

**Community standards of deception:
Deception is perceived to be ethical when it prevents unnecessary harm**

Emma E. Levine

Supplementary Online Materials

Content:

2. Inductive study	4
2.1. Additional details on survey	4
2.2. Additional details on coding process	4
2.3. Additional results	6
3. Study 2	8
3.1. Choice Options for Each Vignette in Study 2	8
3.2. Manipulation check questions and results for each vignette	11
3.3. Study 2 Choice Results, excluding those who failed manipulation check	13
3.4. Mechanism results by vignette	14
3.5. Pooled regression analyses for Study 2	16
4. Study 3	17
4.1. Choice Options for Each Vignette in Study 3	17
4.2. Main analyses in Study 3, Split by sample	18
4.3. Correlation table for Study 3 DVs	21
4.4. Mediation analyses without moral duty	22
5. Study 4	23
5.1. Additional measures in Study 4a	23
5.2. Descriptive Statistics in Studies 4a and 4b	26
5.3. Results of Studies 4a, Split by relationship type	27
6. Details of additional studies	28
6.1. Study S1: Is lying always wrong? pilot	28
6.2. Study S2: Lying in everyday life pilot	30
6.3. Study S3: Targets' reactions to lies in everyday life	32
7. Relationship between immediate harm and instrumental value	35

1. Stopping rules

For all studies, I used a priori stopping rules. Notably, stopping rules were not based on formal power calculations, but rather convenience (e.g., collecting data for the length of one laboratory session) or back-of-the-envelope calculations based on my past research in this space (e.g., 50 participants per cell for a scenario study). For each study, I collected data until the target (whether length of time or number of participants) was hit and then analyzed all available data.

1a. Inductive study

This study featured a combined sample of lab participants and Amazon Mechanical Turk participants. I made the a priori decision to distribute the survey in the laboratory for the length of one laboratory session (3 days) and then stop data collection. This resulted in 187 participants who were recruited from the lab (59% female; $M_{\text{age}} = 24$ years).

I made the a priori decision to collect 100 participants from Amazon Mechanical Turk. Some extra participants completed the survey before it closed, resulting in a recruitment of 117 adults (50% female; $M_{\text{age}} = 37$ years).

1b. Study 1: Theory testing surveys

This study featured three separate surveys, two of which were hosted on Amazon Mechanical Turk and one of which was hosted in a U.S. laboratory.

The laboratory survey was run first. I made the a priori decision to distribute the survey in the laboratory for one week, and then stop data collection. This resulted in 60 participants who were recruited from the lab (47% female; $M_{\text{age}} = 26$ years).

I then made the a priori decision to collect 60 participants for each Amazon Mechanical Turk survey, based on the number of participants I had recruited from the lab. Some extra participants completed one of the surveys before it closed, resulting in 84 adults who were recruited via Amazon Mechanical Turk (51% female; $M_{\text{age}} = 36$ years) in the first MTurk survey and 60 (50% female; $M_{\text{age}} = 36$ years) in the second MTurk survey.

1c. Study 2: Community standard surveys

This study featured three separate surveys, two of which were hosted on Amazon Mechanical Turk and one of which was hosted in a U.S. laboratory.

I made the a priori decision to distribute the survey in the laboratory for the length of one laboratory session (3 days), and then stop data collection. This resulted in 195 participants who were recruited from the lab (52% female; $M_{\text{age}} = 25$).

I made the a priori decision to collect 250 participants for each Amazon Mechanical Turk survey. Some extra participants completed the surveys before they closed, resulting in 267 adults who were recruited via Amazon Mechanical Turk (46.8% female; $M_{\text{age}} = 35$) in the first MTurk survey and 269 (45.4% female; $M_{\text{age}} = 38$) in the second MTurk survey.

1d. Study 3: Ruling out alternative mechanisms

This study featured three separate surveys, two of which were hosted on Amazon Mechanical Turk and one of which was hosted in a U.S. laboratory.

I made the a priori decision to distribute the survey in the laboratory for the length of one laboratory session (3 days), and then stop data collection. This resulted in 142 participants who were recruited from the lab (61% female; $M_{age} = 23$ years).

I then made the decision to recruit the same number of participants on Amazon Mechanical Turk (target was 140 participants). I ended up with 136 adults recruited via Amazon Mechanical Turk (43% female; $M_{age} = 32$ years).

1e. Study 4: Everyday lies

For each study, I made the a priori decision to recruit 300 adults from a U.S. representative sample via Prolific Academic (<https://www.prolific.co/>). I ended up with a final sample of 296 participants who completed Study 4a (50% female; $M_{age} = 45$ years; $M_{work\ experience} = 23$ years), and 285 participants who completed Study 4b (50% female; $M_{age} = 44$ years; $M_{work\ experience} = 21$ years). Study 4b was preregistered at aspredicted.org (<https://aspredicted.org/yw6z2.pdf>). The preregistration, however, had a typo, indicating that each participant would rate 20, rather than 10, situations described by participants in Study 4a. Furthermore, I preregistered regressions in which I clustered standard errors at the participant level to account for within-participant dependence. In the main manuscript, however, I use random-effects models to account for within-participant dependence, based on guidance from the review team. The results from the clustered standard error models and the random effects models are qualitatively identical.

2. Inductive study

2.1. Additional details on survey

Procedure. All participants completed an online survey in which they answered free-response questions about deception. I randomly assigned participants to one of two conditions in a between-subjects design: *Preferences* or *Ethics*. Participants either answered three questions about their preferences for deception (the *Preferences* condition) or the general ethicality of deception (the *Ethics* condition).

In the *Preferences* condition, I first asked participants to “Think about when you would want someone to lie to you.” Then participants answered the following three questions” “In what circumstances would you want someone to lie to you?”, “In what circumstances would you not want someone to be completely honest with you?”, and “Please come up with three concrete examples of instances in which you would want to be lied to.” In other words, they indicated the lies that they would consent to being told.

In the *Ethics* condition, I first asked participants to “Think about when lying is right and when lying is wrong.” Then participants answered the following three questions” “In what circumstances is lying to someone the right thing to do?”, “In what circumstances is being completely honest with someone the wrong thing to do?”, and “Please come up with three concrete examples of instances in which it is ethical to lie.” In both conditions, participants had to respond to each question for at least one minute and write at least 500 characters. Then, I collected demographic information for exploratory purposes.

2.2. Additional details on coding process

My goal in this study was to develop a codified set of community standards and an underlying theory regarding lay perceptions of deception. Specifically, the goal was to identify the standards and underlying mechanisms that describe when people would want to be deceived and to explore whether deception was also perceived to be ethical in these cases. Therefore, the primary goals of my coding scheme were to unearth individuals’ preferences for deception, and to develop an underlying theory that described these preferences. To do this, I adopted an iterative coding procedure (Strauss & Corbin, 1990). I first read through 50 participants’ responses and developed a preliminary coding scheme informed by the present data, related research (DePaulo et al., 1996), and pilot data.

To perform the initial round of coding, I trained two research assistants to independently code all of the responses from both the *Preferences* and the *Ethics* perspectives using the initial coding scheme. The initial coding scheme required coders to read through each participant’s responses to all three questions and then code each participant’s responses according to the expressed justification for deception. The initial coding scheme included 12 possible justifications. I then met with both research assistants to collectively discuss the coding.

