HUMAN
COMMUNICATION RESEARCH

Human Communication Research ISSN 0360-3989

ORIGINAL ARTICLE

The Prevalence of Lying in America: Three Studies of Self-Reported Lies

Kim B. Serota, Timothy R. Levine, & Franklin J. Boster

Department of Communication, Michigan State University, East Lansing, MI 48823, USA

This study addresses the frequency and the distribution of reported lying in the adult population. A national survey asked 1,000 U.S. adults to report the number of lies told in a 24-hour period. Sixty percent of subjects report telling no lies at all, and almost half of all lies are told by only 5% of subjects; thus, prevalence varies widely and most reported lies are told by a few prolific liars. The pattern is replicated in a reanalysis of previously published research and with a student sample. Substantial individual differences in lying behavior have implications for the generality of truth–lie base rates in deception detection experiments. Explanations concerning the nature of lying and methods for detecting lies need to account for this variation.

doi:10.1111/j.1468-2958.2009.01366.x

Humans are ambivalent about deception. On one hand, virtually all human cultures have some prohibition against lying. On the other hand, the ability to deceive well may be essential for polite interaction and, at times, self-preservation. Considerable research exists on the topic of deception, yet surprisingly little is known about the base prevalence of deception. Instead, much of this research has relied on untested assumptions and anecdotal evidence or on a few studies with small and nonrepresentative samples.

The dearth of deception prevalence research is a symptom of a broader systemic concern regarding research in the social sciences. Asch (1952, reprinted 1987) observed that "before we inquire into origins and functional relations, it is necessary to know the thing we are trying to explain." Influenced by Asch, Rozin (2001) argued that social scientific research often emphasizes experimental studies and formal hypothesis testing to the exclusion of more basic descriptive work. In line with Rozin's critique, more than 30 years of experimental detection research has proceeded without much attention to the basic nature of the phenomena itself. We believe that inquiry into deception and related behaviors associated with deception detection requires basic descriptive research examining the extent and distribution of deceptive communication in the population. In the extensive literature on deception,

Corresponding author: Timothy R. Levine; e-mail: levinet@msu.edu

the question of prevalence remains without a clear, well-documented answer. Thus, our research investigates reports of how often people lie.

In order to study the prevalence of lying, it is necessary to consider what constitutes a lie. Simply and broadly put, lying occurs when a communicator seeks knowingly and intentionally to mislead others. Ford, King, and Hollender (1988) suggest the "consciousness of falsity" to distinguish "normal" lies from pathological ones. Thus, it is not sufficient that something is false for it to be a lie; it is the intent that distinguishes the lie. As Bok (1999) observes: "The moral question of whether you are lying or not is not *settled* by establishing the truth or falsity of what you say. In order to settle this question, we must know whether you *intend your statement to mislead*" (p. 6).

Bok (1999) argues for the principle of veracity that involves a moral asymmetry between honesty and lies. Lying requires justification, whereas truth telling does not. Given prohibitions against deceit, people may try to avoid situations in which there is pressure to lie. Research finds that this principle of veracity guides everyday communication and people consider using deception only when the truth is problematic (Levine, Kim, & Hamel, 2009). But this tells us about the situations in which people lie and not how often people lie.

Despite this moral asymmetry, most deception research has presumed the ubiquity of lying and moved past the question of frequency to focus on the behavioral correlates of lying or lie detection. The frequency question remains mostly unanswered. A notable exception, and the best and most cited prevalence research, however, is the 1996 diary study of lying in everyday life by DePaulo, Kashy, Kirkendol, Wyer, and Epstein (1996). Using two small samples, students and recruited members of the local community, DePaulo et al. reported the mean number of lies per day as 1.96 ($SD = 1.63$, $N = 77$) for the students and 0.97 ($SD = 0.98$, $N = 70$) for the subsequent nonstudent sample. Importantly, the aim of the second sample was to replicate findings regarding the nature and reasons for lying with a different but not necessarily representative sample of the population. Nonetheless, a brief synopsis of this study in *Psychology Today* ("The Real Truth About Lying," 1996) reported that DePaulo conducted research to answer the question "how often do people lie . . . ?" and in many subsequent research reports the finding *that people tell one to two lies per day* has been reified.

More recently, two smaller and lesser-known studies have sought to replicate and extend elements of the DePaulo et al. (1996) diary study. Hancock, Thom-Santelli, and Ritchie (2004) examined differences between reports of face-to-face lies and reports of lying through computer-mediated communication. Results from a student sample yielded an average of 1.58 lies per day ($SD = 1.02$, $N = 28$) and a significant difference for the rates of lying (lies as a proportion of interactions) between face-to-face, telephone, instant message, and e-mail interactions; the highest rates occurred during phone conversations and the lowest rates with e-mail. George and Robb (2008) replaced the pencil-and-paper diary with a personal digital assistant (PDA). They report fewer lies per day; $M = 0.59$ ($SD = 0.37$, $N = 25$) with the 10-minute

definition of interaction used by DePaulo et al. and $M = 0.90$ ($SD = 0.54$, $N = 25$) in a second study shortening the interaction definition to 5 minutes (increasing the number of reporting opportunities). Thus, the current literature provides estimates ranging from 0.59 to 1.96 lies per day. Variation in estimates from study to study is expected due to small sample sizes and large standard deviations.

Other broad estimates of prevarication cited by deception researchers have come from nonacademic sources. In a poll conducted for the book *The Day America Told the Truth* (Patterson & Kim, 1991), 90% of the subjects admitted being deceitful about a list of subjects, the most common being true feelings, income, accomplishments, sex life, and age. In a *Reader's Digest* poll (Kalish, 2004) of 2,861 of the magazine's readers, 93% reported one or more kinds of dishonesty at work or school, 93% reported one or more dishonest acts in the market place, and 96% reported lying to or committing other dishonest acts toward family and friends.

Some studies report on various facets of prevalence but provide limited insight into the overall extent of lying because they deal with specific situations such as lying by job applicants, students lying to parents, or lying about spousal infidelity. Levashina and Campion (2007) found that 90% of undergraduate job candidates used some form of deception during job interviews; however, the distinction between impression management and outright lies is often blurred and their report notes that behaviors that are semantically close to lies are difficult to confirm. They estimate actual lies occur in the wide range of 28–75% of job interviews. Jensen, Arnett, Feldman, and Cauffman (2004) examined lying to parents by adolescents and young adults, and quantified the extent of lying by topic over the course of a year. This study found that 82% of all students reported lying to their parents on at least one of six topical issues (money, alcohol/drugs, friends, dating, parties, and sex) with the mean incidence of lying ranging from 0.6 to 2.4 lies depending on the issue. Much of the research that seeks to quantify lying behavior is concentrated in the area of relational communication. Cochran and Mays (1990) studied dating dishonesty among college students and found that 60% of women claimed to have been lied to in order to obtain sex, whereas 34% of the men in the study admitted lying to obtain sex. Knox, Schacht, Holt, and Turner (1993) found that 92% of students (when given the opportunity to report anonymously) admitted to lying to a current or potential sexual partner.

It is not difficult to understand why many scholars believe lying is a frequent event. Life experiences and anecdotal evidence encourage acceptance of the proposition. A typical research report discussion statement illustrates this view: ''Lying is ubiquitous and comes in many forms, from cherished beliefs about Santa Claus to the self-deception commonly encountered in the treatment situation'' (Tosone, 2006).

General acceptance of the ubiquity assumption has implications for studies on lying and deception detection. If everyone lies and lying is an everyday occurrence for most people, this would suggest that individual differences should not have much influence on the identification of lying behaviors. If this is the case, it should be possible to understand the nature of lying and deceptive behavior and find ways to

detect it by studying anyone telling lies. For example, this presumption is recurrent in studies that look for regularities in nonverbal cues, microexpressions, and the leakage of emotions. Individual differences are typically considered of less relevance than situational considerations such as whether the lie is a minor everyday lie or if the lie is a consequential, high-stakes lie. If, on the other hand, base rates for lying (the frequency with which one lies or the ratio of truths to lies) vary across groups of individuals or if the ubiquitous average masks variation that is not normally distributed in the population, researchers looking into the nature of the phenomenon need to take into account this variation. Research designs and sample selection procedures ought to control for this variation, or they should examine the nature of the variation itself. Consistent with this second possibility, recent meta-analysis suggests substantial individual differences in people's ability to lie convincingly (Bond & DePaulo, 2008).

