stakes on knowledge ascription in two separate pairs of cases. However, the effect looks to be quite small in both cases (less than a 1 point scale difference on a 7 point scale).⁶ Moreover, when inspecting the graphs of their results (2012, 15), the mean responses in both the low- and high-stakes version for both pairs of cases do not seem to be significantly different from the midpoint of “neutral.” That is, in the cases used by Sripada and Stanley, participants are actually neutral about whether to ascribe knowledge in both the high- and low-stakes versions of these cases. Stakes may have an effect, a small one at that, only when people are uncertain about ascribing knowledge. We doubt that these results are very encouraging to proponents of epistemic contextualism and interest relative invariantism since “They remain a far cry from the strong flip from “knowledge” to “ignorance” which DeRose, Stanley and many other epistemologists had predicted from the armchair...” (Buckwalter and Schaffer, 2015, 221). Moreover, there is good empirical reason for thinking that the “stakes effect” Sripada and Stanley claim to have uncovered is confounded with salience and that the effect is instead a salience effect and not a stakes effect (see Buckwalter and Schaffer, 2015).

3.2 The Role of Linguistic Evidence

Proponents of contextualism or interest-relative invariantism may object that their theories are not meant to account for everyday knowledge ascriptions. They are, after all, theories of knowledge, rather than linguistic theories about “to know.”

We have two brief responses to this concern. To the extent that contextualism and interest-relative invariantism are not meant to account for knowledge ascriptions—perhaps they are theories of knowledge to be developed largely independently of how people talk about knowledge and what they think about it—

⁶ We say it “looks” small because we can’t actually calculate the effect size since Sripada and Stanley do not report means and standard deviations.

then, we acknowledge, our findings have little to say about contextualism and interest-relative invariantism. They are only relevant for these philosophical views that aim at accounting for everyday knowledge ascription.

We add that as a matter of fact, many contextualists and interest-relative invariantists propose to account for everyday knowledge ascription. DeRose (1992) formulates contextualism as a semantic theory and he engages with the early experimental philosophy literature on stakes effects in bank cases (2011). Ludlow (2005, 11) too explicitly formulates contextualism as a semantic thesis: “According to the thesis of contextualism in epistemology, many of our knowledge attributions (including self-attributions) are context sensitive.”

3.3 A Negative Result Limited to the Bank Cases

One might argue that our results merely show that the bank cases are inappropriate to elicit a stakes effect, not that there is no stakes effect. Other cases would elicit a stakes effect. However, the bank cases were put forward by some advocates of the view that stakes affect ordinary knowledge ascriptions as being the *best* cases for eliciting a stakes effect (Buckwalter and Schaffer, 2015, 222). Furthermore, even Sripada and Stanley, who think that their pine nut cases are the best cases for eliciting a stakes effect on ordinary knowledge ascription, find *at best* a very small effect of stakes, which shifts people from being slightly more to slightly less neutral about whether to ascribe knowledge. Finally, other studies have used a range of cases with each failing to uncover an effect of stakes on ordinary knowledge ascription (see e.g., Feltz and Zarpentine, 2010; Turri, forthcoming; Turri and Buckwalter, forthcoming). Indeed, in light of our findings and a range of research failing to uncover a stakes effect in ordinary knowledge ascriptions, we think that stakes fail to reflect a core aspect of folk epistemology.

4. Conclusion

Epistemic contextualists hold that knowledge ascription shifts across conversational contexts; interest-relative invariantists recognize a distinctive epistemic role for practical factors such as stakes in knowledge ascription. If either of these is correct, serious pressure is put on classical invariantism since the classical invariantist rejects both claims. However, a wide range of empirical evidence now suggests that stakes do not play any role in ordinary knowledge ascription.

We set out to provide what we take to be a stringent test of whether stakes play a role in ordinary knowledge ascriptions. In doing so we pursued three main questions:

- Is the difference in stakes appreciated by participants?
- Is the lack of a stakes effect due to protagonist projection?
- Is the effect of stakes (or lack thereof) in knowledge ascriptions cross culturally robust?

Our results indicate that subjects do indeed appreciate the difference in stakes, that a stakes effect is not being masked by protagonist projection, and that the lack of a stakes effect in ordinary knowledge ascription is cross-culturally robust. In light of our evidence and a wide range of previous empirical findings on the role of stakes in knowledge ascription, one of the core motivations for epistemic contextualism and interest relative invariantism is undercut. Although we won't defend this claim in detail here, we conclude that classical invariantism should be taken seriously, now that its challengers have been undermined.