During this conversation, a single construct – (the prevention of) unnecessary harm – emerged as the overarching justification for deception. When discussing unnecessary harm, participants

discussed the degree to which deception could prevent harm to the target at the moment of communication and the degree to which honesty could yield instrumental benefits to the target, such as enlightenment and growth. That is, participants generally endorsed deception when it prevented *immediate harm* to the target *and* when honesty had no *instrumental value*.

After converging on this overarching justification, we also discussed 20 participant responses in detail and used this discussion to identify new coding categories and to clarify the categorization scheme for the next round of coding. During this discussion, we also realized that participants' responses to the second survey question often repeated content from their response to the first question. Furthermore, some participants misinterpreted the second survey question. Consequently, the final coding procedure focused on analyzing only responses to the first and third questions in the survey ("In what circumstances is lying to someone the right thing to do?/In what circumstances would you want someone to lie to you?" and "Please come up with three concrete examples of instances in which it is ethical to lie./ Please come up with three concrete examples of instances in which you would want to be lied to.").

Whereas the initial coding scheme focused primarily on identifying different reasons that deception is perceived to be ethical (or preferred to the truth), the final coding scheme focused on first categorizing responses along the two proposed components of unnecessary harm, and then categorizing the features of the target, honest information, and context that participants used to explain the existence of unnecessary harm. The final coding scheme also explored the frequency with which participants discussed deontological and utilitarian approaches to deception (see Table 1 in main manuscript).

Criteria for coding community standards. There were three criteria for maintaining an community standard category in the final coding scheme. First, the category had to be represented in more than one participant's response. This cutoff is intentionally low. Because the inductive study captures the salience of different circumstances in which deception may be justified, rather than the strength of the relationship between any particular circumstance and the justification of deception, I wanted to include standards that may not be particularly salient but very closely map onto the proposed dimensions of unnecessary harm.

Second, the category had to reflect a Harm Avoidance framework. There were a few justifications that appeared with some regularity that I did not include in the final coding scheme because they did not pertain to the prevention of harm: for example, lying to create a surprise or to win a game of poker. It is possible that there are other common justifications for deception that do not pertain to the prevention of harm, but that is not the focus of the present investigation.

Third, the coders needed to come to consensus on the meaning of the category. Several categories were dropped from the final coding scheme because they were too vague and did not lead to strong agreement. For example, the initial coding scheme included a category that read, "When the target is looking for something other than truth." However, this category was too broad and could be more easily categorized into the conditions that would lead the target to avoid truth (e.g., when s/he is fragile).

2.3 Additional results

Table S1. More Details of the Coding Categories for Justifications and Community Standards of Deception

Panel A. Broad Justifications for Deception

		Justification	Description for coders	Examples of participant responses	Kappa	Ethics	Pref	Total
Justifications for deception	Unnecessary harm avoidance	Immediate harm of honesty	These justifications include lies that are told to avoid harm to the target at the moment of communication. This type of harm is immediate and not long-lasting.	<ul style="list-style-type: none"> • From my perspective, lying to someone else is the right thing to do when we can avoid hurting others or make others happy / comfortable • Lying may be the right thing to do when telling that person the truth at that particular moment may be harmful to them. 	0.63	78.9%	60.7%	70.8%
		Instrumental value of honesty	These justifications focus on whether or not there are any potential long-term benefits of honesty. Specifically, is the honest information important, actionable, and objective? These responses suggest that lying is ok when honesty does not have the potential to affect future behavior or thinking in a meaningful way or bring about any other benefit.	<ul style="list-style-type: none"> • As long as it isn't something that's incredibly important for them to know, why bother them with it when you can save them from the truth? • I would want to be lied to under certain circumstances where I cannot change the result. 	0.74	65.7%	74.1%	69.4%
		TOTAL	This is a composite category reflecting the presence of either dimension above: Immediate harm or Instrumental value		0.75	92.2%	89.6%	91.0%
	Utilitarian		These justifications incorporate costs and benefits to parties other than the target of the lie. Any responses that mention how a lie will affect the liar, society, or third parties are considered Utilitarian.	<ul style="list-style-type: none"> • Lying to someone else is the right thing to do when it behooves both you and the other person to have them believe the lie. Lying may prevent conflicts... 	0.65	52.4%	17.8%	36.9%
	Never (Deontological)		"Never" indicates that the participants included a statement expressing that lying is never acceptable. "Never" means that the person does not provide any justifications or examples of when/why lying is right.	<ul style="list-style-type: none"> • There is no instance where lying to someone else is the right thing to do. • I would never want someone to lie to me. 	0.79	2.4%	8.1%	5.0%

Panel B. Community Standards of Deception

		Reason to lie	Definition	Examples of participant responses	Kappa	Ethics	Pref	Total
Community Standards of Deception	Attributes of Target	1. Emotionally fragile	When a person is in an emotionally fragile state (bad day, feeling sad, depressed, drunk, etc.)	• When a person is mentally unstable and his or her emotional well-being is at stake.	0.85	4.8%	4.4%	4.7%
		2. Cannot understand truth	When a person cannot cognitively understand the true information (a child, someone with dementia, etc.)	• When children ask questions about things that they should not know	0.86	25.3%	3.0%	15.3%
		3. Death Bed	When a person (the target) is at the end of their life	• I would want someone to lie to me about how long I might have to live if I were terminally ill.	0.92	7.8%	6.7%	7.3%
	Attributes of Topic	4. Subjective	When the truth is subjective (a function of different tastes, individual differences, preferences, a specific instance, etc.).	• I find a piece of clothing or accessory that I really like and makes me feel good, I would prefer not to have the person I'm with tell me he or she does not like what I've chosen	0.72	29.5%	29.6%	29.6%
		5. Trivial	When the topic is trivial (does not matter in any meaningful way to the target or others) or honesty is not the purpose of the exchange (e.g., social conventions or politeness are more important than honesty)	• I would rather have someone lie to me in trivial matters than important ones, because the magnitude of the issue at hand is smaller.	0.81	34.9%	22.2%	29.2%
		6. Uncontrollable	When the truth is about something that can never be changed (e.g., someone's height, a death, a relationship that has ended) or that feels outside of someone's control (e.g., weight, others' misdeeds) OR When the conversation occurs after feedback could be implemented (e.g., the person can no longer change their clothing) or the conversation occurs immediately before an event and there is not enough time to implement feedback or changes (e.g., a person is about to go on stage).	• If someone knew how my mother really felt about me... I would prefer that the person would lie and tell me said good things about me. My mom is deceased now, so nothing could be changed anyway. • If I were out with my friends at a bar and I asked if I looked okay, I would prefer if my friends said yes because if I did not, there would be nothing I could do about it at the bar.	0.83	12.6%	28.4%	19.6%
	Context of Conversation	7. Precedes sacred event	When the truth is hurtful and may upset someone before another unrelated event, such as a wedding, honeymoon, special day.	• Being told there is no bad news before an important event so that the bad news can be postponed.	0.85	4.2%	8.1%	6.0%
		8. In front of others	When the conversation occurs in front of others (and might affect observers' opinions, or embarrass the target)	• If the truth would embarrass me in front of important people	1.00	0.6%	2.2%	1.3%

Note. The tables above reflect the coding scheme for justifications and community standards of deception. Kappa reflects the level of agreement between the two research assistants who coded participant responses, for each coding category. The percentages listed reflect the percentage of participants that listed each justification/community standard in the *Ethics* condition, the *Preferences* condition, and in total (respectively).