Study 1

Examination of the literature reveals few attempts to document the extent to which everyday lies occur. The few studies that offer behavioral rates have performed so as an adjunct to the main objectives of the research and the rates obtained are restricted to the specific conditions of the studies. This situation is exemplified by the DePaulo et al. (1996) studies. Although lie frequency is among the interests that prompted their research, most of the design and analysis is devoted to the topics of what people lie about, to whom they lie, and what motivates them to lie. DePaulo et al. noted that their observations of lying frequency are based on students and an adult sample that was chosen not for representativeness but, instead, to provide a dissimilar sample in order to determine whether or not the results from the student sample could be replicated. Still, DePaulo et al. reported: ''Participants in the community study, on the average, told a lie everyday; participants in the college student study told two.'' So despite being based on data that were (correctly) noted by its authors as lacking generalizability, this research report has become the standard reference for the prevalence of everyday lies in the deception literature. The goal of Study 1 is to test this claim by obtaining data from a large cross-section of the adult population.

Method

Participants and design

In order to examine the proposition that most people tell one to two lies per day with projectable data, an Internet survey of 1,000 American adults (18 years of age or older) was conducted using the Synovate *eNation* omnibus panel.[1] The omnibus panel is a commercial survey research tool used for daily, multiclient studies and approximates a nationally representative sample with some limitations. Panelists are recruited into a pool of more than 1 million subject households using banner advertisements, mailing lists, and related procedures to promote participation; subjects must formally opt-in to confirm awareness that they are participating in research and consent to participation. Each survey day, a new sample ($N = 5{,}500$) is drawn from the pool.

Subjects are randomly selected within strata defined by population characteristics. The response rate is typically 19–20%, accounting for both nonresponse and incomplete surveys. Responses exceeding the 1,000 daily quota are deleted using systematic (random start, *n*th selection) sampling. The panel is matched on age, gender, income, and region to the U.S. Census Bureau's monthly *Current Population Survey* (*CPS*; U.S. Census Bureau, 2008). Results are poststratification weighted (Kish, 1965) to these *CPS* criteria. In addition, Synovate uses weighting in order to adjust partially for underrepresentation of Hispanics and ethnic minorities in the sample. Subjects were included in a prize drawing as the incentive to participate. Table 1 provides the unweighted and weighted sample demographics and compares them to recent U.S. Census data and to *eNation* weighting targets, which are based on U.S. Census data.

This study used a nonexperimental survey design in order to obtain descriptive measures for the incidence of lying in the population. Results from this survey are compared with the popular standard established by the DePaulo et al. (1996) studies.

Procedure and measures

Subjects received an invitation e-mail asking them to participate in an omnibus survey on the date of the study. The invitation was directed to a specific member of the household identified by age and gender. The invitation instructed that individual to click on a link to the survey website. When subjects accessed the site, they were provided with instructions, asked questions confirming participant identification, asked the omnibus survey questions for several unrelated topics, and asked a series of demographic questions. On the day of the lying study, subjects were asked about four topics (in order of presentation): packaged meals, cat litter products, lying behavior, and water softeners.

The DePaulo et al. (1996) diary study (and subsequent diary studies) used subject training to make the topic less sensitive and provide the subjects with a common definition of lying. Training of survey respondents was not possible, so to encourage accurate reporting, the self-report lying question was preceded by a definition of lying that incorporated the elements of foreknowledge and intent described at the beginning of this article. A brief description of types of lies was also included. Both were presented in a nonpejorative manner:

> We are interested in truth and lies in people's everyday communication. Most people think a lie occurs any time you intentionally try to mislead someone. Some lies are big while others are small; some are completely false statements and others are truths with a few essential details made up or left out. Some lies are obvious, and some are very subtle. Some lies are told for a good reason. Some lies are selfish; other lies protect others. We are interested in all these different types of lies. To help us understand lying, we are asking many people to tell us how often they lie.

Table 1 Comparison of the Unweighted and Weighted Demographics of the Study Sample to U.S. Census Data and eNation Demographic Targets

	CPS[a] (January 2007)	*eNation* Target	Unweighted Results	Weighted Results	Weighted ± *CPS*
Gender (%)					
Male	49.1	48.4	48.4	48.3	−0.8
Female	50.9	51.6	51.6	51.7	0.8
Region (%)					
Northeast	18.2	18.7	18.4	18.8	0.5
MidWest	22.1	22.2	21.6	22.3	0.2
South	36.4	36.2	35.9	36.0	−0.4
West	23.3	22.9	24.1	23.0	−0.3
Age (%)					
Under 25	12.8	12.7	5.3	12.5	−0.3
25–34	17.9	17.9	17.5	18.0	0.1
35–44	19.2	19.2	19.7	19.3	0.1
45–54	19.5	19.5	22.6	19.6	0.1
55–64	14.5	14.5	23.3	14.5	0.0
65+	16.2	16.2	11.6	16.2	0.0
Income (%)					
Under $25,000	18.8	18.7	14.3	18.5	−0.3
$25,000–$49,999	24.4	24.5	24.2	24.6	0.2
$50,000–$74,999	19.7	19.6	22.3	19.7	0.0
$75,000+	37.1	37.2	39.2	37.3	0.2
Race (%)[b]					
White	80.1	81.5	91.0	83.0	2.8
Black	12.6	9.3	3.3	9.2	−3.4
Asian	4.4	1.6	2.8	1.6	−2.8
Other	2.9	6.1	2.9	6.2	3.3
Refused/NA		1.5			
Hispanic (%)[b]					
Yes	15.1	7.5	4.6	7.6	−7.5
No	84.9	92.5	95.4	92.4	7.5

Note: CPS = Current Population Survey.

[a]Synovate weights are based on most current monthly *CPS* data at the time of study (March 2008).

[b]Race and Hispanic *eNation* targets are adjustments toward the U.S. Census data, not Census-based targets.

Subjects were then asked, using an open-ended format, how many times they had lied in the past 24 hours. They responded separately for lies to family members, friends, business contacts, acquaintances (''people you do not know but might see occasionally''), and total strangers; for each type of receiver, they were asked about lies in both face-to-face and mediated situations. Response categories were used as

Figure 1 Screen shot of the Internet survey page giving a definition and description of lying.

a mnemonic device and to proved additional detail for the analysis. The results of the 10 receiver-mode combinations were aggregated. Specifically, the question was worded:

> Think about where you were and what you were doing during the past 24 hours, from this time yesterday until right now. Listed below are the kinds of people you might have lied to and how you might have talked to them, either face-to-face or some other way such as in writing or by phone or over the Internet. In each of the boxes below, please write in the number of times you have lied in this type of situation. If you have not told any lies of a particular type, write in "0." In the past 24 hours, how many times have you lied?

The subjects were presented with a response grid showing the five types of people and two modes of communication. Subjects were instructed to enter a number in each box. The Internet questionnaire required a response in each of the 10 boxes before allowing the subject to continue to the next Web page. Figure 1 presents a screen shot of the lying description used in the Web survey; Figure 2 is a screen shot of the survey question. Based on the 5×2 individual categories of responses, the row, column, and grand total frequencies (of lies per day) were aggregated for each subject.

