

EZRI CHERNAK, KURT DIETRICH, ASHLEY RASPOPOVIC,
SARAH TURRI, JOHN TURRI*

LYING BY OMISSION EXPERIMENTAL STUDIES**

Abstract

Leading theories of lying disagree on many points, but they agree on the following assumption: lying essentially involves asserting. The possibility of *lying by omission* poses a challenge to that shared assumption. To lie by omission is to lie by *not* asserting. This paper is the first experimental investigation of whether lying by omission is conceptually possible, according to our ordinary, shared lying concept. Overall, our results support, without proving, that it is not possible. Based on the present findings, we hypothesize that to the extent that people are tempted to call an omission a “lie,” it is for lack of a better word. When provided more flexibility to express themselves, almost no one in our studies counted an omission as a lie.

Keywords: lying, assertion, communication, omission, signaling

One standard definition of lying is that lying is a dishonest statement intended to deceive an audience. This definition is popular in both philosophical

* Ezri Chernak: Philosophy Department, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L3G1, Canada.

Kurt Dietrich: Mathematics Department, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L3G1, Canada.

Ashley Raspopovic: Philosophy Department, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L3G1, Canada.

Sarah Turri: Independent Scholar.

John Turri: Philosophy Department, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L3G1, Canada; Cognitive Science Program, University of Waterloo, corresponding author: john.turri@gmail.com, ORCID: <https://orcid.org/0000-0001-7611-2824>.

** Data and materials available here: <https://osf.io/a6ehs>. For helpful feedback and discussion, we thank members of Philosophy 251 (Winter 2019) at the University of Waterloo, and Angelo and Geno Turri. This research was supported by the Canada Research Chairs Program and the Social Sciences and Humanities Research Council of Canada.

and social scientific research on lying (Mahon 2016; Vrij 2008). Understood as an account of our shared ordinary lying concept, it has been challenged on multiple grounds. For example, according to this account, falsity is not essential to lying. But convergent evidence from experimental cognitive science and philosophical logic shows that falsity is essential to lying (Strichartz and Burton 1990; Turri and Turri 2015, 2021). According to this account, deceptive intent is also essential to lying, but results from recent studies provide some initial evidence against this (Arico and Fallis 2013; Turri and Turri 2016: 34ff). An alternative hypothesis is that lying is asserting something you know is false (Turri 2016: 34ff; Benton 2018). According to the rules of our social practice of assertion, assertions should express knowledge (Turri 2016, 2017; Turri and Park 2018). And to cheat is to knowingly break the rules. Therefore, on this alternative account, lying is akin to cheating at assertion.

The standard definition and the alternative knowledge-based account both view assertion as essential to lying. A lie is an assertion meeting further conditions. Only a couple studies provide evidence relevant to assessing this claim. One study found that lie attributions are affected by the audience's ability to hear an agent's spoken words, and that this effect was mediated by judgments about whether the agent made an assertion (Turri and Turri 2016, experiment 1). A second study found that lie attributions and assertion attributions were significantly correlated (Turri and Turri 2016, experiment 2).

To assert is, roughly, to send an informational signal by an established convention. The prototypical assertion involves uttering a declarative sentence. But not all assertions are vocalized. Many are written on paper, typed on screen, or enacted in gesture. Humans are creative in establishing signaling conventions. Puffs of smoke, beeps, coughs, winks, posture, and the proportionate duration and frequency of telegraphic taps are just a few examples. Even the absence of a tangible signal can itself be a signal.

The "warrant canary" is a vivid example. Governments sometimes use secret orders to obtain people's information from service providers. Service providers are legally prohibited from telling anyone that they received the order. Many people view this as illegitimate and devised a legal workaround. Here is one way this might go. Every day the provider does not receive a secret order, it publicly states that it has not received one. On days when it receives one, it says nothing about secret orders. Everyone paying attention is thereby informed that the provider received a secret order today. Hence, silence itself becomes an informational signal.¹

¹ For an example of a warrant canary, see <https://puri.sm/warrant-canary/>.

The possibility of lying by omission challenges the view that assertion is essential to lying. Sending a signal containing specific information is to commit rather than omit this information. In other words, omitting information requires not asserting it. Accordingly, to lie by omission would be to lie by not asserting. As “warrant canaries” illustrate, lying by omission is not merely remaining silent. Silence can be an “agreed upon signal with others that is equivalent to making a statement” (Mahon 2016).

In this paper, we begin experimentally investigating the possibility of lying by omission. Across five experiments, we varied whether an agent withholds relevant information (omission) or asserts false information (commission), whether an agent agrees or refuses to be responsible for providing information, the perceived stakes of the situation, the conversational context, the topic under discussion, and how participants recorded lie attributions.

