Lying, Uptake, Assertion, and Intent*

Angelo Turri
angelo.turri@gmail.com
John Turri
john.turri@gmail.com

Abstract: A standard view in social science and philosophy is that a lie is a dishonest assertion: to lie is to assert something that you think is false in order to deceive your audience. We report four behavioral experiments designed to evaluate some aspects of this view. Participants read short scenarios and judged several features of interest, including whether an agent lied. We found evidence that ordinary lie attributions can be influenced by aspects of audience uptake, are based on judging that the agent made an assertion (assertion attributions), and, at least in some contexts, are not based on attributions of deceptive intent. The finding on assertion attributions is predicted by the standard view, but the finding on intent attributions is not. These results help to further clarify the ordinary concept of lying and shed light on the psychological processes involved in ordinary lie attributions and related judgments.

Keywords: lying, deception, assertion; social cognition; regression; mediation; causal search

Word count: 5088

Introduction

A standard view in social science and philosophy is that a lie is a dishonest assertion (e.g. Augustine, 395; Aquinas, 1273; Vrij, 2008). On this view, lying is defined as asserting something that you think is false in order to deceive your audience. Otherwise put, making a dishonest assertion

* This is the penultimate version of a paper to appear in *International Review of Pragmatics*. Please cite the final, published version if possible.

is necessary and sufficient for lying.

Recently this view has been challenged on several fronts, at least when understood as an attempt to characterize the ordinary, shared concept of lying. First, results from several studies suggest that the ordinary concept of lying is not properly defined in terms of a "checklist" of necessary and sufficient conditions but instead is a prototype concept (Coleman & Kay, 1981; Lee & Ross, 1997). According to this view, assertions are judged to be lies based on how closely they resemble a prototypical lie, with no particular feature being necessary for lying. Second, results from other studies suggest that objective falsity is centrally important to the ordinary concept of lying (Strichartz & Burton, 1990; Turri & Turri, 2015). Some of these results suggest that when someone makes an assertion she mistakenly believes is false, she did not lie but, rather, *only thinks* she lied. Third, results from another study arguably suggest that deceptive intent is inessential to the ordinary concept of lying (Arico & Fallis, 2013). Finally, some have argued that in addition to semantic properties, the ordinary concept of lying is partly constituted by social and contextual factors, such as the perceived purpose of a conversational exchange (Sweetser, 1987). For instance, people are more willing to classify an assertion as a lie when it is spoken in a context where information-exchange, rather than politeness, is the principal goal (Lee & Ross, 1997).

In this paper, we report a new series of experiments designed to advance knowledge of the ordinary concept of lying, in two broad ways. First, we studied the potential role of *audience uptake* on people's willingness to categorize an assertion as a lie. More specifically, we studied whether lie attributions are affected by an audience's ability to hear what the speaker said, an au-

dience's ability to comprehend the speaker's language, and even whether there is an audience at all. Second, we studied whether people's lie attributions are based on their attributions of assertion and deceptive intent. More specifically, we used mediation analysis to determine whether the effect of audience uptake on lie attributions is mediated by assertion attributions. We also used regression analysis and causal search to investigate whether people's lie attributions are caused by the attribution of deceptive intent. In order to accomplish this, we collected people's judgments about assertion and intent. Surprisingly, prior work on lie attributions has typically assumed that participants interpreted these critical variables in the relevant way.

Experiment 1 tests whether lie attributions are affected by the audience's ability to hear the agent's spoken words, and whether such an effect is mediated by assertion attributions. Experiment 2 tests whether lie attributions are affected by the presence of an audience. Experiment 3 tests whether lie attributions are affected by whether the auditor comprehends the speaker's language. Experiment 4 uses regression analysis and causal search to investigate the relationship among judgments about five factors, including lie attributions and deceptive-intent attributions.

Experiment 1

This experiment tests whether lie attributions are affected by the audience's ability to hear the agent's spoken words. More specifically, we tested whether lie attributions are affected by whether the audience can hear the agent due to factors external or internal to the agent. The experimental design also allowed us to test whether lie attributions are based on attributing an as-

sertion to the agent. To attribute an assertion is not merely to attribute the utterance of words. A verbal assertion requires uttering words and typically uttering a declarative sentence suffices to make an assertion. For instance, saying, "Allen is meeting a client," is typically counted as asserting that Allen is meeting a client, which brings with it a commitment to the truth of that proposition. But this is not always the case. For example, an actor uttering those words on stage is not making an assertion. (For additional theoretical discussion and references, see Turri, 2010; Turri, 2012.)

Method

Participants

One hundred and twenty-five participants were tested (aged 18-83 years, mean age = 33 years; 49 female; 91% reporting native competence in English). Participants were U.S. residents, recruited and tested online using Amazon Mechanical Turk and Qualtrics, and compensated $0.40 for approximately 2-3 minutes of their time. Repeat participation was prevented within and across experiments. We used the same recruitment and compensation procedures were used for all experiments reported here.

Materials and Procedure

Participants were randomly assigned to one of four conditions in a 2 (Ability: unable, able) × 2 (Obstruction: external, internal) between-subjects design. Participants read a brief text and answered three test statements. The story featured Samantha, whose co-worker Allen recently start-

ed taking an extra coffee break in the morning. The boss asks Samantha where Allen is. Samantha thinks (correctly) that he is taking an extra coffee break, but she tells the boss that he is meeting a client. The Ability factor manipulated whether the boss was unable or able to hear Samantha's words. The Obstruction factor manipulated whether the potential obstruction was external or internal to Samantha. The external factor was noise from construction upstairs; the internal factor was a failing larynx.