Results

The results of this national study are consistent with the oft-cited observation that *on average* Americans tell one to two lies per day ($M = 1.65$ lies per day, $SD = 4.45$, $Mdn = 0$, Mode $= 0$, $N = 998$, Max $= 53$ lies, 95% CI $= 1.37–1.93$).[2] But the most intriguing finding is the distribution of responses, not the mean. As Figure 3 illustrates the majority of people report telling no lies during the past 24 hours and most of the reported lies are told by few people. The 40.1% who reported lying told a total of

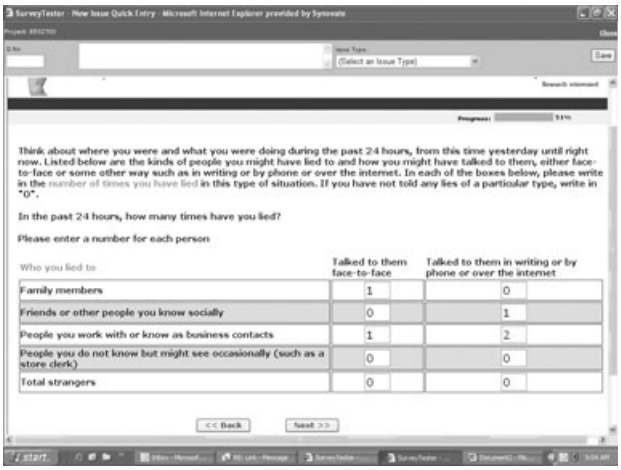

Figure 2 Screen shot of the Internet survey page with the question and response grid.

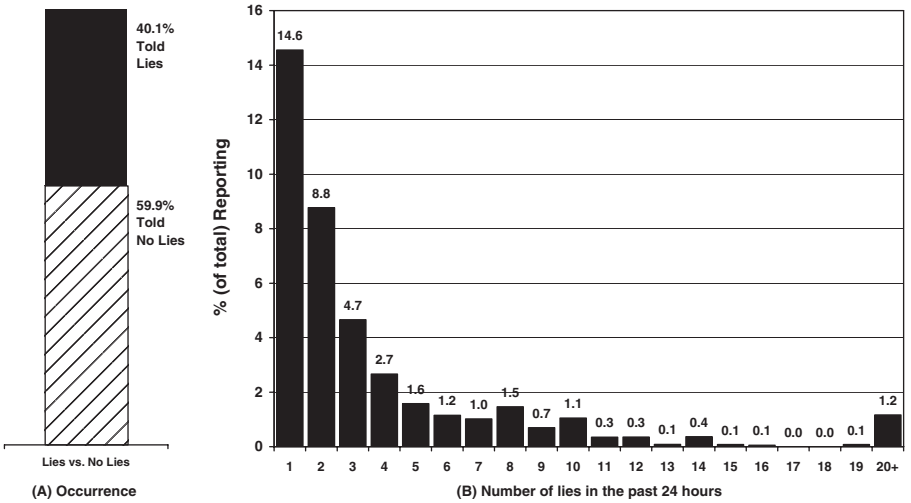

Figure 3 (A) The majority of Americans reported they did not lie in the past 24 hours (59.9%). (B) Percentage distribution by number of lies told; 32.2% told one to five lies and 7.9% reported telling six or more lies.

1,646 lies ($M = 4.11$, $SD = 6.26$, $n = 400$). Of these, 22.7% of all reported lies were told by 1% of the national sample. Results indicate that one-half of all reported lies are told by just 5.3% of American adults ($M = 15.61$, $SD = 11.22$, $n = 53$).

Figure 4 indicates that, among those who reported lying, the proportion of people who report a particular number of lies per day decreases as a function of the number of lies. Moreover, observation of this curve suggests that the decrease is a standard power function. Fitting a power function curve to these data produces

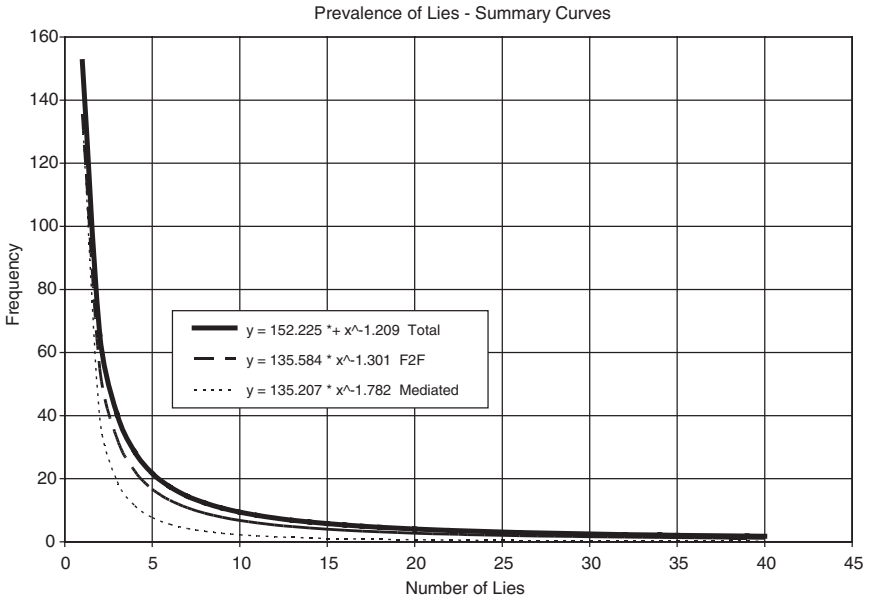

Figure 4 Frequencies of total, face-to-face, and mediated lies are represented by similar standard power functions.

the equation $y = 152.225 \times x^{-1.209}$ where x is the number of lies reported per day and y the frequency of people reporting a given number of lies. The function's sizeable intercept indicates that the majority of respondents who report lying do so in moderation, and the steep negative slope indicates the substantial decrease in frequency as the number of daily lies increases. Figure 4 also indicates that this pattern holds regardless of mode of communication. Although the slopes for both face-to-face lies and mediated lies are even steeper than that of the aggregated data, both retain the power law character of the total data set. Regardless of mode, most people report telling no lies and as the curve-fitting for number of lies reinforces, among those who report lying, most tell very few lies; but in each case, there are a few subjects who account for a large proportion of the lies being told.

Similarly, Figure 5 (among those telling face-to-face lies) and Figure 6 (among those telling mediated lies) indicate that lying behavior replicates the fractal character of power functions observed in other disciplines such as biological systems (Brown et al., 2002) and market segmentation (Anderson, 2008). Although each group of lies told to family members or friends or other types of message receivers represents only a small portion of the total lies, within each of the 10 mode-receiver combinations, the power function pattern is repeated. Most of the variation is among the intercepts and reflects that more lies are typically told to family members or friends than to acquaintances or total strangers.

These data do not include the number of interactions by type; therefore, it is possible that this variation is as much due to the number of opportunities for lying as it

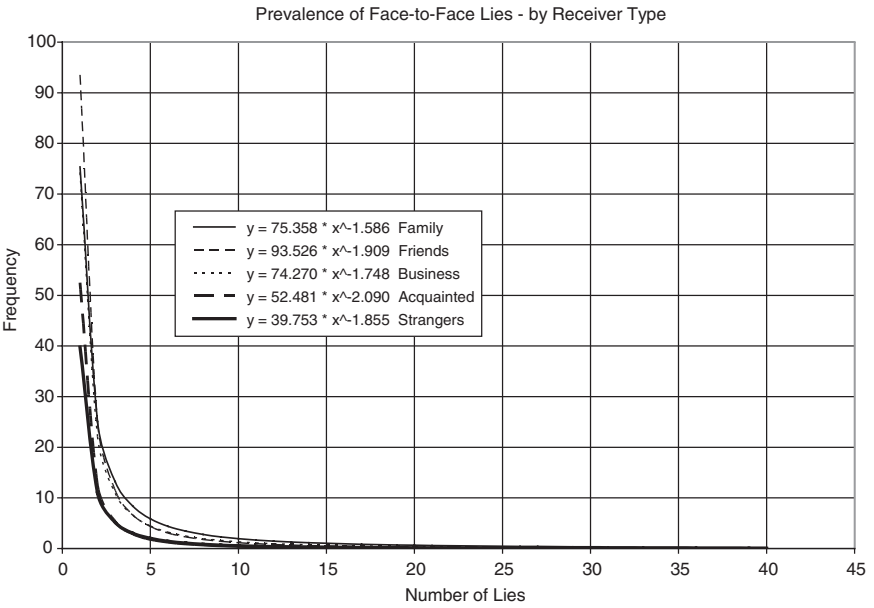

Figure 5 Standard power function curves fit the frequencies of face-to-face lies plotted by number of lies for all receiver types.