3. Study 2

3.1. Choice Options for Each Vignette in Study 2

Emotional fragility vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the employee the truth – the manager should say that the report was not well done
 - Lie to the employee – the manager should say the report was fine

Preferences condition:

- Of the following options, how would you prefer that your manager respond?
 - Tell you the truth – the manager should say that the report was not particularly well done
 - Lie to you – the manager should say the report was fine

Deathbed vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the individual the truth – the spouse should admit that s/he cheated on the individual
 - Lie to the individual – the spouse should not mention that s/he cheated on the individual

Preferences condition:

- Of the following options, how would you prefer that your spouse respond?
 - Tell you the truth – the spouse should admit that s/he cheated on you
 - Lie to you – the spouse should not mention that s/he cheated on you

Ability to understand vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the patient the truth – the doctor should tell the patient s/he is likely to die
 - Lie to the patient – the doctor should not tell the patient that s/he is likely to die

Preferences condition:

- Of the following options, how would you prefer that your doctor respond?
 - Tell you the truth – the doctor should tell you that you are likely to die
 - Lie to you – the doctor should not tell you that you are likely to die

Subjective vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the employee the truth – the colleague should tell the employee that she thinks the employee looks bad in the scarf
 - Lie to the employee – the colleague should tell the employee that she thinks the employee looks fine (or good) in the scarf

Preferences condition:

- Of the following options, how would you prefer that your colleague respond?
 - Tell you the truth – the colleague should tell you that she thinks you look bad in the scarf
 - Lie to you – the colleague should tell you that she thinks you look fine in the scarf

Trivial vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the host the truth – the guest should tell the host that the soup is too salty
 - Lie to the host – the guest should tell the host that the soup is good or fine

Preferences condition:

- Of the following options, how would you prefer that your guest respond?
 - Tell you the truth – the guest should tell you that the soup is too salty
 - Lie to you – the guest should tell you that the soup is good or fine

Uncontrollable (Feature of person) vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the intern the truth – the friend should tell the intern that his stutter decreased the quality of his presentation
 - Lie to the intern – the friend should tell the intern that the presentation was fine (or good)

Preferences condition:

- Of the following options, how would you prefer that your friend respond?
 - Tell you the truth – your friend should tell you that your stutter decreased the quality of your presentation
 - Lie to you – your friend should tell you that the presentation was fine (or good)

Uncontrollable (Time to implement) vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the employee the truth – the colleague should tell the employee that he thinks the suit is inappropriate
 - Lie to the employee – the colleague should tell the employee that he thinks the suit is fine

Preferences condition:

- Of the following options, how would you prefer that your colleague respond?
 - Tell you the truth – the colleague should tell you that he thinks the suit is inappropriate
 - Lie to you – the colleague should tell you that he thinks the suit is fine

Disruption to special moments and event vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the employee the truth – the manager should tell the employee that s/he is getting laid off
 - Lie to the employee – the manager should not tell the employee that s/he is getting laid off

Preferences condition:

- Of the following options, how would you prefer that your manager respond?
 - Tell you the truth – the manager should tell you that you are getting laid off
 - Lie to you – the manager should not tell you that you are getting laid off

The presence of others vignette

Ethics condition:

- Which of the following options is the more ethical response?
 - Tell the employee the truth – the manager should tell the employee that the report was not particularly well done
 - Lie to the employee – the manager should tell the employee that the report was fine

Preferences condition:

- Of the following options, how would you prefer that your manager respond?
 - Tell you the truth – your manager should tell you that your report was not particularly well done
 - Lie to you – your manager should tell you that the report was fine

3.2. Manipulation check questions and results for each vignette

Emotional fragility vignette

Recall question:

- At the time at which the manager could offer feedback, was the employee in a fragile emotional state? [Response options: Yes or No, not particularly]

92% answered correctly

Deathbed vignette

Recall question:

- At the time at which the spouse could admit to cheating, was the individual likely to die soon? [Response options: Yes, probably or No, it was unlikely]

88% answered correctly

Ability to understand vignette

Manipulation check:

- Would the patient in this scenario be able to understand information related to his/her prognosis? [Response options: Yes, probably or No, probably not]

78% answered correctly

Subjective vignette

Recall question:

- Was the colleague's negative opinion of the scarf subjective, or widely shared? [Response options: It was a subjective opinion or It was a widely held opinion]

79% answered correctly

Trivial vignette

Recall question:

- Was whether the soup tasted really good important to the host? [Response options: Yes, it was very important or No, it was probably trivial]

62% answered correctly

Uncontrollable (Feature of person) vignette

Recall question:

- Is the intern able to control his stutter? [Response options: Yes or No]

68% answered correctly

Uncontrollable (Time to implement) vignette

Recall question

- At the time at which the colleague could offer feedback, was the employee able to change his suit? [Response options: Yes or No]

96% answered correctly

Disruption to special moments and event vignette

Recall question:

- Would sharing truthful information in this scenario upset the employee before another important event? [Response options: Yes, definitely or No, not that I know of]

57% answered correctly. Most people in the non-violation condition (81%) still thought the information about potential layoffs would upset the target before another event.

The presence of others vignette

Recall question:

- Would the manager have publicly embarrassed the employee by offering honest feedback? [Response options: Yes, the conversation happened in front of others or No, the conversation happened in private]

82% answered correctly

3.3. Study 2 Choice Results, excluding those who failed manipulation check

Because so many participants did not correctly recall the community standard violation, I also analyzed the main results, examining only those who correctly recalled the violation (see Table S2 below). These effects are notably stronger than the effects reported in the main manuscript, which include all participants.

Table S2. Study 2 Choice Results, including only those who answered recall questions correctly

Community standard (Vignette name)	Control Condition	Community standard Violation Condition	Chi-squared test
1. Emotional fragility	2.5% endorse deception	20.2% endorse deception	$\chi^2 = 18.92, p < .001$
2. Ability to understand	6.5% endorse deception	36.8% endorse deception	$\chi^2 = 42.78, p < .001$
3. Death bed	28.6% endorse deception	64.6% endorse deception	$\chi^2 = 30.21, p < .001$
4. Subjective	39.8% endorse deception	74.5% endorse deception	$\chi^2 = 51.81, p < .001$
5. Trivial	13.8% endorse deception	37.8% endorse deception	$\chi^2 = 18.48, p < .001$
6a. Uncontrollable (Feature of person)	10.4% endorse deception	56.7% endorse deception	$\chi^2 = 61.33, p < .001$
6b. Uncontrollable (Time to implement)	5.5% endorse deception	64.8% endorse deception	$\chi^2 = 98.27, p < .001$
7. Disruption to special moments and events	19.4% endorse deception	54.3% endorse deception	$\chi^2 = 14.66, p < .001$
8. The presence of others	1.5% endorse deception	41.2% endorse deception	$\chi^2 = 117.50, p < .001$

3.4. Mechanism results by vignette

Survey group 1

Perspective was manipulated between subjects. Analyses below reflect results from an OLS regression using *Community standard Violation* (1 = violation, 0 = control), *Perspective* (1 = preferences, 0 = ethics), and their interaction as IVs on either *Immediate Harm* or *Instrumental Value*.