After reading the text, participants rated their agreement with three statements:

1. Samantha lied. (Lied)
2. Samantha asserted that Allen was meeting a client. (Asserted)
3. The boss heard what Samantha said. (Heard)

The first item was the critical measure of lie attributions. The second item measured whether people viewed Samantha as asserting the relevant proposition. The purpose of including this measure was to check whether an effect on lie attributions was mediated by assertion attributions. The third item was a check on the Ability manipulation. Response to these items was collected on a standard 7-point Likert scale, 1 ("strongly disagree") – 7 ("strongly agree"), left-to-right across the participant's screen. Each item appeared on a new screen and always in the order listed above. The story remained visible at the top of the screen throughout testing. The complete text of all stimuli is included in the Appendix. After testing, participants completed a brief demographic questionnaire.

Results

The Ability manipulation was extremely effective, with participants strongly disagreeing that the boss heard Samantha in unable conditions and strongly agreeing that he heard her in the able conditions. (See Table 1.) An analysis of variance on response to the lie attribution revealed a main effect of Ability, $F(1, 121) = 8.72$, $p = .004$, $\eta_p^2 = .067$, a trending effect of Obstruction, $F(1, 121) = 2.99$, $p = .086$, $\eta_p^2 = .024$, and an interaction effect of Ability and Obstruction, $F(1, 121) = 7.95$, $p = .006$, $\eta_p^2 = .062$. (See Figure 1.) Because of the interaction, we conducted separate follow-up independent samples t-tests for external and internal obstructions. When the obstruction was external, lie attributions were unaffected by whether the boss was able to hear Samantha, $p = .912$, but when the obstruction was internal, lie attributions were significantly lower when the boss was unable to hear Samantha, $t(60) = -3.62$, $p < .001$, $d = 0.93$. (See Figure 1.) An analysis of variance on response to the assertion attribution revealed a main effect of Ability, $F(1, 121) = 3.91$, $p = .050$, $\eta_p^2 = .031$, a trending effect of Obstruction, $F(1, 121) = 3.11$, $p = .080$, $\eta_p^2 = .025$, and no interaction, $F(1, 121) = 1.53$, $p = .219$. Assertion attribution was lower when the boss was unable to hear Samantha ($M = 5.32$, $SD = 1.93$) than when he was able to hear her ($M = 5.89$, $SD = 1.34$).

Table 1. Experiment 1. Descriptive statistics.

	Unable				Able			
	External (N = 32)		Internal (N = 30)		External (N = 31)		Internal (N = 32)	
Measure	**M**	**SD**	**M**	**SD**	**M**	**SD**	**M**	**SD**
lied	6.06	1.34	4.90	2.09	6.10	1.11	6.38	0.94
asserted	5.75	1.74	4.87	2.05	5.97	1.33	5.81	1.38
heard	1.56	0.80	1.50	1.14	6.06	1.39	5.69	1.91

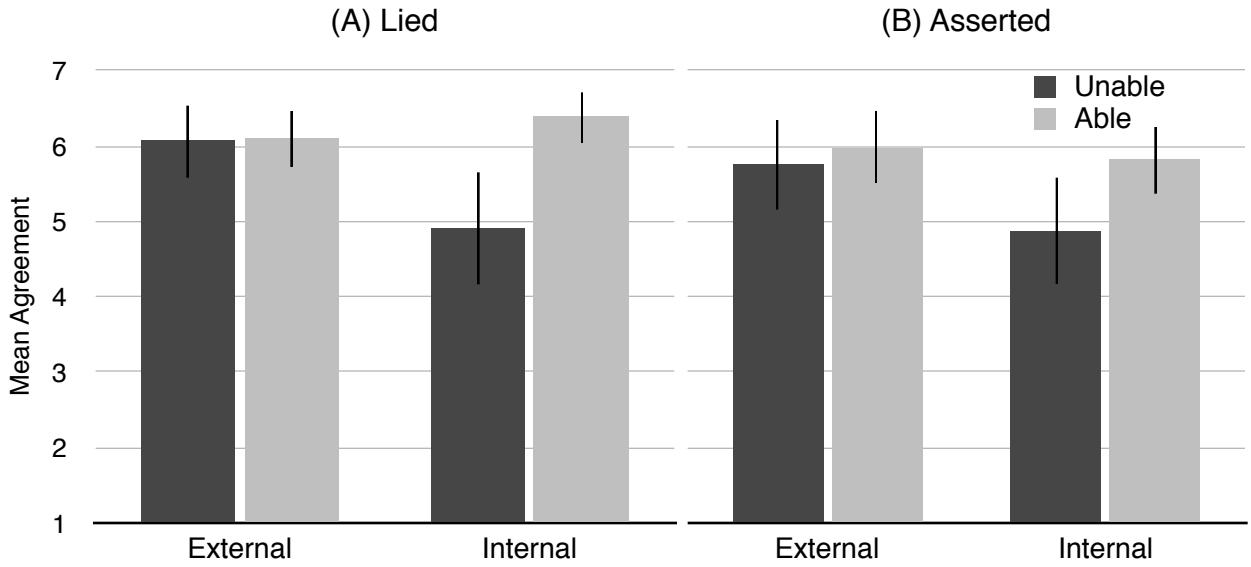

Fig. 1. Experiment 1. Mean attribution of lying (Panel A) and assertion (Panel B) in the four conditions. Scales ran 1 (strongly disagree) - 7 (strongly agree). Error bars represent 95% confidence intervals.