If lying by omission is conceptually possible, then in a case that clearly involves withholding relevant, requested information, people will tend to attribute a lie. The stronger the central tendency, the stronger the evidence. Perhaps the most convincing evidence would be to detect a strong central tendency to count an omission as a lie and to detect no difference between the omission condition and a closely matched control condition involving a commission. In short, if lying by omission is possible, then we should observe high lie attribution for omissions and, ideally, detect no difference from commission.

By contrast, if lying by omission is not possible, then even in a case that clearly involves withholding relevant, requested information, people will withhold lie attribution. The stronger the central tendency, the stronger the evidence. Of course, people might withhold lie attribution for many reasons. Accordingly, in order for the evidence to be even minimally persuasive, people must also attribute a lie in a closely matched control condition involving a commission. In short, if lying by omission is not possible, then we should observe low lie attribution for omission and high attribution for commission.

To preview the results, in Experiment 1, we found that lie attributions were highly sensitive to the difference between commission and omission. The strong central tendency was to count the commission as a lie. But lie attribution for omission was at chance rate. Accordingly, the results were inconclusive. In Experiment 2, we followed up on the possibility that lie attribution for the omission was inflated because some people interpreted the agent’s silence as a signal that had been antecedently agreed upon, thus failing to view the case as a genuine omission. We found that lie attributions were highly sensitive to the difference between the agent agreeing or refusing to be responsible for the relevant information. When the agent agreed to be responsible but then remained silent, lie attribution was high. By contrast,

when the agent refused to be responsible and then remained silent, lie attribution again went to chance. In Experiment 3, we followed up on the possibility that lie attribution for omissions was inflated by the earlier questioning procedures. The basic idea is that people might use an isolated “yes/no” attribution to convey something other than their judgment about whether the agent lied. Using a more flexible check-all-that-apply task, we found strong initial evidence that omissions are not counted as lies, whereas commissions are. In Experiment 4, we followed up on the possibility that lying by omission would be attributed only when the stakes were high. Stakes did not affect lie attributions, and people denied lying by omission at similarly high rates for both low and high stakes. In Experiment 5, we generalized this result for a range of scenarios that varied the conversational context and topic under discussion. We continued to observe a central tendency to deny lying for omissions but attribute it for closely matched commissions.

To forestall potential confusion, we emphasize that we do not intend our use of “omission” to closely track the use of that term in technical philosophical disputes regarding, for instance, causation, responsibility, or agency (e.g., Clarke 2014; Sartorio 2009). Nor do we intend to leverage the present results by applying them to such disputes. Instead, we limit ourselves to the sense of “omission” relevant to central theoretical questions regarding lying specifically. In this sense, the omission of interest consists of “omitting to make a statement” (Mahon 2016: section 1.1), or more generally, omitting to send an informational signal. We could even avoid the word “omission” and its cognates entirely and instead speak of *withholding information*, in contrast to providing information.

1. GENERAL METHODS

All manipulations, measures, and exclusion criteria are reported. All participants were adult residents of the United States. We recruited and tested people using an online platform of Amazon Mechanical Turk (<https://www.mturk.com>), TurkPrime (Litman, Robinson, and Abberbock 2017), and Qualtrics (<https://www.qualtrics.com>). We decided in advance to recruit fifty participants per condition, plus some extra as a precaution against attrition. Participants completed a brief demographic questionnaire after testing. We used R 3.6.3 for all analyses (R Core Team 2018). All stimuli, data, and code are available through an Open Science Foundation project (osf.io/a6ehs). All studies were pre-registered.

2. EXPERIMENT 1

METHOD

PARTICIPANTS

One hundred five people participated in the study. Their mean age was 35.69 years (range = 22-72, $sd = 10.25$), 30% (32 of 105) were female, and 93.3% reported native competence in English.

MATERIALS AND PROCEDURE

Participants were randomly assigned to one of two conditions (omission, commission), read a brief scenario, and responded to a lie attribution while the scenario remained visible. The scenario differed in whether the agent withheld information (omission) or provided false information (commission). Response options were randomly rotated.

Mary and Bob co-own a painting supplies store. Mary asks Bob to let her know when they are running short on lime-green paint so they can buy more. However, Bob does not want to buy more lime-green paint once their current supply runs out. When Bob sees they are running low on lime-green paint, [he does not inform Mary of it / he tells Mary that they have plenty].

Did Bob lie to Mary? [No/Yes]

RESULTS

Our principal research question was whether “lies of omission” are counted as lies. To answer this question, we conducted proportion tests on lie attributions, comparing rates across conditions and against chance (test proportion = .5). A Fisher’s exact test revealed that assignment to condition affected lie attributions, $p < .001$ (see Table 1 and Figure 1). Binomial tests revealed that attribution was significantly above chance rate in the commission condition (see Table 1). By contrast, attribution did not differ from chance in the omission condition (see Table 1).