References

- Bach, K. (2005). The emperor's new 'knows'. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: knowledge, meaning, and truth* (pp. 51–89). Oxford: Oxford University Press.
- Buckwalter, W. (2010). Knowledge isn't closed on Saturday: A study in ordinary language. *Review of Philosophy and Psychology*, 1(3), 395–406.
- Buckwalter, W. (2014). The mystery of stakes and error in ascriber intuitions. In J. R. Beebe (Ed.), *Advances in experimental epistemology* (pp. 145–173). London: Bloomsbury Academic.
- Buckwalter, W. (2014). Factive verbs and protagonist projection. *Episteme*, 11, 391-409.
- Buckwalter, W., Rose, D., & Turri, J. (2015). Belief through thick and thin. *Nous*, 49, 748-775.
- Buckwalter, W. & Schaffer, J. (2015). Knowledge, stakes and mistakes. *Nous*, 201-234.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
- Cohen, S. (2013). Contextualism defended. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 69–75). Malden, Mass.: Wiley-Blackwell.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913–29.
- DeRose, K. (2009). *The case for contextualism*. Oxford: Oxford University Press.
- DeRose, K. (2011). Contextualism, contrastivism, and X-Phi surveys. *Philosophical studies*, 156, 81-110.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford University Press.
- Feltz, A., & Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical Psychology*, 23(5), 683–706.
- Ganson, D. (2007). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139, 441–458
- Holton, R. (1997). Some telling examples: A reply to Tsohatzidis. *Journal of Pragmatics*, 28, 624–628.
- Lawlor, K. (2013). *Assurance: an Austinian view of knowledge and knowledge claims*. Oxford: Oxford University Press.
- Ludlow, P. (2005). Contextualism and the new linguistic turn in epistemology. In G. Preyer and G. Peter, (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (pp. 11–50). Oxford: Oxford University Press.

- Machery, E., Stich, S. P., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., Sirker, S., Usui, N., & Hashimoto, T. (2015). Gettier across cultures. *Nous* doi: 10.1111/nous.12110.
- Machery, E., Stich, S. P., Rose, D., Chatterjee, A., Karasawa, K., Struchiner, N., Sirker, S., Usui, N., & Hashimoto, T. (2015). Gettier was framed! MS
- May, J., Sinnott-Armstrong, W., Hull, J. G., & Zimmerman, A. (2010). Practical interests, relevant alternatives, and knowledge attributions: an empirical study. *Review of Philosophy and Psychology*, 1(2), 265–273.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2), 279–294.
- Pinillos, Á. (2012). Knowledge, Experiments and Practical Interests. In (eds.) Brown & Gerken, *Knowledge Ascriptions* (pp. 192–221). Oxford University Press.
- Pinillos, Á., & Simpson, S. (2014). Experimental evidence supporting anti-intellectualism about knowledge. In J. Beebe (Ed.), *Advances in experimental epistemology* (pp. 9-43). London: Bloomsbury.
- Rose, D. (2015). Belief is prior to knowledge. *Episteme*, 12, 385-399.
- Rose, D. & Schaffer, J. (2013). Knowledge entails dispositional belief. *Philosophical Studies*, 166, 19-50.
- Rysiew, P. (2001). The context-sensitivity of knowledge attributions. *Nous*, 35, 477–514.
- Schaffer, J. & Knobe, J. (2009). Contrastivism surveyed. *Nous*, 46, 675-708.
- Sripada, C., & Stanley, J. (2012). Empirical Tests of Interest-Relative Invariantism. *Episteme* 9, 3–26.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stich, S. (2013). Do different groups have different epistemic intuitions? A Reply to Nagel. *Philosophy and Phenomenological Research*, 87, 151-178.
- Turri, J. & Buckwalter, W. (forthcoming). Descartes' schism, Locke's reunion: Completing the pragmatic turn in epistemology. *American Philosophical Quarterly*.
- Turri, J. (forthcoming). Epistemic contextualism: An idle hypothesis. *Australasian Journal of Philosophy*.
- Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19, 417–443.
- Weinberg, J. M., Nichols, S., & Stich, S. (2001). Normativity and epistemic intuitions. *Philosophical Topics*, 29, 429-460.