If lie attributions are based on assertion attributions, then one might expect assertion attributions to mediate the observed effect on lie attributions. To test the mediating role of assertion attributions on lie attributions, we conducted a bootstrap mediation analysis (Hayes, 2013). The analysis focused on the Internal conditions. We used assignment to Ability condition as the inde-

pendent variable (coded: 0 = unable, 1 = able), lie attribution as the outcome, and assertion attribution as potential mediator. This analysis showed that assertion attributions significantly mediated the effect of Ability on lie attributions: indirect effect = 0.39 [0.03, 1.05], direct effect = 1.09 [0.32, 1.85]. (See Figure 2.)

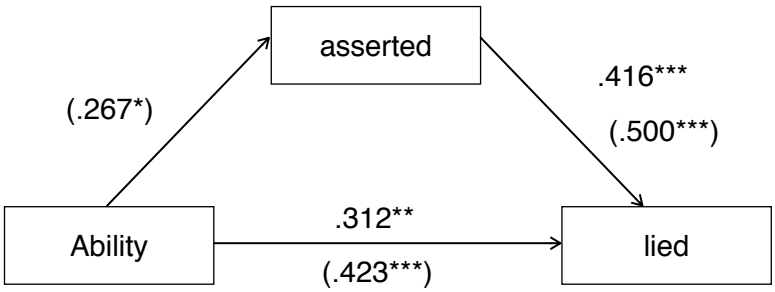

Fig. 2. Experiment 1. Mediation results. Parenthetical values represent the strength of a simple regression between the two variables; values outside the parentheses represent the strength of the relationships in a model used to test for mediation. *p < .05, **p < .01, ***p < .001.

Discussion

This experiment tested whether lie attributions are affected by the audience's ability to hear the agent's spoken words. When the source of an audience's inability was external to the speaker, the inability did not affect lie attributions. More specifically, when ambient noise prevented the audience from hearing the speaker, lie attributions did not differ from a control condition in which the audience was able to hear the speaker. By contrast, when the source of the audience's inability was internal to the speaker, the inability did affect lie attributions. More specifically, when a failure of the speaker's larynx prevented the audience from hearing, lie attributions were signifi-

cantly lower than in a control condition in which the audience was able to hear. Mediation analysis showed that this effect on lie attributions was mediated by participants' judgments about whether the speaker made an assertion. That is, the effect on lie attributions was mediated by assertion attributions. This provides some evidence that lie attributions are based on assertion attributions, which supports the standard view that a lie is a type of assertion.

Experiment 2

This experiment tests whether lie attributions are affected by whether an audience is present to hear the agent speak.

Method

Participants

Ninety-one new participants were tested (aged 19-64 years, mean age = 33 years; 40 female; 97% reporting native competence in English).

Materials and Procedure

The basic testing procedures were very similar to those used in Experiment 1. Participants were randomly assigned to one of two conditions (no audience, audience) in a between-subjects design. Participants read a brief text and answered three test statements. The story featured William, who works on a sales team whose current objective is to get new corporate customers.

William gets two new corporate customers in the morning then at lunchtime says that he got four new ones. The two conditions differed in whether everyone else had the day off, in which case no-one was around to hear William, or no-one had the day off, in which case there were people around to hear him. After reading the text, participants rated their agreement with three statements:

1. William lied. (Lied)
2. William's statement was intended to deceive. (Intended)
3. William asserted that he got four new corporate customers. (Asserted)

The first item was the critical measure of lie attributions. The second item measured whether people viewed William's statement as deceptively motivated. The purpose of including this measure was to check whether lie attributions correlated with the attribution of deceptive intent. The third item measured whether people viewed William as making an assertion. The complete text of all stimuli is included in the Appendix.

Results

Independent samples t-tests revealed that assignment to condition (no audience, audience) marginally affected lie attributions, had a large effect on deceptive-intent attributions, and did not affect assertion attributions. (See Figure 3 and Table 2.) When there was no audience, mean lying attribution ($M = 5.60$, $SD = 1.75$) was higher than mean deceptive-intent attribution ($M = 4.04$, $SD = 1.82$), paired samples t-test, $t(44) = 5.63$, $p < .001$, $d = 0.84$. Nevertheless, lie attributions and deceptive-intent attributions were strongly positively correlated, $r = .57$, $n = 91$, $p < .001$. Lie

attributions and assertion attributions were also strongly positively correlated, $r = .51$, $n = 91$, $p < .001$. Deceptive-intent attributions and assertions attributions were moderately positively correlated, $r = .32$, $n = 91$, $p = .002$.