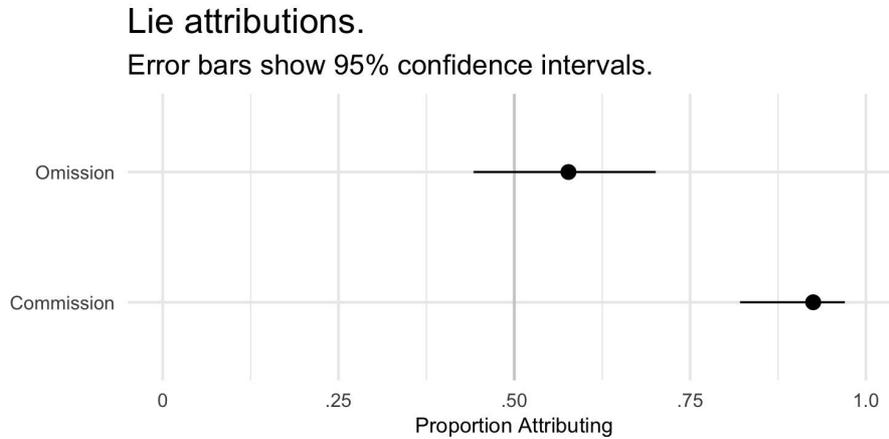


Figure 1. Experiment 1. Proportion of participants attributing a lie in the two conditions

Condition	n	k	prop	conf.low	conf.high	p	test.prop	h
Omission	52	30	.577	.442	0.701	.332	.5	0.154
Commission	53	49	.925	.821	0.970	<.001	.5	1.014

Table 1. Experiment 1. Descriptive statistics and binomial tests

DISCUSSION

This experiment asked whether the difference between withholding relevant information (omission) and providing false information (commission) affects lie attributions. We observed a large significant effect. We also found that lie attribution significantly exceeded chance rates for a commission (92%), whereas it did not differ from chance for a closely matched omission (58%). These results support the conclusion that lie attributions are partly based on whether the speaker made an assertion. But the chance rate of attribution for the omission makes it difficult to interpret the results as either supporting or undermining the hypothesis that assertion is essential to lying. The large effect of condition suggests that assertion is essential to lying. But if assertion were essential to lying, then one might expect the central tendency to be to deny that an omission was a lie.

The next two studies each explore a possible explanation for the chance rate of attribution observed here.

3. EXPERIMENT 2

One concern about the results of Experiment 1 is that people might disagree about the lie attribution because they make different inferences about the scenario. In particular, people might differently interpret the communicative exchange's background assumptions. Suppose people interpreted the speaker as agreeing, perhaps implicitly, to provide information under specified circumstances. Accordingly, if the speaker does not say anything, then people could interpret silence as sending a signal by an agreed convention, effectively telling the recipient something. Participants reasoning this way count a non-vocalized assertion as a lie. This does not constitute a "lie of omission" in the relevant sense. Better evidence regarding "lies of omission" would address this possibility, which is what the present experiment investigates.

METHOD

PARTICIPANTS

One hundred five people participated in the study. Their mean age was 33.51 years (range = 19-67, $sd = 9.69$), 35% (37 of 105) were female, and 92.4% reported native competence in English.

MATERIALS AND PROCEDURE

Participants were randomly assigned to one of two conditions (refuse, agree), read a brief scenario, and responded to a lie attribution while the scenario remained visible. The scenario differed in whether Bob refused or agreed to be responsible for providing requested information. Response options were randomly rotated.

Mary and Bob co-own a painting supplies store. Mary asks Bob to let her know when they are running short on lime-green paint so they can buy more. Bob explicitly [refuses/agrees] to be responsible for providing Mary with this information.

Bob does not want to buy more lime-green paint once their current supply runs out. When Bob sees they are running low on lime-green paint, he does not inform Mary of it.

Did Bob lie to Mary? [No/Yes]

RESULTS

Our principal research question was whether the refuse/agree manipulation would affect lie attributions, with attribution higher in the agree condition. To answer this question, we conducted proportion tests on lie attributions, comparing rates across conditions and against chance (test proportion = .5). A Fisher's exact test revealed that assignment to condition affected lie attributions, $p < .001$ (see Table 2 and Figure 2). Binomial tests revealed that attribution was significantly above the chance rate in the agree condition (see Table 2). But attribution did not differ from chance in the refuse condition (see Table 2).

Lie attributions.

Error bars show 95% confidence intervals.