is that people are more likely to lie to others with whom they are familiar. Establishing the proportion of interactions involving lies may be more difficult than establishing the rate of lying for a prescribed time period. All of the diary studies to which these results are compared are flawed with regard to the interaction ratio. In order to capture as many lies as possible, DePaulo et al. (1996) specified that subjects were to record all interactions of 10 minutes or more. Subjects were told to record all lies that occurred during these interactions and, importantly, record lies occurring during shorter interactions as well. As a result, the ratio of total lies, regardless of interaction duration, to 10-minute interactions distorts the true relationship. If the number of 10-minute interactions varies across respondents or modes of communication, these comparisons may be more misleading than comparisons of the number of lies in each category during the fixed 24-hour time frame. Subsequent diary studies incorporate the lie per interaction distortion created by the DePaulo et al. methodology.

In order to consider more fully the possibility of individual differences in the propensity to lie, a multiple regression analysis was performed. Initially, a natural logarithm transformation was applied to the continuous lying measure as a means of reducing its nonnormality. Although not eliminating the nonnormality completely, this transformation had the effect of decreasing skewness by a factor of approximately 4 and kurtosis by a factor of approximately 18. To assess the impact of the demographic measures on lying, the natural logarithm transformed lying measure was regressed onto all demographic measures. Trivial predictors were dropped, the analysis was iterated, and two important, albeit modest in magnitude, predictors

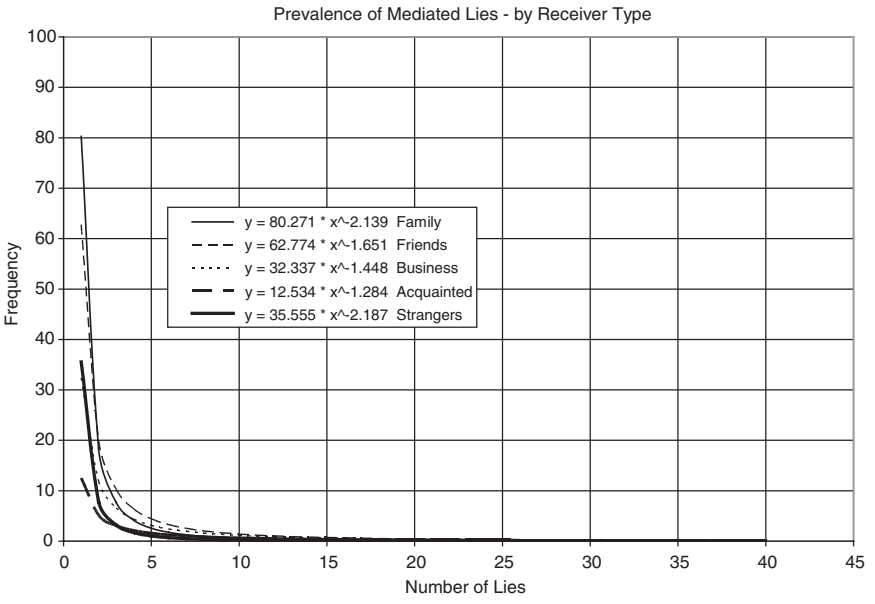

Figure 6 Standard power function curves fit the frequencies of mediated lies plotted by number of lies for all receiver types.

emerged ($R = .21, F(2, 979) = 22.82, p < .001$). First, age was an important predictor ($\beta = -.18$, $t(979) = -5.71$, $p < .001$), and the effect was such that a decrease in lying was associated with increasing age. Second, race (Caucasian vs. other) was an important predictor ($\beta = -.09$, $t(979) = -2.99$, $p = .003$), and the effect was such that Caucasians reported lying less than other racial groups. This analysis was replicated with the dichotomous not lie/lie measure, with the same two important predictors emerging in the subsequent logistic regression analysis.

No sex differences were observed when controlling for other demographic predictors. Bivariate analysis showed that the overall rate of lying by men ($M = 1.93$, $SD = 4.81, n = 482$) and women ($M = 1.39, SD = 4.08, n = 516$) is directionally but not significantly different when using conventional criteria for statistical significance ($t(996) = 1.89, p = .059$/ns, two-tailed, $d = 0.12$).[3] This apparent gender difference is in the opposite direction observed by DePaulo et al. (1996) but is consistent with the finding that women find lying less acceptable than men (Levine, McCornack, & Avery, 1992).

Discussion

The results of this national study are consistent with the DePaulo et al. (1996) diary study and suggest that *on average* Americans lie once or twice a day. However, the important findings are that many people do not lie on a given day, the majority of lies are told by a few prolific liars, and because the distribution is highly skewed, the mean number of lies per day is misleading. This pattern is consistent across modes of

communication and varies little on the basis of who is being lied to. Examination of individual differences suggests some variation but in most cases the differences are small.

The representativeness of online panel data is debatable. Use of poststratification (such as employed by Synovate) and propensity (likelihood of response) weighting schemes are the usual solutions to matching panel samples to the population. In general, the use of a large number of small and internally homogeneous strata will enhance the proportional fit of the sample to the population. However, non-coverage and self-selection can create sampling problems in Web-based research that weighting does not solve (Loosveldt & Sonck, 2008). These representativeness issues are of particular concern when the measured values are correlated with the underlying reasons for the selection bias. However, when measuring socially undesirable behavior, assessment of representativeness is confounded. Loosveldt and Sonck obtained measures for validation of questions influenced by social desirability using face-to-face interviews; this introduces potential mode effects. Even with the selection problem, the social distance advantage of the Internet may actually produce better data. Birnbaum (2004) provides an argument for representativeness when the results are homogeneous across strata. Because selection bias tends to vary across the levels of a stratified sample, small individual differences for a measure across the stratification variables is evidence that the aggregate finding transcends the sampling issues and is indicative that the sample result is representative of the phenomenon in the population as a whole. In Study 1, confidence in representativeness is enhanced by the homogeneity of the results across strata, and convergent validity is subsequently established by consistency with the results in Studies 2 and 3, and by the advantage of the Internet for creating social distance in the measurement of a sensitive topic.

Although the findings were generally homogenous across the sample of adult Americans, some small demographic differences were apparent. Notably, age and race/ethnicity account for small but statistically significant variation. Further, the difference between reports of lying by men and women approached statistical significance. These findings may have theoretical, social, and cultural implications.

Perhaps the most interesting individual difference is the negative association between lying and age. Lying is acquired by children in early childhood and the ability to lie is correlated with the acquisition of perspective-taking, theory of mind, and communication skills (Vasek, 1986; see Knapp, 2008, pp. 91–116 for a summary discussion of lying and development). As the child reaches adolescence, lying skill is perfected, but lying declines in acceptance in early adulthood (Jensen et al., 2004). The difference between rates of lies reported by the DePaulo et al. (1996) student and (adult) community samples suggests that maturity tempers the usage of lying as a strategy for goal attainment, and the current findings of the national data in Study 1 are consistent with the claim that the lying declines with age.

With regard to the finding that Caucasians report fewer lies than those of other ethnic or racial groups, it would be irresponsible to simply conclude that White people are more honest in general than those of other races. Research on race and

deception is limited, and more research is required to make sense of this finding. The current study also found a marginally significant trend toward men reporting more lies than women. Some studies suggest men lie more than woman while others suggest the opposite (e.g., DePaulo et al., 1996; Levine et al., 1992). Other research finds that sex differences vary by the topic of the lie (Haselton, Buss, Oubaid, & Angleitner, 2005).