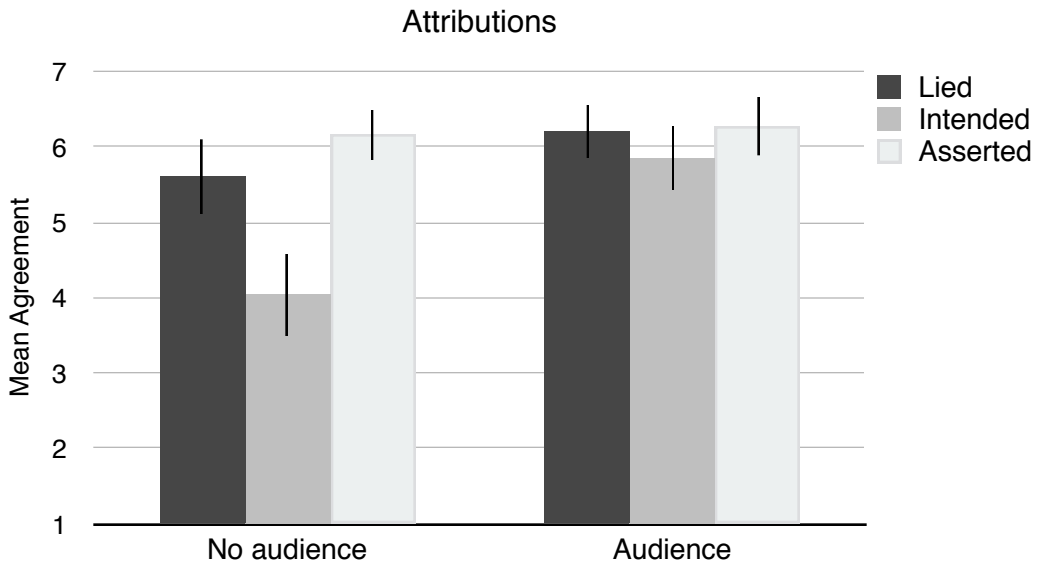

Fig. 3. Experiment 2. Mean attribution of lying, deceptive-intent, and assertion (within-subjects) in the two conditions (no audience, audience). Scales ran 1 (strongly disagree) - 7 (strongly agree). Error bars represent 95% confidence intervals.

Table 2. Experiment 2. Independent samples t-tests.

	No audience (N = 45)		Audience (N = 46)							
Measure	**M**	**SD**	**M**	**SD**	**t**	**df**	**p**	**MD**	**95% CI**	**d**
lied	5.60	1.75	6.20	1.22	-1.89	78.52	.063	-0.60	-1.22, 0.03	0.43
intended	4.04	1.82	5.85	1.35	-5.36	81.09	<.001	-1.80	-2.47, -1.13	1.20
asserted	6.16	1.11	6.26	1.12	-0.45	89	.654	-0.11	-0.57, 0.36	0.10

Discussion

This experiment tested whether lie attributions are affected by an audience's presence. When there was an audience to hear the speaker's false statement, participants judged that the speaker lied and intended to deceive. By contrast, when there was no audience to hear the speaker's false statement, participants judged that the speaker lied but were ambivalent about the speaker's intent. Although the attributions of lying and deceptive intent were positively correlated, when there was no audience, mean lie attribution was much higher than mean deceptive-intent attribution.

Experiment 3

This experiment tests whether lie attributions are affected by whether the auditor can comprehend the language spoken by the agent.

Method

Participants

Ninety-two new participants were tested (aged 19-68 years, mean age = 33 years; 42 female; 98% reporting native competence in English).

Materials and Procedure

The basic testing procedures were very similar to those used in Experiment 2. Participants were randomly assigned to one of two conditions (unable, able) in a between-subjects design. Participants read a brief text and answered three test statements. The story featured Fiona, who is traveling abroad and who has not visited an art museum. During an elevator ride with a local person, Fiona tells him that she visited an art museum today. The two conditions differed in whether the local can comprehend English. In the unable condition, Fiona knows that the local cannot comprehend English and, thus, that there is no chance he could be deceived by her speech. In the able condition, Fiona knows that the local comprehends English very well and, thus, that he might be deceived by her speech. After reading the text, participants rated their agreement with three statements:

1. Fiona lied. (Lied)
2. Fiona's statement was intended to deceive. (Intended)
3. The local person understood Fiona's statement. (Comprehended)

The first item was the critical measure of lie attributions. The second item measured whether

people viewed Fiona's statement as deceptively motivated. The third item was a check on the comprehension manipulation. The complete text of all stimuli is included in the Appendix.

Results

The manipulation was effective, with participants strongly disagreeing that the local understood in the unable condition, and strongly agreeing in the able condition. (See Table 3.) Independent samples t-tests revealed that assignment to condition (unable, able) did not affect lie attributions, but it had a large effect on deceptive-intent attributions, with attributions lower when the agent was unable to comprehend. (See Figure 4 and Table 3.) When the agent was unable to comprehend, mean lie attribution ($M = 5.46$, $SD = 1.82$) was higher than mean deceptive-intent attribution ($M = 3.76$, $SD = 1.96$), paired samples t-test, $t(45) = 5.75$, $p < .001$, $d = 0.85$. Nevertheless, lie attributions and deceptive-intent attributions were strongly positively correlated, $r = .59$, $n = 92$, $p < .001$.