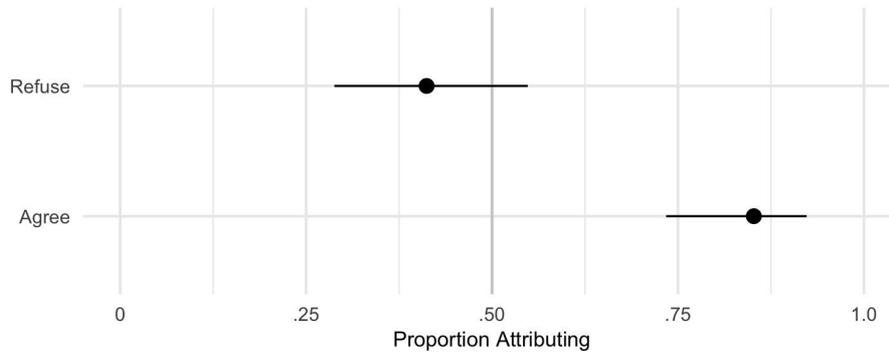


Figure 2. Experiment 2. Proportion of participants attributing a lie in the two conditions

Condition	n	k	prop	conf.low	conf.high	p	test.prop	h
Refuse	51	21	.412	.288	0.548	.262	.5	-0.177
Agree	54	46	.852	.734	0.923	<.001	.5	0.781

Table 2. Experiment 2. Descriptive statistics and binomial tests

DISCUSSION

This experiment addressed the concern that some participants in Experiment 1 interpreted silence as sending a signal by an agreed convention, rather than an omission. We manipulated whether the speaker explicitly refused or agreed to be responsible for providing information. When the speaker agreed,

most participants counted silence as a lie (85%). By contrast, when the speaker refused, lie attributions again did not differ from chance (41%). These findings again fail to provide clear evidence for or against the possibility of lying by omission. Ultimately, we conclude that the present study failed to clarify the interpretation of Experiment 1.

There are at least two possible explanations of the effect of condition on lie attributions. One explanation is that when Bob agreed to be responsible for the information, participants interpreted his later silence as sending a signal by an agreed convention. This explanation coheres with the view that assertion is essential to lying. An alternative explanation is that there were more candidates for a lie in the agree condition than in the refuse condition. When Bob refused to provide the information, there was only one candidate for a lie, a lie by omission, enacted by withholding the information about the supplies. But when Bob agreed to provide the information, there were multiple candidates for a lie: the same lie by omission, but also a potential lie by commission, enacted by agreeing in bad faith to provide the information. The experimental design does not allow us to confidently choose between the two explanations. Moreover, the explanations are consistent with each other and could each capture part of the truth.

4. EXPERIMENT 3

Another concern about the results of Experiment 1 is that the questioning procedures might have artificially inflated lie attributions. When Bob withholds relevant information, and participants rate a “yes/no” lie attribution in isolation, some might answer “yes” despite thinking that he did not lie, because “yes” comes closer to conveying something salient to them. For example, they might use the lie attribution to express disapproval of Bob’s conduct. To address this, the present experiment used a different questioning procedure.

METHOD

PARTICIPANTS

One hundred five people participated in the study. Their mean age was 34.05 years (range = 18-71, $sd = 10.73$), 34% (36 of 105) were female, and 92.4% reported native competence in English.

MATERIALS AND PROCEDURE

Participants were randomly assigned to one of two conditions (omission, commission), read a brief scenario, and completed a check-all-that-apply task while the scenario remained visible atop. The scenario and manipulation were taken verbatim from Experiment 1. Here is the check-all-that-apply task (option order rotated randomly):

In what follows, please select all that apply to Bob. You can select all options, none, or a mix.

Bob:

- withheld information about whether they have enough (withheld)
- told Mary that they have enough (told)
- lied to Mary about whether they have enough (lied)
- misled Mary about whether they have enough (misled)

The option labels are for expository convenience; participants did not see them.

RESULTS

Our principal research question was whether lie attributions will be high for commission but low for omission, using the current questioning procedures. To answer this question, we conducted proportion tests on lie attributions, comparing rates across conditions and against chance (test proportion = .5). Separate Fisher's exact tests revealed that assignment to condition affected all four dependent measures (see Table 3 and Figure 3). Binomial tests revealed that lie attribution exceeded chance for commission (see Table 4 and Figure 3). By contrast, lie attribution was significantly below chance for omission.