Table S3. Mechanism results for Survey Group 1 (Study 2)

Community standard (Vignette name)	Perspective condition	Control condition		Community standard Violation		Main effect of Community standard Violation	Main effect of Perspective	Standard violation x Perspective
Immediate harm		<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>			
Emotional fragility	Ethics	3.31	1.07	4.38	1.34	$b = 1.07, p < .001$	$b = -.18, p = .399$	$b = -.72, p = .021$
	Preferences	3.12	1.13	3.48	1.47			
	Total	3.22	1.10	3.92	1.47			
Death bed	Ethics	5.37	1.21	5.79	1.26	$b = .42, p = .075$	$b = -.21, p = .377$	$b = .23, p = .486$
	Preferences	5.16	1.31	5.81	1.54			
	Total	5.26	1.26	5.80	1.40			
Uncontrollable (Time to implement)	Ethics	3.39	1.49	4.68	1.48	$b = 1.29, p < .001$	$b = -.55, p = .027$	$b = .14, p = .692$
	Preferences	2.84	1.11	4.27	1.54			
	Total	3.12	1.34	4.48	1.51			
Instrumental value		<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>			
Emotional fragility	Ethics	5.47	1.42	4.30	1.53	$b = -1.17, p < .001$	$b = .31, p = .202$	$b = .33, p = .341$
	Preferences	5.78	1.21	4.94	1.49			
	Total	5.63	1.32	4.62	1.54			
Death bed	Ethics	4.39	1.66	3.29	1.74	$b = -1.10, p < .001$	$b = .31, p = .307$	$b = -.47, p = .259$
	Preferences	4.70	1.68	3.12	1.74			
	Total	4.55	1.67	3.20	1.74			
Uncontrollable (Time to implement)	Ethics	5.38	1.17	3.83	1.53	$b = -1.55, p < .001$	$b = .07, p = .781$	$b = .11, p = .742$
	Preferences	5.45	1.30	4.02	1.58			
	Total	5.41	1.23	3.92	1.55			

Survey groups 2 and 3

Perspective was manipulated within subjects, but participants only answered questions about *Immediate Harm* and *Instrumental Value* once in each scenario. Analyses below reflect results from an OLS regression using *Community standard Violation* (1 = violation, 0 = control) as an IV on either *Immediate Harm* or *Instrumental Value*, clustering standard errors at the participant level.

Table S4. Mechanism results for Survey Groups 2 and 3 (Study 2)

Community standard (Vignette name)	Control condition		Community standard Violation		Main effect of Community standard Violation
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	
Immediate Harm					
Ability to understand	4.14	1.33	5.18	1.05	$b = 1.04, p < .001$
Subjective	4.40	1.35	4.77	1.21	$b = .37, p = .019$
Trivial	3.63	1.30	4.22	1.34	$b = .60, p = .002$
Uncontrollable (Feature of person)	4.03	1.01	4.77	1.35	$b = .74, p < .001$
Disruption to special moments and events	4.28	1.23	5.25	1.30	$b = .97, p < .001$
The presence of others	3.62	1.04	4.44	1.27	$b = .82, p < .001$
Instrumental Value					
Ability to understand	5.53	1.30	4.87	1.61	$b = -.66, p = .002$
Subjective	4.21	1.41	3.26	1.36	$b = -.96, p < .001$
Trivial	5.26	1.30	4.77	1.33	$b = -.49, p = .010$
Uncontrollable (Feature of person)	5.51	1.25	3.71	1.79	$b = -1.80, p < .001$
Disruption to special moments and events	5.14	1.40	4.56	1.53	$b = -.58, p < .006$
The presence of others	6.06	0.94	5.45	1.12	$b = -.62, p < .001$

3.5. Pooled regression analyses for Study 2

As in Study 1, I ran a mixed effects logit model (i.e., melogit) using Stata. I used a fixed-effects approach to control for dependence in observations among vignettes, and a random-effects approach to control for person dependency (i.e., observations nested within persons). Pseudo- R^2 was calculated using the method described in Tjur (2009).

Table S5. Pooled regression analyses for Study 2

Dependent variable = Endorsement of deception; 1 = lie, 0 = tell the truth								
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Intercept	-0.46	-1.20***	-1.14***	-1.17***	-5.05***	-5.51***	-2.56**	-3.31***
Gender ^a	0.26*							
Age	-0.00							
Community standard Violation ^b		2.09***	2.09***	2.13***		1.08***		1.05***
Perspective ^c			-0.11	-0.06				
Perspective x Community standard Violation				-0.09				
Immediate Harm of Truth					1.51***	1.44***	1.00***	0.99***
Instrumental Value of Truth					-0.93***	-0.84***	-1.52***	-1.36***
Immediate Harm x Instrumental Value							0.12***	0.11**
Vignette Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant Random Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3585	3585	3585	3585	3585	3585	3585	3585
R ²	0.24	0.37	0.37	0.37	0.62	0.63	0.62	0.63

Note. *, **, *** denote significance at $p \leq .05$, $< .01$, and $< .001$ respectively

^aGender is coded as 1 = female, 0 = male

^bViolation is coded as 1 = community standard violation, 0 = no standard violation

^cPerspective is coded as 1 = target, 0 = observer

4. Study 3

4.1. Choice Options for Each Vignette in Study 3

Ability to understand vignette

In the course of this conversation, which of the following options is the more ethical response?

- Tell Jeff the truth - Jeff's caregiver should tell him about his dead daughter
- Lie to Jeff - Jeff's caregiver should not tell him about his dead daughter

Uncontrollable (Time to implement) vignette

In the course of this conversation, which of the following options is the more ethical response?

- Tell Jeff the truth - Jeff's friend should tell Jeff about the errors
- Lie to Jeff - Jeff's friend should not tell Jeff about the errors

The presence of others vignette

In the course of this conversation, which of the following options is the more ethical response?

- Tell Jeff the truth - Jeff's friend should tell Jeff his presentation went poorly
- Lie to Jeff - Jeff's friend should not tell Jeff his presentation went poorly

4.2. Main analyses in Study 3, Split by sample

Study 3 was run with two different samples. The purpose of recruiting two samples was to ensure robustness across different populations. Below are the main regressions from Study 3, split by sample. The results are qualitatively identical among both MTurk and student (university behavioral lab) populations.