Study 2

The results of Study 1 replicate, with a large and nationally representative sample, the often repeated conclusion that people lie, on average, once or twice per day. The results also document that the distribution of lies per day is substantially skewed. Most reported lies were told by a few prolific liars.

These findings have important implications. Most importantly, the nature of the distribution makes conclusions drawn from sample means misleading. Although the mean lies per day reported in the literature appear reflective of aggregate reality, the mean as a central tendency does not reflect the lying behavior of the typical person. Instead, most people reported telling no lies at all on a given day, with the median and mode both being zero. We are not the first to note this shape for the lie frequency distribution. A similar pattern, with many telling a few lies and a few telling most of the lies, was reported by Feldman, Forrest, and Happ (2002) in a laboratory study on self-presentation. But in their analysis of the data, the skew was treated as a methodological limitation rather than as a finding. The clarity of the results observed in the national study raises a question of whether or not this pattern existed in previous studies reporting lying frequency.

Method

The raw data were obtained from the student phase of the DePaulo et al. (1996) study and from both phases of the George and Robb (2008) study. A distribution of lie frequency was partially reconstructed from the results reported by Feldman et al. (2002). The shape of each of the four distributions (excluding those reporting no lies) is examined, and the resulting distributions are compared with the overall results from the national survey. Data from the DePaulo et al. community sample and Hancock et al. (2004) studies were not available.

Results

DePaulo et al. (1996)

Of the 77 students sampled by DePaulo et al., 76 reported telling at least one lie over the period of 1 week ($M_{\text{week}} = 13.74$, $SD_{\text{week}} = 11.40$; this is equivalent to $M_{\text{day}} = 1.96$, $SD_{\text{day}} = 1.63$); the total number of lies told was 1,058; and the most lies told by one person was 46 (equivalent to 6.6 lies per day). Curve-fitting yields a power function for these data of $y = 5.366 \times x^{-0.290}$ ($n = 76, r^2 = .289$). This function has a poor fit; however, no equation provided a fit better than $r^2 = .350$. Even so, the

data exhibit the overall distributional properties of a few lies told by most of the subjects and most of the lies told by a few liars. Of the total students in this study, 66.2% told the equivalent of two lies per day or less. Conversely, the seven most frequent liars (9.2% of the sample) told more than the equivalent of five lies per day, or 26.2% of all lies reported. Although raw data for the community sample are not available, DePaulo et al. reported 64 of 70 subjects told 477 lies with $M_{week} = 6.79$ ($SD_{week} = 6.86$) and $Mdn_{week} = 4.5$; these measures suggest a positive skew similar to that observed with the student sample.

George and Robb (2008)

Two studies were conducted following the diary methodology used by DePaulo et al. (1996) and Hancock et al. (2004). The objective of the two studies was to examine variations in deceptive behavior by media use. A key difference between these studies and the prior diary studies was the use of PDAs instead of paper and pencil to record interactions and lies. In the first study, George and Robb used the 10-minute minimum time established in the prior diary studies to define an interaction; in the second study, the minimum interaction time was reduced to 5 minutes. In both studies, subjects were instructed to record lies even when the interaction was shorter than the prescribed time.

Results of this research are notably different from the prior diary studies. In each study, the mean number of lies was small. In the first study, 24 of 25 subjects reported lying over the period of a week ($M_{week} = 4.16$, $SD_{week} = 2.59$; this is equivalent to $M_{day} = 0.59$, $SD_{day} = 0.37$; $Mdn = 4$); the total number of lies told was 104; and the most lies told by one individual was 11 (the equivalent of 1.6 lies per day). Despite the low average, only one subject who reported lying told fewer than two lies; consequently, a power function could not be fit to the distribution (if this individual is eliminated a power function with a modest $r^2 = .758$ can be fit to the remaining data). Even with this poor fit, the distributional properties of the study demonstrated a positive skew similar to the national survey and the DePaulo et al. (1996) diary study. Half of the subjects (48% of the sample) told just 25% of the lies while the three subjects (12%) reporting the most lies told 26.9% of the lies.

In George and Robb's (2008) second study, the length of interaction was shortened in order to encourage more reporting. As a result, 23 of 24 subjects reported lying over the period of a week and the lying frequency increased from the first study ($M_{week} = 6.33$, $SD_{week} = 3.78$; this is equivalent to $M_{day} = 0.90$, $SD_{day} = 0.54$; $Mdn = 5$). The total number of lies reported was 152 and the most lies told by one individual was 14 (the equivalent of two lies per day). Similar to the first study, and despite the overall low number of lies told, those reporting lies told a minimum of two lies; consequently, a power function could not be fit to the data. Nonetheless, the positive skew is again apparent in the pattern of responses. Just three subjects (12.5% of the sample) told 40 of 152 lies (26.3% of the total).

Feldman et al. (2002)

Although addressing a much shorter time frame in an experimental setting, Feldman et al. (2002) observe the same distinct skew pattern in their lying data as was observed in the daily national study and the weekly student diary studies. This study examined lying as a component of self-presentation; a sample of 121 students was divided into two induction groups and an experimental control. Across all subjects, a total of 211 lies were analyzed with a mean of 1.75 lies per subject and 2.92 lies per subject who lied.[4] Of the 121 subjects, 49 told no lies (40.5%), 23 told one lie (19.0%), and 18 told two lies (14.9%). The maximum number of lies was 12. Thus, 74.4% of the total sample accounted for only 28% of the lies told. The remaining 31 subjects (25.6%) told 3–12 lies, accounting for 72% of the lies ($M = 4.9$ lies per subject).

When prior research reporting the frequency of lies is reexamined, results show the small diary samples, the experimental self-presentation study, and our large, national self-report survey have similarly skewed distributions. In all cases, the infrequent liars are a large part of the sample and account for a disproportionately small share of the lies reported. Conversely, each study includes a small number of individuals who account for a disproportionately large share of the lies.

Study 3

Several issues are of concern despite the consistency of the findings across studies and the face validity obtained by reanalysis of prior studies reporting lie frequency. Primarily, these concerns have to do with the accuracy of the study findings using the survey approach to obtaining self-report data. Further, the apparent discrepancy between the numbers of nonliars when the time frame is 1 day versus the number when the time frame is 1 week needs to be resolved. For these reasons, the national survey was replicated using student samples and additional measures.

Accuracy of reporting lies

An obvious concern with self-report, mass survey research is accuracy of reporting. Bias in the national study data, if it exists, would likely manifest itself as under-reporting. Given the pervasive cultural prohibitions against lying, self-presentation motives favor under- rather than overreporting. Methodological research on other sensitive topics such as drinking (Lemmens, Tan, & Knibbe, 1992) and sexual behavior (Ramjee, Weber, & Morar, 1999) indicates that diary studies produce higher mean scores than self-report questionnaires. The diary method used by DePaulo et al. (1996) is expected to be less susceptible to underreporting bias even though the study had the added limitations of a smaller and nonrepresentative sample. Because the mean number of lies per day in the national survey data was greater than the mean observed in the diary research, concern that the national study may be underreporting due to the use of the self-report questionnaire method is not consistent with results of most comparisons of survey research to an alternative diary method. Thus, the

observed frequency of lying reported is not likely attributable solely to the survey methodology used to collect the national study data.

Nonetheless, a survey of lying behavior invites the question: "How do you know the subjects are not lying?" One answer resides in the methods used to address sensitive questions. Tourangeau and Yan (2007) identify four techniques for asking sensitive questions that might contribute to improved prevalence measurement: Use of self-administered rather than interviewer-administered questions, forgiving wording, randomized response techniques (RRT), and use of the bogus pipeline (BPL) approach. The national survey of lying was conducted as a self-administered study and used a forgiving wording preamble. The meta-analysis by Tourangeau and Yan (2007) provides clear evidence that self-administration yields more behavioral reporting for sensitive topics than does interviewer-administration. There is limited evidence that forgiving wording may also be helpful. Catania et al. (1996) found increased reporting of sexual activity but a series of experiments by Holtgraves, Eck, and Lasky (1997) suggest that forgiving wording is more likely to improve attitudinal responses than behavioral reporting.