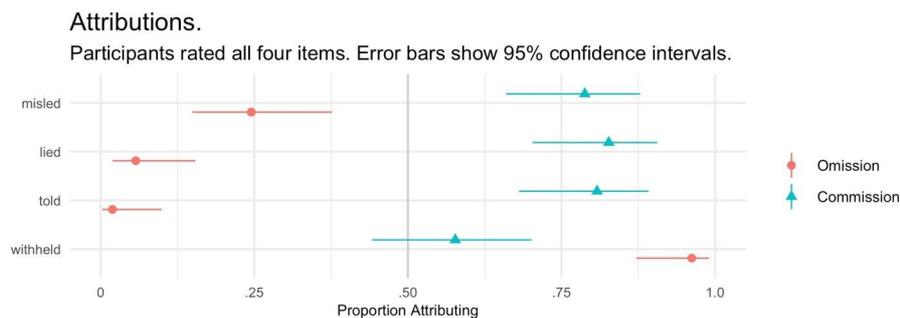


Figure 3. Experiment 3. Proportion of participants attributing four statuses in the two conditions

Condition	Judgment	n	k	prop	conf.low	conf.high	p	test.prop	h
Omission	withheld	53	51	.962	.872	.990	<.001	.5	1.180
Omission	told	53	1	.019	.003	.099	<.001	.5	-1.295
Omission	lied	53	3	.057	.019	.154	<.001	.5	-1.090
Omission	misled	53	13	.245	.149	.376	<.001	.5	-0.535
Commission	withheld	52	30	.577	.442	.701	.332	.5	0.154
Commission	told	52	42	.808	.681	.892	<.001	.5	0.663
Commission	lied	52	43	.827	.703	.906	<.001	.5	0.713
Commission	misled	52	41	.788	.660	.878	<.001	.5	0.615

Table 3. Experiment 3. Descriptive statistics and binomial tests

Judgment	odds.ratio	conf.low	conf.high	p
withheld	0.055	0.006	0.248	.001
told	200.709	28.111	8417.601	.001
lied	73.830	18.308	449.102	.001
misled	11.135	4.229	31.794	.001

Table 4. Experiment 3. Fisher's exact tests for the four dependent measures

DISCUSSION

One concern about the results of Experiment 1 is that an isolated “yes/no” lie attribution might have misleadingly inflated lie attributions. The present experiment addressed this concern by using a check-all-that-apply task to provide more flexibility to participants. If the concern is well founded, then lie attribution should be low for the omission while remaining high for the commission. The results are consistent with this prediction. When Bob provided false information (commission), most participants attributed to him a lie (83%). By contrast, when Bob withheld information (omission), nearly no participants attributed to him a lie (6%). These findings inform the interpretation of Experiment 1, which did not provide clear evidence for or against the possibility of lying by omission. By the same token, the present findings also provide initial evidence that omissions do not count as lies in ordinary social cognition.

5. EXPERIMENT 4

One concern about the results of Experiment 3 is that we missed the opportunity to detect lie attribution for omissions. In particular, it is theoretically possible that lying by omission can occur only in combination with other factors. For example, the case we tested does not seem to involve “high stakes,” but perhaps that is the missing ingredient. To address this, the present experiment manipulated the potential costs of failing to provide accurate information.

METHOD

PARTICIPANTS

Two hundred ten people participated in the study. Their mean age was 37.19 years (range = 18-71, $sd = 12.4$), 47% (98 of 210) were female, and 91% reported native competence in English.

MATERIALS AND PROCEDURE

Participants were randomly assigned to one of four conditions in a 2 x 2 between-subjects design (Action: omission, commission; Stakes: low, high). Participants read a brief scenario and completed a check-all-that-apply task while the scenario remained visible. Here is how the scenario began for all participants, with the low/high manipulation bracketed:

Mary and Bob co-own a hardware store. Mary asks Bob to let her know if their stock of paint [begins running low, so that together they can restock the shelves / tests positive for toxins, so that together they can pull it from the shelves]. However, Bob does not want to go through the trouble. When Bob sees that the [stock of paint is running low / paint tested positive for toxins].

In omission conditions, the scenario then ended with, “he simply remains silent and does not inform Mary of it.” In commission conditions, the scenario ended with Bob providing Mary with salient false information: “he tells Mary that [they have plenty of paint left in stock / the test results show the paint is safe].” Participants then completed a check-all-that-apply task, using the same instructions from Experiment 3, and with details that differed appropriately between the low and high stakes conditions:

Bob:

- withheld information about whether [they have enough paint / the paint is safe]

- told Mary that [they have enough paint / the paint is safe]
- lied to Mary about whether [they have enough paint / the paint is safe]
- misled Mary about whether [they have enough paint / the paint is safe]

Participants then went to a new screen and responded to a manipulation check:

Please rate your agreement or disagreement with the following statement:

The situation with the paint is potentially very serious.

Responses were collected on a standard 7-point Likert scale, 1 (strongly disagree) – 7 (strongly agree), arranged vertically on the participant's screen. The purpose of including this was to verify that our stakes manipulation was effective.