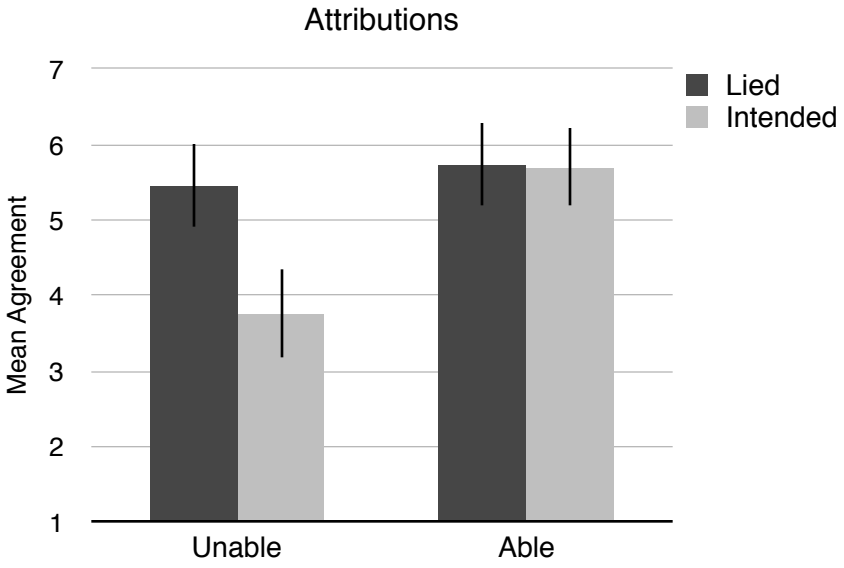

Fig. 4. Experiment 3. Mean attribution of lying and deceptive intent (within-subjects) in the two conditions (unable, able). Scales ran 1 (strongly disagree) - 7 (strongly agree). Error bars represent 95% confidence intervals.

Table 3. Experiment 3. Independent samples t-tests.

	Unable (N = 46)		**Able** (N = 46)							
Measure	**M**	**SD**	**M**	**SD**	**t**	**df**	**p**	**MD**	**95% CI**	**d**
lied	5.46	1.82	5.74	1.90	-0.73	90	.469	-0.28	-1.05, 0.49	0.15
intended	3.76	1.96	5.70	1.88	-4.83	90	<.001	-1.94	-2.73, -1.14	1.02
comprehended	1.67	0.76	6.15	1.15	-21.97	90	<.001	-4.48	-4.88, -4.07	4.63

Discussion

This experiment tested whether lie attributions are affected by whether the auditor can comprehend the language spoken by the agent. When the auditor could comprehend, participants strongly attributed lying and deceptive intent to the agent. When the auditor could not comprehend,

participants strongly attributed lying but did not attribute deceptive intent. This replicates a finding from Experiment 2 whereby participants strongly attributed lying but, overall, were ambivalent about deceptive intent.

Experiment 4

This experiment uses regression analysis and causal search to investigate the relationship among judgments about lying, deceptive intent, auditor comprehension, truth-value, and ethical evaluation.

Method

Participants

Ninety-one new participants were tested (aged 22-70 years, mean age = 33 years; 37 female; 95% reporting native competence in English).

Materials and Procedure

Participants were randomly assigned to one of two conditions (false, true) in a between-subjects design. Participants read a brief text and answered five test statements. The basic story was very similar to the one used for the unable condition in Experiment 3. The two conditions differed in whether Fiona's statement is false (she said she visited a museum but did not) or true (she said she visited a museum and did). After reading the text, participants rated their agreement with five

statements:

1. Fiona lied. (Lied)
2. Fiona's statement was intended to deceive. (Intended)
3. It was unethical for Fiona to make that statement. (Unethical)
4. The local person understood Fiona's statement. (Comprehended)
5. Fiona's statement was false. (False)

In contrast to the experiments reported above, the statements appeared all at once beneath the story in a matrix table and their order was randomized. The point of these procedural differences was to avoid order effects. Responses were collected using the same 7-point Likert scale as in previous experiments. The complete text of all stimuli is included in the Appendix.

Results

Independent samples t-tests revealed that assignment to condition (false, true) affected response to all dependent measures except for the comprehension attribution. (See Table 4.) We conducted a multiple linear regression analysis to assess which judgments significantly predicted lie attributions. The model included deceptive-intent attributions, ethical evaluations, comprehension attributions, and falsity evaluation as predictors of lie attributions. (We did not include the independent variable, which was a truth-value manipulation. Instead, we included the participants' own evaluation of the relevant proposition's truth-value.) The model was statistically significant and explained over 90% of the variance in lie attributions, $F(4, 86) = 204.44$, $p < .001$, $R^2 = .905$. (See Table 5.) Deceptive-intent attributions and ethical evaluations did not significantly predict

lie attributions, but comprehension attributions and falsity evaluations did. Comprehension attributions explained less than 1% of variance in lie attributions; falsity evaluations explained 44%.

Table 4. Experiment 4. Independent samples t-tests.

	False (N = 46)		True (N = 45)							
Measure	**M**	**SD**	**M**	**SD**	**t**	**df**	**p**	**MD**	**95% CI**	**d**
lied	5.65	1.75	1.84	1.43	11.34	89	<.001	3.81	3.14, 4.48	2.40
intended	4.07	1.96	2.18	1.60	5.04	86.3	<.001	1.89	1.14, 2.63	1.09
unethical	3.87	1.96	2.20	1.60	4.44	89	<.001	1.67	0.92, 2.42	0.94
comprehended	1.91	1.15	1.73	1.03	0.78	89	.435	0.18	-0.28, 0.64	0.16
false	5.96	1.62	1.80	1.33	13.39	89	<.001	4.16	3.54, 4.77	2.84

Table 5. Experiment 4. Multiple linear regression predicting lie attributions.