Table S6. Main analyses in Study 3, Split by sample

Dependent variable = Endorsement of deception; 1 = lie, 0 = tell the truth

[illegible]

MTurk sample										
Model number:	1-MT	2-MT	3-MT	4-MT	5-MT	6-MT	7-MT	8-MT	9-MT	10-MT
Intercept	-1.54***	-3.13***	-3.16***	-3.33***	.45	-.53	5.70	4.61	1.75	.57
Gender ^a	0.52*									
Age	-.02*									
Community standard Violation ^b		1.58***	1.58***	1.83***		1.11**		1.11**		1.16**
Communicator ^c			0.07	0.39						
Target ^d			0.02	0.15						
Liar x Community standard Violation				-0.50						
Target x Community standard Violation				-0.21						
Immediate Harm of Truth					.93***	.93***	-0.004	0.02	.77***	.74***
Instrumental Value of Truth					-1.41***	-1.39***	-2.70**	-2.65**	-1.01***	-.94***
Imm Harm x Instr Value							0.23	0.22		
Self-interest									0.29 ⁺	0.29
Autonomy									-0.06	-.04
Probability of Detection									-0.06	0.001
Societal Harm									0.02	0.02
Moral Duty									-.76***	-.79***
Vignette Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Participant Random Effect	No ^e	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	408	408	408	408	408	408	408	408	408	408

Note. ⁺, *, **, *** denote significance at $p \leq .10$, .05, .01, and .001 respectively

^aGender is coded as 1 = female, 0 = male; ^bViolation is coded as 1 = community standard violation, 0 = no standard violation; ^cCommunicator is coded as 0 = target or observer perspective, 1 = communicator perspective; ^dTarget is coded as 0 = communicator or observer perspective, 1 = target perspective; ^eTo account for multiple observations per participant, Model 1 used a clustered standard error approach rather than participant random effect, because an LR test confirmed that the random effects model was not appropriate.

4.3. Correlation table for Study 3 DVs

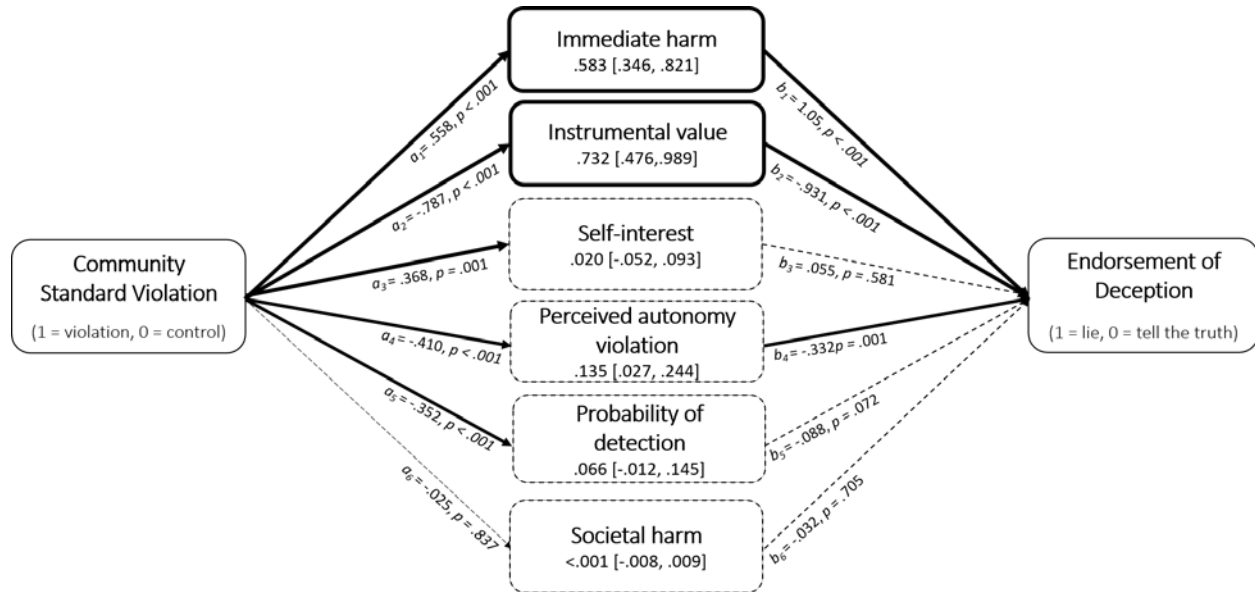
Variables	Variable number								
	1)	2)	3)	4)	5)	6)	7)	8)	9)
1) Endorsement of deception	1.00								
2) Community Standard Violation	0.27 <i>p</i> < .001	1.00							
3) Immediate Harm of Truth	0.41 <i>p</i> < .001	0.19 <i>p</i> < .001	1.00						
4) Instrumental Value of Truth	-0.45 <i>p</i> < .001	-0.27 <i>p</i> < .001	-0.27 <i>p</i> < .001	1.00					
5) Self-interest	0.20 <i>p</i> < .001	0.12 <i>p</i> < .001	0.53 <i>p</i> < .001	-0.10 <i>p</i> < .001	1.00				
6) Probability of Detection	-0.21 <i>p</i> < .001	-0.14 <i>p</i> < .001	-0.03 <i>p</i> = .322	0.35 <i>p</i> < .001	0.05 <i>p</i> = .162	1.00			
7) Autonomy	-0.27 <i>p</i> < .001	-0.13 <i>p</i> < .001	-0.05 <i>p</i> = .158	0.36 <i>p</i> < .001	-0.04 <i>p</i> = .265	0.22 <i>p</i> < .001	1.00		
8) Societal Harm	-0.18 <i>p</i> < .001	-0.01 <i>p</i> = .860	-0.13 <i>p</i> < .001	0.17 <i>p</i> < .001	-0.16 <i>p</i> < .001	0.02 <i>p</i> = .662	0.43 <i>p</i> < .001	1.00	
9) Moral Duty	-0.48 <i>p</i> < .001	-0.17 <i>p</i> < .001	-0.16 <i>p</i> < .001	0.50 <i>p</i> < .001	-0.02 <i>p</i> = .644	0.34 <i>p</i> < .001	0.46 <i>p</i> < .001	0.20 <i>p</i> < .001	1.00

4.4 Mediation analyses without moral duty

The moral duty item is conceptually quite similar to the preference for deception. Therefore, I reran the mediation analysis without this item. As in the main manuscript, I performed a multilevel logistic mediation model using Stata's `gsem` function. A random intercept was included in each equation at the participant level to account for within-person dependence of observations, and fixed effects were included to control for scenario and perspective. The results of this analysis are below.

Although the significance of some of the effects did change, perceptions of immediate harm and instrumental value continued to have significantly larger indirect effects than any other potential mediators.

Figure S1. Mediation analyses in Study 3 without moral duty item



5. Study 4

5.1. Additional measures in Study 4a

In Study 4a, I also measured *the target's anticipated anger* (If [insert target name] found out you had lied to them in this situation, how angry do you think they would be?; 1 = not at all, 7 = extremely) and the *target's anticipated understanding* (If [insert target name] found out you had lied to them in this situation, to what extent do you think they would believe you had good intentions when lying to them?; 1 = not at all, 7 = extremely).

The results of these measures are shown on the next pages, in Tables S7 and S8. For both Tables, ⁺, *, **, *** denote significance at $p \leq .10$, .05, .01, and .001 respectively.

Target's anticipated anger. Relative to judgments of immediate harm, judgments of instrumental value were more predictive of communicators' beliefs that targets would be angry upon learning they were deceived.

Judgments of self-interest, societal harm, and the probability of detection also predicted judgments of a target's anticipated anger. However, perceptions of instrumental value remained significant after controlling for these judgments.