A second answer can be obtained through the use of additional measures in order to assess the social desirability bias (SDB) in the self-report measures (Fisher & Katz, 2000). This is sometimes done directly by obtaining measures of SDB using a scale of SDB traits (Crowne & Marlowe, 1960; Reynolds, 1982) and adjusting for the subjects' level of bias. Alternatively, Fisher and Katz suggest the validity of reports of a socially undesirable behavior such as lying can be assessed indirectly, for example, by comparing the self-report of that behavior to subjects' estimates of the extent of that behavior in others. Subjects in Study 3 were asked to report the total number of times others had lied to them in the past 24 hours.

Minimum observable differences

A second key issue raised by the results of the national study is the minimum observable difference of one lie per day. Those who lie but do so less than once a day may be recorded as having told no lies. Because the frequency of lies is typically reported as the rate of lying in a fixed interval of time (i.e., one to two lies per day), expanding the duration in which the behavior can be observed is likely to increase the overall reported incidence of the behavior; measures taken over a longer period and converted to a daily rate will also increase precision. In order to observe those who lie once a week or once a month or even less frequently, it is necessary to use a wider time aperture. In the student sample of the DePaulo et al. (1996) study, 30% of diaries recorded six or fewer lies per week (less than once per day). If the number of lies told in a week is evenly distributed across the days of the week (and there is no reason to believe an even distribution is a good assumption), we might expect that as many as 17% (those reporting four, five, or six lies in a week) have an above average likelihood of being included as liars in a 1-day study. But the other 13% (reporting one, two, or three lies in a week) would be more likely to report not lying on the survey day. Identifying infrequent liars by repeating the daily survey on 2 or

more days with the same sample, or by asking those who report no lies to identify other times when they had told a lie, we would expect to narrow substantially the gap between the national study observation that 40% told a lie on one given day and the DePaulo et al. result that 95% of subjects (both studies combined) reported at least one lie over the period of a week.

Method

Participants and design

Students were recruited from communication, advertising, and marketing classes at two large universities in the Midwestern United States. Of the 229 subjects, two failed to provide complete information regarding the number of lies told and two provided answers that were determined to be statistical outliers using a discordancy test. These subjects were excluded from the analysis. Of the remaining 225 respondents, 140 were female (62.2%) and 85 were male (37.8%). Participation was voluntary.

The primary purpose of this study was to cross-validate the results of the national study using a separate sample. A nonexperimental survey design similar to that of the national study was used in order to obtain comparable descriptive measures for the incidence of lying in the student sample. Results from this survey are compared with those of Study 1 and the reanalysis of prior studies in Study 2.

Procedures and measures

The study was administered in five separate classroom settings using a paper questionnaire. The introduction and frequency of lies question were identical to that of the online survey used with the national sample except that students could leave parts of the 5×2 grid blank (the online survey forced subjects to place zeroes in all cells for which no lies were told before they could continue with the questionnaire). Blank cells were coded as zero lies. In addition, those subjects who reported telling no lies in the past 24 hours were asked a follow-up closed-ended question regarding when they last told a lie. Response categories included, ''more than 24 hours ago but within the last 2 days,'' ''more than 2 days ago but within the last week,'' ''more than a week ago but within the last month,'' ''more than a month ago,'' and ''never.'' Subjects were asked how many lies they had been told by others in the past 24 hours and what percentage of the U.S. adult population lies on a given day. One group of subjects was asked about the difficulty of the lie reporting task.

Results

The pattern of lie frequency for the total student sample is consistent with the distribution of lies for the national survey using the same frequency measure. The data fit a power function of $y = 53.891 \times x^{-1.012}$ with a goodness-of-fit of $r^2 = .894$. The pattern of many telling few lies versus a few telling many lies is reproduced in the student sample frequencies. Of the total subjects in the student sample, 68.4% told one, two, or no lies (accounting for 24.5% of 526 total lies); conversely, the 13 most prolific liars (5.8% of the total sample) told 22.4% of the lies. On the basis of the relationship

of the means in the DePaulo et al. (1996) student and community diary studies, we expected the mean of the student survey to be higher than the mean of the national adult sample. Results from the student sample ($M = 2.34$ lies per day, $SD = 2.94$, $Mdn = 1$, Mode $= 0$, $N = 225$, Max $= 21$ lies, 95% CI $= 1.95–2.72$) indicate higher frequency than in the national survey (nonoverlapping 95% confidence intervals; $t(1221) = 2.62$, $p < .01$, $d = 0.18$).

Those subjects who reported telling no lies in the past 24 hours (28.9%) were asked the follow-up question regarding when they last told a lie. Combining those who indicated their most recent lie was within the past week (though not in the past 24 hours) with those who reported at least one lie in the past 24 hours, a total of 92.4% reported telling a lie in the past week (204 of 221 subjects; 4 did not answer the follow-up question). This result is reasonably consistent with the DePaulo et al. (1996) result of 95%.

Since there is a legitimate concern about SDB in the self-report of lying, a check on the subjects' reported rate of lying (lies per day) was obtained by asking subjects in the student sample how many times they believe they were lied to by others in the past 24 hours. Subjects reported that others told an average of 2.79 lies in the past 24 hours ($SD = 2.82$, $N = 198$). Although 28.9% reported telling no lies, 17.2% reported not being lied to by others. Own lies ($M = 2.34$, $SD = 2.94$) are significantly lower than others' lies; the paired sample $t(197) = 2.18$, $p < .05$. Despite this significant difference, the distribution of others' lies is strikingly similar to that of own lies following the long-tail pattern observed in the self-report of lies (the power function for others' lies is $y = 44.552 \times x^{-0.790}$ with a modest goodness-of-fit of $r^2 = .700$). It is important to note that individual subject's reports of others' lies did not mimic their self-report of lying. The correlation of the two measures is a meager $r = .073$ ($p = .152$, *ns*).

Participants were also asked what proportion of the U.S. adult population (18 years or older) told at least one lie on any given day. The mean estimate is 75.8% ($SD = 21.7$). The proportion of students (who are a subsegment of the adult population) telling at least one lie in the past 24 hours is 71.1% (CI $= \pm 5.9\%$ points). Although not directly comparable due to the different ways in which the proportions were derived, in both instances one percentage falls within the confidence interval of the other, suggesting that the two proportions would not be significantly different.

Discussion

The results of Study 3 fit a power function similar to that of the national survey and further indicated that, based on self-reports of lying behavior, most people tell few or no lies in a given day but a few prolific liars tell a disproportionately large share of the daily lies. Survey results, diary studies, and the distribution of lies in an experimental setting share this pattern, and the consistency of results provides strong evidence that the frequency of lying has a strong positively skew. The similarity of subjects' estimates of others' lies and the distribution of number of lies told provides convergent validity for the self-report measure of lying.

With regard to the accuracy of mean prevalence figures, Study 3 provides mixed results. The mean reported by this student sample is 2.34 lies per day and expands the range of responses across all studies to 0.59–2.34 lies per day. The mean for the student survey is significantly higher than the mean of the national survey of American adults. This is similar to the relationship observed by DePaulo et al. (1996) for their student and community samples. Consistent with the Study 1 finding that younger people tend to lie more than older people, the predominantly young student sample has a higher frequency of lying. Two measures were obtained to assess the validity of the reported lying frequency. One measure estimated the proportion of the population that lies on a given day; the result of 75.8% is similar to the Study 3 result showing that 71.1% of the subjects told at least one lie in the past 24 hours. However, an estimate of others' lies to the subjects ($M = 2.79$) is significantly higher than the average of the subjects' own lies. Either the subjects truly believe that other people lie more than they do or the subjects are slightly but systematically underestimating their own lying behavior.