Predictor	**B**	**SE (B)**	**β**	**t**	**p**
intended	.094	.065	.076	1.44	.153
unethical	.087	.063	.069	1.37	.176
comprehended	.198	.079	.087	2.50	.014
false	.826	.042	.848	19.84	<.001
Constant	-.371	.203		-1.83	.071

To better understand the psychological processes involved in making these judgments, we also conducted a causal search with the Greedy Equivalence Search (GES) algorithm. Similar to regression-based mediation analysis and structural equation modeling, GES is used to make causal inferences from correlations and independence relations in a dataset. However, GES has certain advantages. Whereas mediation analysis and structural equal modeling assume a causal direction, GES does not. GES provides an overall measure of model fit and, given enough data,

will return the true causal model (Chickering, 2002). And GES is not limited to considering a single pre-specified model, which typically depends on theoretical assumptions. GES considers all possible models available given the different variables. Each variable is treated as a node. GES assigns an information score to the model in which all the nodes are disconnected — the "null model." GES then evaluates adding causal arrows — "edges" — between the nodes (Meek 1997 provides the edge orientation rules). GES adds edges that best improve the model's information score, if such edges exist, until adding more edges does not improve the information score. At this point, GES evaluates whether deleting any edges will further improve the information score, and it deletes any such edges until deleting more edges does not improve the information score. (For more on the general theory behind causal search, see Spirtes, Glymour & Scheines, 2000; Pearl, 2000.)

The causal search was conducted using Tetrad 5.1.0 (http://www.phil.cmu.edu/projects/tetrad/). We entered assignment to condition and response to the five dependent measures (lied, intended, unethical, comprehended, false) into a causal search using GES. The penalty discount was set to .75 and the model was constrained so that assignment to condition could not be caused by any other variable in the model. The algorithm returned a good fitting model for the data, $\chi^2(10) = 11.46$, $p = .323$, BIC = -33.65. Figure 5 depicts the best fitting model. In the model, lie attributions are caused by falsity evaluations and comprehension attributions; lie attributions cause deceptive-intent attributions, which in turn cause ethical evaluations. The causal link from falsity evaluations to lie attributions was the strongest relationship among any dependent variables in the model.

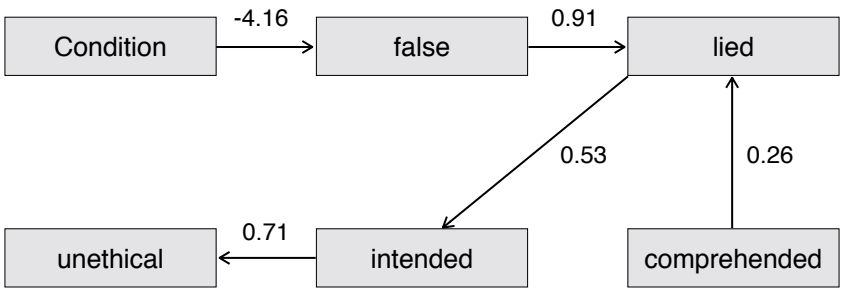

Fig. 5. Experiment 4. Causal search results. Arrows represent directional causal relations; path coefficients represent the strength of the causal relation. Reference class for condition: false.

Discussion

This experiment used regression analysis and causal search to investigate the relationships among judgments about lying, deceptive intent, auditor comprehension, truth-value, and ethical evaluation. The results from the two procedures were broadly similar. In the regression analysis, lie attributions were significantly predicted by judgments about auditor comprehension and evaluations of truth-value, but not by attributions of deceptive intent or ethical evaluations. In the best fitting causal model for the data, lie attributions were caused by judgments about auditor comprehension and evaluations of truth-value; moreover, lie attributions caused attributions of deceptive intent, which reverses the direction one might expect given the standard view of lying.

General Discussion

Results from four experiments shed further light on the ordinary concept of lying and lie attributions. The research took place against the backdrop of a standard definition of lying, according to

which a lie is a dishonest assertion (i.e. you lie if and only if you make an assertion that you think is false in order to deceive your audience). In Experiment 1, lie attributions were affected by the audience's ability to hear the agent's spoken words. When the source of the audience's inability was internal to the speaker (a debilitated larynx), lie attributions were significantly lower than in a control condition in which the audience was able to hear. This effect on lie attributions was mediated by participants' judgments about whether the speaker made an assertion (assertion attributions). In Experiment 2, when an agent was alone and no-one could hear his words, participants strongly attributed lying but did not attribute deceptive intent. In Experiment 3, when the auditor could not comprehend the speaker's language, participants again strongly attributed lying but did not attribute deceptive intent. In Experiment 4, lie attributions were caused by judgments about auditor comprehension and evaluations of truth-value, but not by attributions of deceptive intent. Instead, lie attributions caused attributions of deceptive intent.

These findings advance knowledge of the ordinary concept of lying and lie attributions in several ways. First and foremost, they suggest that deceptive intent is less central to the ordinary concept of lying than previously thought. When an agent falsely asserted a proposition with no-one around to hear, or to an auditor who could not comprehend the language, participants strongly judged that the agent lied without also attributing deceptive intent to the agent (Experiments 2 and 3). Moreover, in other contexts, deceptive-intent attributions did not significantly predict lie attributions. Instead, on the best fitting causal model for the data, deceptive-intent attributions were actually a *consequence* of lie attributions, rather than a cause (Experiment 4). This supports some theories of lying, which claim that deceptive intent is inessential to it (Sorenson, 2007;

Carson, 2006; Fallis, 2009). It also fits with recent experimental work, which found that participants attributed lies in contexts where, according to the researchers, the agent lacked an intention to deceive (Arico & Fallis, 2013, Experiment 1). However, one concern about this result, noted by the researchers themselves, is its assumption that participants shared the researchers' interpretation of the case. Participants still might have attributed deceptive intent. By contrast, the present research avoided this concern by having participants judge deceptive intent for themselves.