Target's anticipated understanding. Judgments of immediate harm were more predictive of communicators' beliefs that targets would understand their intentions than were judgments of instrumental value. Judgments of self-interest, societal harm, and the probability of detection also predicted judgments of a target's anticipated understanding.

Table S7. OLS Regression on *target's anticipated anger* in Study 4a

Independent variables	Model 1		Model 2		Model 3		Model 4		Model 5	
	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>
Immediate Harm	0.058	0.056	0.227	0.138	0.012	0.054	0.149	0.127	0.153	0.128
Instrumental Value	0.451***	0.065	0.615***	0.138	0.272***	0.072	0.407**	0.134	0.399**	0.135
Immediate Harm x Instrumental Value			-0.048	0.036			-0.039	0.033	-0.039	0.033
Self-interest					0.334***	0.059	0.332***	0.059	0.337***	0.060
Societal Harm					0.235**	0.074	0.234**	0.074	0.228**	0.075
Perceived Autonomy Violation					0.095	0.066	0.100	0.066	0.102	0.066
Probability of Detection					0.138*	0.054	0.133*	0.054	0.131*	0.054
Decision to Lie									-0.088	0.198
Constant	2.229***	0.294	1.671***	0.508	0.141	0.405	-0.296	0.546	-0.241	0.561
R ²	0.153		0.159		0.294		0.297		0.298	

Table S8. OLS Regression on *target's anticipated understanding* in Study 4a

Independent variables	Model 1		Model 2		Model 3		Model 4		Model 5	
	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>
Immediate Harm	0.435***	0.057	0.376**	0.141	0.436***	0.059	0.400**	0.14	0.387**	0.14
Instrumental Value	-0.267***	0.067	-0.325*	0.141	-0.188*	0.079	-0.223	0.147	-0.195	0.148
Immediate Harm x Instrumental Value			0.017	0.036			0.010	0.036	0.008	0.036
Self-interest					-0.130*	0.065	-0.130*	0.065	-0.147*	0.066
Societal Harm					-0.181*	0.081	-0.180*	0.081	-0.158+	0.083
Perceived Autonomy Violation					0.056	0.072	0.055	0.072	0.048	0.072
Probability of Detection					-0.079	0.059	-0.078	0.059	-0.070	0.060
Decision to Lie									0.321	0.218
Constant	3.460***	0.300	3.658***	0.52	4.386***	0.445	4.501***	0.602	4.301***	0.616
R ²	0.182		0.183		0.211		0.211		0.217	

In Study 4a, I also asked the following three exploratory questions:

1. To what extent did you have time to plan what you were going to say, before actually telling the truth or lying in the situation above? (1 = not at all, 7 = extremely)
2. To what extent did you have to communicate on the spot, with little preparation, in the situation above? (1 = not at all, 7 = extremely)
3. Did this conversation (in which you lied or told the truth) happen face-to-face, over email or text, or over the phone? (Choices: Face-to-face, Over email or text. Over the phone, Other)

The first two items were combined into a single measure, with the second item reverse-scored ("Time to plan," $r = .64$, $p < .001$). The descriptive results of this measure are depicted in Table S9. The descriptive results of the medium question (question 3 above) are depicted in Table S10.

5.2. Descriptive Statistics in Studies 4a

Table S9. Descriptive Statistics of all scale measures in Studies 4a, split by the decision to lie

	Ultimate decision participant made when facing temptation to lie							
	Lie		Truth		Other		Overall	
Study 4a: Communicators	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Moral judgment of lie	4.31	1.67	2.93	1.64	3.88	1.66	3.81	1.77
Target's anticipated anger	3.93	1.78	4.56	1.65	3.50	1.90	4.09	1.78
Target's anticipated understanding	4.28	1.79	3.13	1.74	4.44	1.83	3.92	1.85
Immediate harm	3.64	1.74	2.73	1.53	3.46	1.86	3.32	1.73
Instrumental value	3.35	1.35	4.38	1.49	3.45	1.55	3.70	1.49
Self-interest	5.28	1.42	4.51	1.62	4.92	1.48	4.99	1.53
Societal Harm	2.19	1.17	3.10	1.61	2.70	1.49	2.54	1.42
Perceived Autonomy Violation	2.19	1.48	2.56	1.83	2.09	1.44	2.30	1.61
Probability of Detection	2.77	1.67	3.41	1.81	3.34	1.84	3.04	1.76
Time to Plan	3.37	1.89	3.28	1.69	3.45	2.07	3.35	1.84
	<i>n</i> = 167		<i>n</i> = 97		<i>n</i> = 32		<i>N</i> =296	

Table S10. Communication medium results in Studies 4a

	Ultimate decision participant made when facing temptation to lie			
	Lie	Truth	Other	Overall
Study S4: Medical practitioners	<i>Freq.</i>	<i>Freq.</i>	<i>Freq.</i>	<i>Freq.</i>
Face to face	92	72	19	183
Email or text	39	9	5	53
Phone	30	10	4	44
Other	6	6	4	16

5.3. Results of Studies 4a, Split by relationship type

Participants in Study 4a indicated the nature of their relationship with the target. Below, I explore whether perceptions of immediate harm and instrumental value differentially impacted the decision to lie across relationships. Although these analyses are exploratory, they suggest the interesting possibility that decisions to lie to friends and family members hinge more on perceptions of instrumental value than immediate harm, whereas decisions to lie to strangers and significant others hinge more on perceptions of immediate harm than instrumental value.

Table S11. The effects of Immediate Harm and Instrumental Value on the Decision to Lie Across Relationships

Relationship	Independent variables	B	S.E.	Sig.
Friend, <i>n</i> = 66	Immediateharm	0.357	0.432	0.409
	Instrumentalvalue	-0.912	0.533	0.087
	ImmHarmxInstValue	0.040	0.111	0.721
	Constant	2.065	1.626	0.204
Family Member, <i>n</i> = 90	Immediateharm	-0.323	0.349	0.355
	Instrumentalvalue	-1.196	0.452	0.008
	ImmHarmxInstValue	0.178	0.100	0.075
	Constant	3.143	1.557	0.044
Coworker, <i>n</i> = 22	Immediateharm	1.117	1.469	0.447
	Instrumentalvalue	-0.477	0.899	0.595
	ImmHarmxInstValue	-0.160	0.300	0.593
	Constant	1.506	3.536	0.670
Stranger, <i>n</i> = 32	Immediateharm	3.757	2.082	0.071
	Instrumentalvalue	1.287	0.858	0.133
	ImmHarmxInstValue	-0.811	0.443	0.067
	Constant	-6.195	3.765	0.100
Significant Other, <i>n</i> = 42	Immediateharm	1.936	0.801	0.016
	Instrumentalvalue	1.105	0.815	0.175
	ImmHarmxInstValue	-0.369	0.196	0.061
	Constant	-5.274	2.803	0.060
Other, <i>n</i> = 41	Immediateharm	0.937	1.062	0.377
	Instrumentalvalue	-0.730	0.562	0.194
	ImmHarmxInstValue	-0.051	0.223	0.819
	Constant	1.228	2.227	0.581

6. Details of additional studies

6.1. Study S1: Is lying always wrong? pilot

Study S1 was another open-ended study that explored people's general beliefs about the ethicality of lying, as well as their preferences for being lied to.