RESULTS

The stakes manipulation was effective (see Table 5). Our principal research question concerned the potential effect of stakes on lie attribution. To address this question, we constructed a generalized linear model predicting lie attribution with Action, Stakes, and participant age and sex as predictors, followed up by appropriate proportion tests. The generalized linear model revealed a main effect of Action only (see Table 6 and Figure 4). Binomial tests revealed that lie attribution exceeded chance in commission conditions (see Table 6, Figure 4). But lie attribution was significantly below chance in omission conditions.

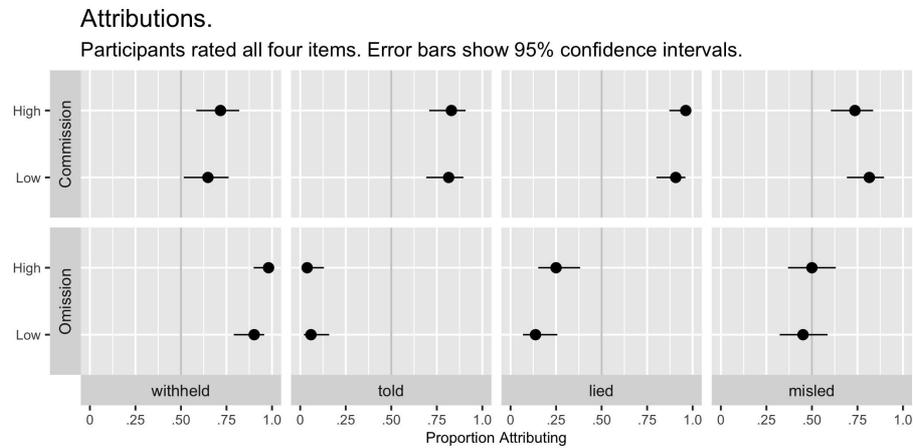


Figure 4. Experiment 4. Proportion of participants attributing four statuses in the four conditions

Low	High	MD	t	df	p.value	d
4.85	6.32	-1.47	-7.75	185.37	<.001	-1.07

Table 5. Experiment 4. Independent samples t-test comparing seriousness ratings in the low/high conditions

Term	estimate	std.error	z	p.value
(Intercept)	-1.853	0.743	-2.496	.013
Commission	4.122	0.623	6.620	<.001
High	0.740	0.518	1.427	.153
Female	0.216	0.431	0.500	.617
Age	-0.002	0.017	-0.139	.889
Commission: High	0.220	1.006	0.218	.827

Table 6. Experiment 4. Generalized linear model predicting lie attributions. Reference class for Commission is Omission. Reference class for High is Low. Reference class for female is male.

Action	Stakes	Judgment	n	k	prop	conf.low	conf.high	p	h
Omission	Low	lied	51	7	.137	.068	.257	<.001	-0.812
Omission	High	lied	52	13	.250	.152	.382	<.001	-0.524
Commission	Low	lied	54	49	.907	.801	.960	<.001	0.952
Commission	High	lied	53	51	.962	.872	.990	<.001	1.180

Table 7. Experiment 4. Descriptive statistics and binomial tests (test proportion = .5)

DISCUSSION

One concern about the earlier results is that they failed to detect high lie attribution for omissions because the perceived stakes were low. This experiment addressed the concern by manipulating stakes and testing for an effect. The results undermine the concern. Stakes did not significantly affect lie attribution. Lie attributions remained high for commissions but low for omissions.

6. EXPERIMENT 5

It is possible that the earlier results were due to specific features of the basic scenario we tested. For example, they could have been due to the conversational context or topic under discussion, which might interact in a very specific and unexpected way with the difference between withholding relevant information and providing false information. To assess whether the findings generalize, this experiment tests a range of different scenarios.

METHOD

PARTICIPANTS

Six hundred sixteen people participated in the study. Their mean age was 37.39 years (range = 0-83, *sd* = 12.33), 44% (274 of 616) were female, and 92.4% reported native competence in English.

MATERIALS AND PROCEDURE

Participants were randomly assigned to one of twelve conditions in a 2 (Action: omission, commission) x 6 (Scenario) between-subjects design. Participants read a brief scenario and completed a check-all-that-apply task while

the scenario remained visible atop. The options were the same as in Experiments 3 and 4 (withheld, told, lied, misled) (order randomized). The Action manipulation was the same as from Experiments 3 and 4. The Scenario factor was not of independent theoretical interest and was included to support generalization of the results beyond the specific stimuli studied here (Clark 1973; Baayen, Davidson, and Bates 2008; Judd, Westfall, and Kenny 2012). The scenarios all involved Mary and Bob conversing, as in previous studies, but the scenarios varied the conversation's context and topic. All the stimuli are available through the OSF project for this paper (osf.io/a6ehs).