Regardless, the difference is one of degree and not magnitude, and the relevant finding is the consistent skew of the distribution and not the average number of lies told. In fact, the mean of a long-tail (power) distribution is in part a function of sample size. As the sample size increases, the probability of rare but legitimate extreme values that are not statistical outliers being reported also increases. In Study 1, the maximum number of lies by one person was 46 ($N = 998$). Study 3 consisted of five subsamples (separate classroom administrations of the survey) and the maximum number of lies told was closely associated with sample size. In the largest subgroup ($n = 118$), the maximum number of lies was 21. In the next largest group ($n = 31$), the most lies told was 12. In the three smallest subsamples ($n = 26$, 25, and 25, respectively), the most lies told by one person were 10, 8, and 6.

A question related to accuracy may be raised with regard to the difficulty of the task (since it might be expected that a more difficult task would produce increased variance). We note that researchers familiar with the topic of deception tend to struggle more with the question asked in the self-administered survey than do naive subjects. Those familiar with the deception literature wrestle with the explication of lying and, perhaps because of this cognitive effort, their own ability to recall precisely over a 24-hour period. Naive subjects seem to have no such problem. As a check, 25 students in one Study 3 subsample completed the recall of lies task and were subsequently asked to rate the difficulty of answering the question. On a 1–10 scale (1 = not at all difficult; 10 = extremely difficult), the subjects reported $M = 3.48$ ($SD = 3.03$) with the $Mdn = 2$ and the Mode = 1. Most subjects seemed to feel that they were able to complete the task with little difficulty. During a verbal debriefing of this subject group, subjects reiterated that most had made an earnest effort to estimate the number of times they had lied and most felt the task was not difficult.

Finally, Study 3 appears to have resolved the discrepancy in the proportion of nonliars reported in the daily survey and weekly diary results. All of the diary studies report that more than 90% of the subjects told at least one lie in the period of 1

week. However, many of those subjects lie at a rate that translates to less than one lie per day. When Study 3 subjects who did not report lying in the 24 hours preceding the survey were asked when they last told a lie, results indicate that most did lie in the previous week. Therefore, the apparent discrepancy is due to the precision of the question being asked. The daily questions do not allow for fractional responses, whereas the conversion of weekly data to daily data inherently creates fractional reporting. When the precision issue is accounted for by additional measurement, the discrepancy disappears.

General discussion

The results of Study 1 reproduce with a large, national sample the conclusion that people lie, on average (arithmetic mean), once or twice per day. More importantly, the results document that the distribution of self-reported lies per day is substantially skewed. Most people report telling few or no lies on a given day and most reported lies are told by a few prolific liars. The reanalyses in Study 2 and the replication of the survey methodology in Study 3 provide strong evidence in support of this finding. All the data reported here are consistent with the claim that most lies are told by a few prolific liars.

The highly skewed, long-tail distribution that results from reports of lying may be emblematic of a larger class of behaviors. Although much of what we measure and observe relies on the principle of randomness and results in the well-known bell-shaped curve of observations distributed around a central tendency for a phenomenon, scientists continue to find new evidence for a self-organizing principle in nature that produces distributions that are scale-free, or lacking in a characteristic tendency (Barabási, 2003; Barabási & Albert, 1999). This principle is found in the examination of atomic structure, biological systems, economic theory (the Pereto Law or ''80/20 Rule''), and the emergence of nodes on the Internet. The latter describes a pattern consisting of a few very large nodes that are connected with extremely high frequency (e.g., Google, eBay, and Amazon.com) that are surrounded by over 100 million smaller and, in many cases isolated, Websites. But it is the human tendency to search, use, and link to these few extreme large sites while ignoring so many others that creates this apparent scale-free pattern in the first place. It may be less the case that lying is some unique form of communicative behavior that divides people into those who do it (with vigor) and those who do not than an indicator that we need to reexamine a broad range of social and symbolic behaviors, look for scale-free distributions, and consider broadly the implications for all forms of communication research.

Our finding with regard to the distribution of reported lies has specific implications for the conclusions drawn from deception detection accuracy experiments. A meta-analysis shows that people do only slightly better than chance (54%) when distinguishing between truths and lies (Bond & DePaulo, 2006). In the experiments

cited by Bond and DePaulo, the deception base rate is almost always 50% and subjects are typically truth-biased. Truth bias leads to the veracity effect, meaning that accuracy is higher for detecting honest messages than for detecting lies (Levine, Park, & McCornack, 1999). As a consequence, honesty base rates impact accuracy such that people are increasingly accurate as the proportion of truthful to deceptive messages increases (Levine, Kim, Park, & Hughes, 2006). The finding that most people are honest most of the time suggests that experiments employing nonrepresentative base rates (i.e., a lower proportion of truths to lies than is typical of everyday interactions) will underestimate accuracy. Findings from this general population study indicate that truth bias may be functional. Since self-report data suggest most people do not lie or tell very few lies in the course of a typical day, it is reasonable for people to believe others (that is, to be truth-biased) most of time.

However, although laboratory deception detection experiments may underestimate overall accuracy due to nonrepresentative base rates, lie detection rates may be overestimated. The current findings suggest substantial individual differences for truth telling (or lying) with an honest majority and a few prolific liars. Bond and DePaulo (2008) report substantial variation in both people's demeanor and people's ability to lie. People's ability to lie successfully likely impacts how often they lie (Kashy & DePaulo, 1996). We speculate that the prolific liars are likely those people with especially honest demeanor and that unusually transparent liars avoid lying. If most lies outside the laboratory are told by people who are usually believed, lie detection rates would be lower than those observed in randomized experiments.

One caution relates to Rozin's (2001) concern for achieving "context- and culture-sensitive scientific social psychology." The large sample of American adults used in the current research is more representative and allows greater generalizability than previous studies using students and other convenience samples. Nevertheless, caution should be taken before assuming that these results will hold for other cultures, individuals under 18 years of age, or subsets of the population with characteristics or beliefs that are sharply discrepant from the norm. Anecdotal evidence suggests that base rates may well vary across cultures and studies of lying involving young children and adolescents show considerable variation in the ways lying occurs (Knapp, 2008). Criminals, political extremists, and those at the far ends of the socioeconomic spectrum may well behave differently than the vast majority of the American population. The domain of deception research would benefit from cross-cultural, subcultural, and cross-generational examination of the prevalence and distribution of lying behavior.

Conclusion

Over time, most Americans probably lie at least occasionally. But the inference of pervasive daily lying, drawn from the statistical average of one to two lies per day, and reinforced by media coverage of corporate deception and political malfeasance, as well as pop culture portrayals of deception detectors, belies the basic honesty

present in most people's everyday communication. Self-report data for the U.S. adult population show the average rate of lying is around 1.65 lies per day; but the data are not normally distributed. On any given day, the majority of lies are told by a small portion of the population, and nearly 6 out of 10 Americans claim to have told no lies at all. As researchers continue to examine the nature of lying and look for ways to detect deception effectively, both the theories of deception, and the methods used to test those theories, necessarily must take into account that veracity is not a constant across the population and that the propensity to lie can be an important moderator.

Acknowledgments

The authors thank Deborah Kashy and Joey George for providing the original data from their studies for additional analysis.

Notes

1 Synovate, Inc. (Chicago) conducts marketing research and strategic studies for business, industry, and government agencies. White papers on the omnibus methodology are available at www.synovate.com. The omnibus panel was surveyed on April 30, 2008.

2 Two subjects, one with an extreme value of 134 lies (20.8 standard deviations above the mean of 1.80 lies) and a corresponding subject with no lies, were eliminated from the sample by α-trimming. Using procedures outlined by Barnett and Lewis (1978), a discordancy test appropriate for gamma distributions was applied to the data. Even though the potential for extreme values is inherent in the tail of the power function, the most extreme value recorded in these data is well beyond the limits of statistical probability.