Suppose it turns out to be a robust finding that lie attributions cause deceptive-intent attributions, rather than vice versa. This fact, all by itself, is still consistent with the ordinary concept of lying being partly constituted by its relation to deceptive intent. For instance, some researchers distinguish between "containment" and "inferential" models of conceptual structure (Laurence & Margolis, 1999, p. 5). On the containment model, the concept of lying would be constituted by other concepts, including the concepts of assertion and falsity and, perhaps also, deceptive intent. Given a concept structured this way, a lie attribution would be partly constituted by a deceptive-intent attribution, and it would not be possible for the former to cause the latter. By contrast, on the inferential model, a concept is constituted by a set of inferential dispositions to apply other concepts. On this model, a lie attribution need not be informed by a prior deceptive-intent attribution. That is, a lie attribution could occur without a deceptive-intent attribution occurring. Instead, a lie attribution need only generate a disposition to attribute deceptive intent. And this is consistent with our finding that lie attributions caused deceptive-intent attributions. Indeed, one possibility that fits well with our results is that the concept of lying is a containment/inferential

hybrid. For instance, it could be that the concept of lying is partly constituted by a containment relation with *falsity* and *assertion*, and partly constituted by an inferential relation with *deceptive intent*. This would explain the causal modeling results from Experiment 4.

Of course, we are not confidently proposing this hybrid model based on our findings here. Our primary goal is merely to flag a potentially fruitful avenue of future research on the concept of lying, based on more general theoretical distinctions researchers have proposed about the nature of concepts. Also consistent with our results is the hypothesis that the relationship between lying and deceptive intent is merely correlational rather than conceptual.

Second, prior work has assumed that lies were a species of assertion (e.g. Coleman & Kay, 1981), but little if any evidence exists that this accurately reflects the ordinary concept. We found evidence that people's lie attributions are sensitive to their assertion attributions (Experiment 1). As far as we are aware, the present research is the first to measure assertion attributions and provide evidence that lie attributions are based on them.

Third, prior work has shown that lie attributions are sensitive to contextual features, such as the purpose of the conversational exchange (informative versus politeness) (Lee & Ross, 1997). Our results suggest that lie attributions are also sensitive to contextual features related to audience uptake. In particular, lie attributions are sometimes based on whether the audience comprehends the speaker's language (Experiment 4). The observed influence of comprehension attributions on lie attributions was very small, explaining less than 1% of variance in lie attributions. However, this occurred in a context where we did not manipulate audience comprehension. One possibility for future work, then, is to manipulate this factor in order to see whether it can play a

larger role.

Fourth, theorists have argued that there are different types of lie, including "straightforward lies," "bald-faced lies," "confused lies," "proviso lies," and "white lies" (e.g. Arico & Fallis, 2013; Peterson, Peterson & Seeto, 1983). Our results suggest a potentially surprising addition to this list: *private lies*. When an agent asserted a proposition that he presumably knew was false, people strongly judged that it was a lie regardless of whether he did so publicly with others around, or privately when he was alone. Not only does this result suggest that intending to deceive your audience is inessential to lying, but it also suggests that no audience is needed. Of course, whenever you assert, *you* are present. Future work could investigate whether people interpret private lies as attempts to deceive oneself.

Finally, existing findings suggest that objective falsity is essential to the ordinary concept of lying. But no clear picture has emerged about how important objective falsity is. For instance, some results suggest that objective falsity is the least important feature of a lie (Coleman & Kay, 1981, p. 35). By contrast, other results suggest that objective falsity is enormously important. In one study, the switch from an objectively false to an objectively true statement caused lie attributions to drop from ceiling (90%) to floor (10%) (Turri & Turri, 2015). The present results provide further evidence that objective falsity is essential and centrally important to the ordinary concept of lying. In the best-fitting causal model of data from one study, evaluations of a statement's falsity were by far the strongest cause of lie attributions (Experiment 4) and uniquely explained nearly half of the variance in lie attributions.

Acknowledgments — This research was supported by the Social Sciences and Humanities Research Council of Canada and an Early Researcher Award from the Ontario Ministry of Economic Development and Innovation.

Appendix: Stimuli

Experiment 1

External obstacle (unable/able)

Samantha's co-worker Allen recently started taking an extra coffee break in the morning. One morning the boss visits Samantha's unit and asks, "Where is Allen?" Samantha thinks that Allen is taking an extra coffee break, so she says, "He's meeting a client." (In fact, Allen is out on a coffee break.) Just as Samantha spoke, the volume of noise from the construction upstairs changed suddenly and the boss was [able/completely unable] to hear what she said.

Internal obstacle (unable/able)

Samantha's co-worker Allen recently started taking an extra coffee break in the morning. One morning the boss visits Samantha's unit and asks, "Where is Allen?" Samantha thinks that Allen is taking an extra coffee break, so she says, "He's meeting a client." (In fact, Allen is out on a coffee break.) Just as Samantha spoke, her severely weak larynx suddenly failed [but the boss was able/and the boss was completely unable] to hear what she said.