RESULTS

Our principal research question was whether the previously observed pattern for lie attributions (low for omission, high for commission) is robust across scenarios. To address this question, we constructed a generalized linear mixed model predicting lie attribution. The model included fixed effects for Action and participant age and sex, and it included a random intercept for Scenario. The model revealed a main effect of Action and the same basic pattern is observed across all six scenarios (see Table 8 and Figure 5). Lie attribution exceeded chance in commission conditions (.867) (see Table 9). By contrast, lie attribution was below chance in omission conditions (.15).

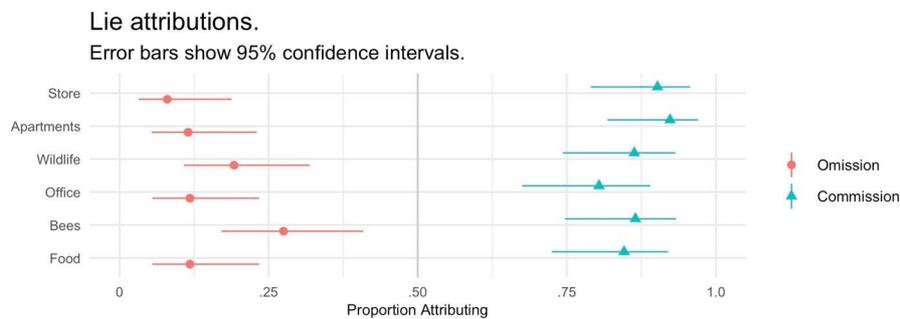


Figure 5. Experiment 5. Proportion of participants attributing four statuses in the four conditions

Term	estimate	std.error	z	p.value
(Intercept)	-1.764	0.389	-4.532	<.001
Commission	3.616	0.235	15.405	<.001
Female	-0.007	0.236	-0.030	.976
Age	0.001	0.010	0.085	.932

Table 8. Experiment 5. Fixed effects of the generalized linear mixed model predicting lie attributions

Action	n	k	prop	conf.low	conf.high	p	test.prop	h
Omission	307	46	.150	.114	.194	<.001	.5	-0.776
Commission	309	268	.867	.825	.901	<.001	.5	0.825

Table 9. Experiment 5. Descriptive statistics and binomial tests for lie attributions in the two Action conditions

DISCUSSION

One concern about the earlier results is that they could have been influenced by the peculiarities of the basic scenario tested. This experiment addressed the concern by testing a range of different scenarios. Across all six scenarios tested, lie attributions remained high for commission but low for omission.

7. GENERAL DISCUSSION

Results from five pre-registered studies advance our understanding of the ordinary lying concept. Leading theories of lying disagree on many things, such as whether lies are essentially false or intended to deceive. But, as we saw, one point of agreement is that lying essentially involves asserting, or sending a signal. Call this *the assertion hypothesis*. Two previous studies provide some correlational evidence supporting the assertion hypothesis, but firm conclusions are not yet warranted. An open challenge to the assertion hypothesis is the possibility of “lying by omission,” or lying by not asserting. Our goal here was to begin assessing this challenge in earnest.

In Experiment 1, we manipulated whether Bob withholds information (omission) or provides false information (commission). In line with the assertion hypothesis, switching from omission to commission had a large effect on lie attributions. Also in line with the assertion hypothesis, lie attribution was high for the commission. The assertion hypothesis also predicts low lie attribution for the omission, but we instead observed a chance rate. So absent a plausible alternative explanation for this discrepancy, the results do not clearly support the assertion hypothesis.

In Experiments 2 and 3, we began assessing alternative explanations for the discrepancy. Experiment 2 pursued one alternative based on the idea that some participants might have interpreted the silence as a signal that had been antecedently agreed upon. But the results were once again inconclusive, neither clearly supporting nor undermining the assertion hypothesis. Experiment 3 pursued another alternative based on the idea that the earlier questioning procedures led to task substitution — that is, the phenomenon whereby participants use a test item to perform a different task from the one requested. The concern was that some participants might answer “yes” to an isolated “yes/no” lie attribution in order to express their recognition that the agent’s communicative behavior was uncooperative or otherwise inappropriate. To address this, we used a more flexible check-all-that-apply task. In line with the assertion hypothesis, lie attribution was low for omission and high for commission.

In Experiment 4, we examined whether raising the stakes affected whether omissions are counted as lies. We manipulated the potential consequences of failing to provide accurate information. Once again, in line with the assertion hypothesis, lie attribution was low for omission and high for commission. Stakes did not significantly affect the rate of lie attribution.

In Experiment 5, we examined whether the critical pattern of lie attributions being low for omissions but high for commissions, observed in earlier studies, was due to contingent features of the conversational context or topic under discussion. Across six different scenarios, we observed the same basic pattern.

Overall, we interpret these results as strong initial evidence for the assertion hypothesis. In light of our findings, a reasonable conjecture is that when people are tempted to call an omission a “lie,” it is for lack of a better word.