Method and Results

Participants. I intended to recruit 100 adults from Prolific (an online research platform, <https://prolific.ac/>) and ended up with a final sample of 99 adults (48% female; $M_{\text{age}} = 32$ years).

Procedure. Participants were randomly assigned to answer questions about when lying is ethical or when they would want to be deceived. In the *Ethics* condition, participants were asked to select the statement they most agreed with (choices: "Lying is always wrong", "Lying is usually wrong", "Lying is often wrong", "Lying is sometimes wrong", "Lying is occasionally wrong", "Lying is rarely wrong", or "Lying is never wrong"). Then they were asked to explain their answer (free response).

In the *Preferences* condition, participants were also asked to select the statement they most agreed with (choices: "I would never want someone to lie to me", "I would rarely want someone to lie to me", "I would occasionally want someone to lie to me", "I would sometimes want someone to lie to me", "I would often want someone to lie to me", "I would usually want someone to lie to me", or "I would always want someone to lie to me"), and to explain their answer (free response). Participants then provided demographic information and were dismissed.

The goal of this study was to simply explore the proportion of people who took an absolute stance against deception in the *Ethics* and *Preferences* condition (by selecting "Lying is always wrong" or "I would never want someone to lie to me") or endorsed deception at some level. Below (in Table S12) are the frequencies with which participants selected each choice.

Table S12. Frequency of different levels of endorsement for deception

Preferences condition			Ethics condition		
Statement	Count	%	Statement	Count	%
Total absolute stance:			Total absolute stance:		
I would never want someone to lie to me	18	36%	Lying is always wrong	3	6%
Total-non-absolute stance	31	62%	Total-non-absolute stance	46	94%
I would rarely want someone to lie to me	22	44%	Lying is usually wrong	28	57%
I would occasionally want someone to lie to me	7	14%	Lying is often wrong	12	24%
I would sometimes want someone to lie to me	2	4%	Lying is sometimes wrong	6	12%
I would often want someone to lie to me	1	2%	Lying is occasionally wrong	0	0%
I would usually want someone to lie to me	0	0%	Lying is rarely wrong	0	0%
I would always want someone to lie to me	0	0%	Lying is never wrong	0	0%
Total-total	50		Total-total	49	

6.2. Study S2: Lying in everyday life pilot

In Study S2, I examined how people judged their own lies.

Method and Results

Participants. I recruited 100 participants from a U.S. university laboratory (50% female; $M_{\text{age}} = 33$ years).

Procedure. In a randomly assigned order, participants were asked to describe the three most recent lies they told (recent lies condition), and one lie they told that was ethical (ethics condition). Regardless of order, participants were asked to provide new lies for each response; therefore, all participants provided four lies in total.

Recent lies condition. For each of their recent lies, participants were asked to rate their agreement with the following statements (1 = strongly disagree, 7 = strongly agree): “In this situation, lying was ethical”, “In this situation, I would want to be lied to,” “I told this lie because I wanted to harm an enemy,” “I told this lie because I thought it would benefit me”, “I told this lie because it was trivial”, “I told this lie because I didn't think anyone would find out”, “I told this lie because I thought it would benefit society”, “I told this lie because I wanted to benefit the person being lied to”, and “I told this lie because I didn't think the person I was talking to deserved to know the truth.”

I expected that participants' judgments of a lie's ethicality would be predicted by beliefs about whether the lie benefitted the target of the lie (more so than any other justification). Indeed, this was the case. A regression analysis (with standard errors clustered at the participant-level) revealed that judgments of ethicality were predicted by beliefs about whether the lie benefitted the target of the lie ($b = .277, p < .001$), more so than any other justification (harm an enemy: $b = -.179, p = .133$; benefit me: $b = -.004, p = .947$; trivial: $b = .133, p = .051$; nobody would find out: $b = .058, p = .447$; benefit society: $b = .108, p = .139$; deserve to know the truth: $b = .180, p = .009$).

Ethics condition. Participants were asked to indicate why their ethical lie was ethical. Specifically, they indicated their agreement with the following justifications (1 = strongly disagree, 7 = strongly agree): “This lie is ethical because the person lied to didn't deserve to know the truth,” “This lie is ethical because it benefitted society,” “This lie is ethical because it was told to an enemy”, “This lie is ethical because nobody would find out about it”, “This lie is ethical because it was trivial”, “This lie is ethical because it benefitted me”, and “This lie is ethical because it benefitted the person being lied to.” Participants also rated their agreement with the statement: “I would want to be told this lie.”

I expected that participants would indicate strongest agreement with the justification, “This lie is ethical because it benefitted the person being lied to,” consistent with the findings of the inductive study (i.e., people primarily care about the consequences of lying for the target), and that agreement with this justification would be correlated with the statement, “I would want to be told

this lie.” Indeed, there was strongest agreement with the target-benefit justification (see Table S13).

A regression analysis (with standard errors, clustered at the participant-level) revealed that the desire to be told the lie was predicted by beliefs about whether the lie benefitted the target of the lie ($b = .266, p = .013$), more so than any other justification (harm an enemy: $b = -.082, p = .565$; benefit me: $b = -.037, p = .717$; trivial: $b = .122, p = .221$; nobody would find out: $b = .022, p = .868$; benefit society: $b = -.024, p = .826$; deserve to know the truth: $b = .112, p = .053$).

Table S13. Descriptive statistics for justifications of deception

Recent Lies Condition	<i>Mean</i>	<i>SD</i>
"In this situation, lying was ethical"	3.90	1.94
"In this situation, I would want to be lied to"	3.29	1.74
"I told this lie because I didn't think anyone would find out"	4.58	2.10
"I told this lie because I didn't think the person I was talking to deserved to know the truth."	3.41	2.12
"I told this lie because I thought it would benefit society"	2.38	1.69
"I told this lie because it was trivial"	4.36	2.02
"I told this lie because I thought it would benefit me"	4.90	2.00
"I told this lie because I wanted to harm an enemy"	1.64	1.35
"I told this lie because I wanted to benefit the person being lied to"	3.59	2.22
Ethics condition		
"In this situation, I would want to be lied to"	3.79	1.98
"This lie is ethical because nobody would find out about it"	3.57	1.99
"This lie is ethical because the person lied to didn't deserve to know the truth"	3.34	2.04
"This lie is ethical because it benefitted society"	3.26	2.07
"This lie is ethical because it was trivial"	4.08	2.20
"This lie is ethical because it benefitted me"	3.62	2.19
"This lie is ethical because it was told to an enemy"	1.64	1.31
"This lie is ethical because it benefitted the person being lied to"	4.86	2.08

6.3. Study S3: Targets' reactions to lies in everyday life

Study S3 was similar to Study 4a, except it focused on the perspective of the target, rather than the communicator.

Method. As in Study 4a, I recruited 300 adults from a U.S. representative sample via Prolific Academic (<https://www.prolific.co/>). I ended up with a final sample of 285 participants who completed the entire survey (51% female; $M_{\text{age}} = 45$ years; $M_{\text{work experience}} = 22$ years).