1. Samantha lied.

2. Samantha asserted that Allen was meeting a client.
3. The boss heard what Samantha said.

Experiment 2

(No audience/audience) William is a part of a sales team at work. Their major objective this month is to get new corporate customers. This morning in his cubicle at work, William made many calls and got two new corporate customers. At the water cooler during lunch, William says, "I got four new corporate customers this morning." [It's Saturday and everyone else had the day off: William was alone and no-one was around to hear him. / It's Saturday but no-one had the day off: William was not alone and people were around to hear him.]

1. William lied.
2. William's statement was intended to deceive.
3. William asserted that he got four new corporate customers.

Experiment 3

Unable

Fiona is traveling abroad. She speaks only English, not the local language. Fiona has not been to an art museum. This evening, a local person is the only one in an elevator with Fiona. Fiona knows that this local person clearly cannot speak or comprehend a word of English. Fiona says to him, in English, "Today I visited an art museum." Because this person cannot comprehend a word of English, Fiona knows that he will not understand her and, thus, there is no chance of de-

ceiving him.

Able

Fiona is traveling abroad. She speaks only English, not the local language. Fiona has not been to an art museum. This evening, a local person is the only one in an elevator with Fiona. Fiona knows that this local person clearly can speak and comprehend English fluently. Fiona says to him, in English, "Today I visited an art museum." Because this person speaks and comprehends English fluently, Fiona knows that he will understand her and, thus, there is a chance of deceiving him.

1. Fiona lied.
2. Fiona's statement was intended to deceive.
3. The local person understood Fiona's statement.

Experiment 4

(False/true) Fiona is traveling abroad. She speaks only English, not the local language. Fiona [did not visit/visited] an art museum today. This evening, a local person is the only one in an elevator with Fiona. Fiona knows that this local person clearly cannot speak or comprehend a word of English. Fiona says to him, in English, "Today I visited an art museum." Because this person cannot comprehend a word of English, Fiona knows that he will not understand her and, thus, that he will not believe what she said.

1. Fiona lied.
2. Fiona's statement was intended to deceive.

3. It was unethical for Fiona to make that statement.
4. The local person understood Fiona's statement.
5. Fiona's statement was false.

References

Aquinas, T. 1273. *Summa Theologica*. Amazon Digital Services.

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. http://doi.org/10.1080/09515089.2012.725977

Augustine. 395 CE. On lying. In P. Schaff & K. Knight (Eds.), *Nicene and post-Nicene Fathers*, First Series (Vol. 3). Christian Literature Publishing. Retrieved from http://www.newadvent.org/fathers/1312.htm

Carson, T. L. 2006. The definition of lying. *Noûs*, 40(2): 284–306. http://doi.org/10.1111/j.0029-4624.2006.00610.x

Chickering, D. M. 2002. Optimal structure identification with greedy search. *The Journal of Machine Learning Research*, 3(1): 507–554. http://doi.org/10.1109/69.494161

Coleman, L., & Kay, P. 1981. Prototype semantics: the English word lie. *Language*, 57(1): 26–44.

Fallis, D. 2009. What is lying? *Journal of Philosophy*, 106(1): 29–56.

Hayes, A. F. 2013. *Introduction to mediation, moderation, and conditional process analysis: A*

regression-based approach. New York: Guilford Press.

Laurence, S., & Margolis, E. 1999. Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: core readings* (pp. 3–81). MIT Press.

Lee, K., & Ross, H. J. 1997. The concept of lying in adolescents and young adults: Testing Sweetser's folkloristic model. *Merrill-Palmer Quarterly*, 43(2): 255–270.

Meek, C. 1997. Graphical models: selecting causal and statistical models. Carnegie Mellon University.

Pearl, J. 2000. *Causality: models, reasoning, and inference*. Cambridge: Cambridge University Press.

Peterson, C. C., Peterson, J. L., & Seeto, D. 1983. Developmental changes in ideas about lying. *Child Development*, 54(6): 1529–1535.

Sorensen, R. 2007. Bald-faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly*, 88(2): 251–264.

Spirtes, P., Glymour, C., & Scheines, R. 2000. *Causation, prediction, and search*. Cambridge, Mass.: MIT Press.

Strichartz, A. F., & Burton, R. V. 1990. Lies and truth: A study of the development of the concept. *Child Development*, 61(1): 211–220.

Sweetser, E. E. 1987. The definition of lie: an examination of the folk models underlying a semantic prototype. In D. Holland (Ed.), *Cultural models in language and thought*. Cambridge: Cambridge University Press.

Turri, J. 2010. Epistemic invariantism and speech act contextualism. *Philosophical Review*,

119(1): 77–95. http://doi.org/10.1215/00318108-2009-026

Turri, J. 2012. Pyrrhonian skepticism meets speech-act theory. *International Journal for the Study of Skepticism*, 2(2): 83–98. http://doi.org/10.1163/221057011X588037

Turri, A., & Turri, J. 2015. The truth about lying. *Cognition*, 138(C): 161–168. http://doi.org/10.1016/j.cognition.2015.01.007

Vrij, A. 2008. *Detecting lies and deceit: pitfalls and opportunities* (2nd ed.). Hoboken, NJ: John Wiley & Sons.