The findings do not conclusively prove the assertion hypothesis: we cannot rule out that the range of scenarios we tested omitted a crucial ingredient. To be sure, the scenarios we tested are simple, straightforward, and seemingly well suited to capture the phenomenon of lying by omission. We gave lying by omission a fair shot. It is possible that other scenarios give it an even

better shot, but we submit that if lying by omission is possible, then it would be surprising if it did not show up in the cases we tested.

Regardless of whether this limitation turns out to be critical, our findings can help guide future research. In particular, the most convincing studies will take seriously the possibility of task substitution (and, of course, include proper omission/commission control comparisons). If researchers observe high lie attribution for omissions using an isolated “yes/no” test item, the next step should be to check whether results replicate when using more flexible questioning procedures, such as those we used in Experiments 3-5. Convincing studies will also take seriously the possibility of silence being interpreted as a non-vocalized assertion. In retrospect, we cannot endorse our way of handling this possibility in our inconclusive Experiment 2. Fortunately, and again in retrospect, resolving these problems in future studies will not be difficult.

BIBLIOGRAPHY

- Arico A. J., Fallis D. (2013), “Lies, Damned Lies, and Statistics: An Empirical Investigation of the Concept of Lying,” *Philosophical Psychology* 26(6), 790-816. <https://doi.org/10.1080/09515089.2012.725977>
- Baayen R. H., Davidson D. J., Bates D. M. (2008), “Mixed-Effects Modeling with Crossed Random Effects for Subjects and Items,” *Journal of Memory and Language* 59(4), 390-412. <https://doi.org/10.1016/j.jml.2007.12.005>
- Benton M. A. (2018), “Lying, Belief, and Knowledge” [in:] *The Oxford Handbook of Lying*, J. Meibauer (ed.), Oxford: Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780198736578.013.9>
- Clark H. H. (1973), “The Language-as-Fixed-Effect Fallacy: A Critique of Language Statistics in Psychological Research,” *Journal of Verbal Learning and Verbal Behavior* 12, 335-359. [https://doi.org/10.1016/S0022-5371\(73\)80014-3](https://doi.org/10.1016/S0022-5371(73)80014-3)
- Clarke R. (2014), *Omissions: Agency, Metaphysics and Responsibility*, Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199347520.001.0001>
- Judd C. M., Westfall J., Kenny D. A. (2012), “Treating Stimuli as a Random Factor in Social Psychology: A New and Comprehensive Solution to a Pervasive but Largely Ignored Problem,” *Journal of Personality and Social Psychology* 103(1), 54-69. <https://doi.org/10.1037/a0028347>
- Litman L., Robinson J., Abberbock T. (2017), “TurkPrime.com: A Versatile Crowdsourcing Data Acquisition Platform for the Behavioral Sciences,” *Behavioral Research Methods* 49(2), 1-10. <https://doi.org/10.3758/s13428-016-0727-z>
- Mahon J. E. (2016), “The Definition of Lying and Deception” [in:] *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/win2016/entries/lying-definition/>.
- R Core Team (2018), *R: A Language and Environment for Statistical Computing*, Vienna: R Foundation for Statistical Computing.

- Sartorio C. (2009), "Omissions and Causalism," *Noûs* 43(3), 513-530. <https://doi.org/10.1111/j.1468-0068.2009.00716.x>
- Strichartz A. F., Burton R. V. (1990), "Lies and Truth: A Study of the Development of the Concept," *Child Development* 61(1), 211-220. <https://doi.org/10.2307/1131060>
- Turri A., Turri J. (2015), "The Truth about Lying," *Cognition* 138(C), 161-168. <https://doi.org/10.1016/j.cognition.2015.01.007>
- Turri A., Turri J. (2016), "Lying, Uptake, Assertion, and Intent," *International Review of Pragmatics* 8(2), 314-333. <https://doi.org/10.1163/18773109-00802006>
- Turri A., Turri J. (2021), "Lying, Fast and Slow," *Synthese* 198, 757-775. <https://doi.org/10.1007/s11229-018-02062-z>
- Turri J. (2016), *Knowledge and the Norm of Assertion: An Essay in Philosophical Science*, Cambridge: Open Book Publishers. <https://doi.org/10.11647/OBP.0083>
- Turri J. (2017), "Experimental Work on the Norms of Assertion," *Philosophy Compass* 12(7), e12425. <https://doi.org/10.1111/phc3.12425>
- Turri J., Park Y. (2018), "Knowledge and Assertion in Korean," *Cognitive Science* 42(6), 2060-2080. <https://doi.org/10.1111/cogs.12621>
- Vrij A. (2008), *Detecting Lies and Deceit: Pitfalls and Opportunities*, 2nd ed., Hoboken, NJ: John Wiley & Sons.