As in Study 4a, participants first read a passage explaining that people often face situations in which they are tempted to lie. Then, participants were asked to “Think of the last time you **learned that someone lied to you.**” Participants then described the situation (free response) and answered a series of questions.

Participants then answered all the same attitudinal questions that participants had answered in Study 4a, except the question pertaining to the decision to lie (Moral judgment of deception, Immediate harm & Instrumental value, Self-interest, Societal harm, Perceived autonomy violation, Probability of detection, as well as Target's anticipated anger and Target's anticipated understanding, see SOM 5.1). However, these questions were written from the perspective of the target, rather than the communicator. For example, “To what extent would honesty cause pain to [insert target name] at the moment of communication?” became “To what extent would honesty have caused pain to you at the moment of communication?” After participants submitted their responses, I collected demographic information.

Results

Table S14 depicts the descriptive statistics associated with each measure.

Table S14. Descriptive Statistics of all scale measures in Study S3

Study S3: Target's Perspective	α/r	M	SD
Moral judgment of lie	N/A	2.01	1.27
Target's anger	N/A	4.95	1.88
Target's understanding	N/A	2.56	1.89
Immediate harm	$\alpha = .77$	2.31	1.28
Instrumental value	$\alpha = .76$	4.65	1.53
Self-interest	$\alpha = .71$	4.56	1.71
Societal Harm	$r = .48$	3.61	1.70
Perceived Autonomy Violation	N/A	3.42	2.05
Probability of Detection	N/A	5.13	1.90

Table S15 depicts the results of OLS regressions predicting moral judgment (Panel A), Target's anger (Panel B), and Target's understanding (Panel C) using perceptions of unnecessary harm, perceptions of self-interest, societal harm, autonomy violation, and probability of detection as independent variables.

Table S15. OLS Regression on perceived ethicality of lying (moral judgment), target's anger, and target's understanding in Study S3

Panel A: OLS regression on Perceived Ethicality of Lying

	Model A1		Model A2		Model A3		Model A4	
	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>
Immediate Harm	0.142*	0.058	0.243	0.221	0.153**	0.056	0.276	0.209
Instrumental Value	-0.210***	0.049	-0.170+	0.097	-0.036	0.055	0.011	0.096
Imm. H x Instr. V			-0.021	0.045			-0.026	0.042
Self-interest					-0.086*	0.042	-0.086*	0.042
Societal Harm					-0.153**	0.049	-0.153**	0.049
Autonomy Violation					-0.103*	0.045	-0.101*	0.045
Prob. of Detection					-0.108**	0.037	-0.112**	0.038
Constant	2.664***	0.251	2.481***	0.462	3.682***	0.322	3.474***	0.469
R ²	0.069		0.070		0.199		0.200	

Panel B: OLS regression on Target's Anger

	Model B1		Model B2		Model B3		Model B4	
	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>
Immediate Harm	0.127	0.078	0.376	0.298	0.109	0.073	0.332	0.273
Instrumental Value	0.557***	0.066	0.655***	0.131	0.273***	0.072	0.359**	0.125
Imm. H x Instr. V			-0.052	0.060			-0.047	0.055
Self-interest					0.146**	0.055	0.146**	0.055
Societal Harm					0.145*	0.064	0.146*	0.064
Autonomy Violation					0.226***	0.059	0.229***	0.059
Prob. of Detection					0.190***	0.049	0.183***	0.049
Constant	2.062***	0.338	1.610**	0.621	0.489	0.421	0.111	0.613
R ²	0.227		0.230		0.373		0.374	

Panel C: OLS regression on Target's Understanding

	Model C1		Model C2		Model C3		Model C4	
	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>	<i>b</i>	<i>SE</i>
Immediate Harm	0.299***	0.087	0.837*	0.330	0.295***	0.088	0.954**	0.324
Instrumental Value	-0.186*	0.073	0.025	0.145	-0.059	0.087	0.196	0.149
Imm. H x Instr. V			-0.112+	0.066			-0.139*	0.066
Self-interest					-0.115+	0.065	-0.114+	0.065
Societal Harm					-0.072	0.076	-0.070	0.076
Autonomy Violation					-0.066	0.071	-0.055	0.070
Prob. of Detection					-0.173**	0.058	-0.192***	0.059
Constant	2.738***	0.375	1.762*	0.687	4.053***	0.504	2.936***	0.728
R ²	0.052		0.061		0.109		0.123	

Note. +, *, **, *** denote significance at $p \leq .10$, .05, .01, and .001 respectively.

Moral judgment of deception. Judgments of instrumental value and immediate harm both predicted targets' moral judgment of deception when entered together in a model (see Model A1). However, perceptions of self-interest, societal harm, autonomy violations, and the probability of detection also predicted targets' moral judgment, and controlling for these perceptions attenuated the effect of instrumental value on moral judgment (see Model A3).

Target's anger. Only judgments of instrumental value (but not immediate harm) predicted targets' anger about being deceived (see Models B1-B4). Perceptions of self-interest, societal harm, autonomy violations, and the probability of detection also predicted targets' anger. However, controlling for these perceptions did not attenuate the effect of instrumental value on targets' anger (see Models B3 and B4).

Target's understanding. Judgments of instrumental value and immediate harm both predicted targets' belief that the communicator had good intentions when lying to them (see Model C1). Perceptions of self-interest and the probability of detection also predicted targets' understanding. Controlling for these perceptions attenuated the effect of instrumental value on targets' understanding but did not attenuate the effect of immediate harm on targets' understanding (see Models C3 and C4).

Taken together, these results suggest that perceptions of immediate harm may be more central for targets' moral judgments and the belief that the communicator had good intentions, whereas perceptions of instrumental value may be more central to target's anger about being deceived.

7. Relationship between immediate harm and instrumental value

I predicted that *deception is most likely to be seen as more ethical than honesty* when honesty is low in instrumental value *and* high in immediate harm. I test this prediction by testing whether the endorsement of deception in the low instrumental value/high immediate harm cell is greater than all other cells.

Furthermore, I predicted that there will only be *consensus* that deception is more ethical than honesty when honesty is low in instrumental value and high in immediate harm. I test this prediction by testing whether the percentage of participants who believe deception is more ethical than honesty is greater than 50% in the low instrumental value/high immediate harm cell.

It is important to note that these results could be supported by two main effects, or by an interaction between instrumental value and immediate harm. To illustrate this, I created new figures depicting the data from Study 1 (see Figure S2 on the following page). Across all perspectives (communicators, third parties, and targets), the endorsement of deception in the low instrumental value/high immediate harm cell is significantly greater than the endorsement of deception in all other cells ($ps < .01$) and greater than 50% (though not significantly, among targets). However, there is only a significant high immediate harm x low instrumental value interaction among third parties.

Figure S2. Clarifying the relationship between immediate harm and instrumental value

Among communicators, there is a main effect of Immediate harm ($p < .001$), a main effect of Instrumental value ($p < .001$), and no interaction. Among third parties, there is a main effect of Immediate harm ($p < .001$), a main effect of Instrumental value ($p < .001$), and a significant interaction ($p < .01$). Among targets, there is a main effect of Immediate harm ($p < .001$), a main effect of Instrumental value ($p < .001$), and no interaction.