3 The application of a *t*-test for establishing the significance of differences for nonnormal distributions is problematic. Nonetheless, in those situations where the shapes of the distributions are similar, the sample sizes are similar, and the samples are sufficiently large the *t*-test will provide an acceptable test of significance even though the assumption of normality has been violated (Johnson, 1978). Although most of the small student sample studies fail to meet all of these conditions, the power distributions in this large-scale national study replicate for most subsamples and do so with large enough samples to allow for robust statistical testing. An alternative is to make the nonparametric assumption and apply difference testing accordingly. Calculating chi-square for age (under 45 years old vs. 45 years and older) and lying (did vs. did not lie) yields $\chi^2(1, N = 998) = 36.59, p < .001$; this leads to a conclusion similar to the one reported in the text, specifically that younger people are more likely to lie than older people. Calculating chi-square for gender and lying yields $\chi^2\ (1, N = 998) = 2.84, p < .092$; again the result is consistent with the reported *t*-test showing that men are directionally but not significantly more likely to lie than are women.

4 Despite including the induction groups, using a student sample, and being constrained to only those who told lies, the mean of 2.92 lies has been misinterpreted in the media (online program description, *About Lie to Me*, 2009, http://www.fox.com/lietome/about) and in scholarly work as claiming "in the course of a 10-minute interaction . . . the average person tells two to three lies" (Harrington, 2009, p. 4).

References

About *Lie to Me*. (2009). Online program description. Retrieved April 10, 2009, from http://www.fox.com/lietome/about.

Anderson, C. (2008). *The long tail*. New York: Hyperion.

Asch, S. E. (1987). *Social psychology* (Original work published 1952). New York: Oxford University Press.

Barabási, A.-L. (2003). *Linked*. New York: Plume.

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, **286**, 509–512.

Barnett, V., & Lewis, T. (1978). *Outliers in statistical data*. Chichester, UK: Wiley.

Birnbaum, M. H. (2004). Human research and data collection via the Internet. *Annual Review of Psychology*, **55**, 803–832.

Bok, S. (1999). *Lying: Moral choice in public and private life*. New York: Vintage.

Bond, C. F., Jr., & DePaulo, B. M. (2006). Accuracy of deception judgments. *Review of Personality and Social Psychology*, **10**, 214–234.

Bond, C. F., Jr., & DePaulo, B. M. (2008). Individual differences in judging deception: Accuracy and bias. *Psychological Bulletin*, **134**, 477–492.

Brown, J. H., Gupta, V. K., Li, B.-L., Milne, B. T., Restrepo, C., & West, G. B. (2002). The fractal nature of nature: Power laws, ecological complexity, and biodiversity. *Philosophical Transactions of the Royal Society B*, **357**, 619–626.

Catania, J. A., Binson, D., Canchola, J., Pollack, L. M., Hauck, W., & Coates, T. J. (1996). Effects of interviewer gender, interviewer choice, and item wording on responses to questions concerning sexual behavior. *Public Opinion Quarterly*, **60**, 345–375.

Cochran, S. D., & Mays, V. M. (1990). Sex, lies and HIV. *New England Journal of Medicine*, **322**(11), 774–775.

Crowne, D. P., & Marlowe, D. (1960). A new scale of social desirability independent of psychopathy. *Journal of Consulting Psychology*, **24**, 349–354.

DePaulo, B. M., Kashy, D. A., Kirkendol, S. E., Wyer, M. M., & Epstein, J. A. (1996). Lying in everyday life. *Journal of Personality and Social Psychology*, **70**, 979–995.

Feldman, R. S., Forrest, J. A., & Happ, B. R. (2002). Self-presentation and verbal deception: Do self-presenters lie more? *Basic and Applied Social Psychology*, **24**(2), 163–170.

Fisher, R. J., & Katz, J. E. (2000). Social desirability bias and the validity of self-reported values. *Psychology & Marketing*, **17**(2), 105–120.

Ford, C. V., King, B. H., & Hollender, M. H. (1988). Lies and liars: Psychiatric aspects of prevarication. *American Journal of Psychiatry*, **145**, 554–562.

George, J. F., & Robb, A. (2008). Deception and computer-mediated communication in daily life. *Communication Reports*, **21**, 92–103.

Hancock, J. T., Thom-Santelli, J., & Ritchie, T. (2004). Deception and design: The impact of communication technology on lying behavior. *CHI Letters*, **6**(1), 129–134.

Harrington, B. (Ed.). (2009). *Deception: From ancient empires to Internet dating*. Stanford, CA: Stanford University Press.

Haselton, M. G., Buss, D. M., Oubaid, V., & Angleitner, A. (2005). Sex, lies, and strategic interference: The psychology of deception between the sexes. *Personality and Social Psychology Bulletin*, **31**, 3–23.

Holtgraves, T., Eck, J., & Lasky, B. (1997). Face management, question wording, and social desirability. *Journal of Applied Social Psychology*, **27**, 1650–1671.

Jensen, L., Arnett, J., Feldman, S., & Cauffman, E. (2004). The right to do wrong: Lying to parents among adolescents and emerging adults. *Journal of Youth and Adolescence*, **33**(2), 101–112.

Johnson, N. J. (1978). Modified *t* tests and confidence intervals for asymmetric populations. *Journal of the American Statistical Association*, **73**, 536–544.

Kalish, N. (2004, January). How honest are you? *Reader's Digest*, 114–119.

Kish, L. (1965). *Survey sampling*. New York: John Wiley & Sons.

Kashy, D. A., & DePaulo, B. M. (1996). Who lies? *Journal of Personality and Social Psychology*, **70**, 1037–1051.

Knapp, M. L. (2008). *Lying and deception in human interaction*. Boston: Pearson Education.

Knox, D., Schacht, C., Holt, J., & Turner, J. (1993). Sexual lies among university students. *College Student Journal*, **26**, 269–272.

Lemmens, P., Tan, E. S., & Knibbe, R. A. (1992). Measuring quantity and frequency of drinking in a general population survey: A comparison of five indices. *Journal of Studies on Alcohol*, **53**, 476–486.

Levashina, J., & Campion, M. (2007). Measuring faking in the employment interview: Development and validation of an interview faking behavior scale. *Journal of Applied Psychology*, **92**(6), 1638–1656.

Levine, T. R., Kim, R. K., & Hamel, L. M. (2009). *People lie for a reason: An experimental test of the principle of veracity*. Unpublished manuscript, Michigan State University.

Levine, T. R., Kim, R. K., Park, H. S., & Hughes, M. (2006). Deception detection accuracy is a predictable linear function of message veracity base-rate: A formal test of Park and Levine's probability model. *Communication Monographs*, **73**, 243–260.

Levine, T. R., McCornack, S. A., & Avery, P. B. (1992). Sex differences in emotional reactions to discovered deception. *Communication Quarterly*, **40**, 289–296.

Levine, T. R., Park, H. S., & McCornack, S. A. (1999). Accuracy in detecting truths and lies: Documenting the "veracity effect." *Communication Monographs*, **66**, 125–144.

Loosveldt, G., & Sonck, N. (2008). An evaluation of the weighting procedures for an online access panel survey. *Survey Research Methods*, **2**(2), 93–105.

Patterson, J., & Kim, P. (1991). *The day America told the truth*. New York: Prentice-Hall.

Ramjee, G., Weber, A. E., & Morar, N. S. (1999). Recording sexual behavior: Comparison of recall questionnaires with a coital diary. *Sexually Transmitted Diseases*, **26**, 374–380.

Reynolds, W. M. (1982). Development of reliable and valid short forms of the Marlowe-Crowne social desirability scale. *Journal of Clinical Psychology*, **38**, 119–125.

Rozin, P. (2001). Social psychology and science: Some lessons from Solomon Asch. *Personality and Social Psychology Review*, **5**(1), 2–14.

The Real Truth About Lying (1996, September/October). *Psychology Today*, **29**, 16.

Tosone, C. (2006). Living everyday lies: The experience of self. *Clinical Social Work Journal*, **34**(3), 335–348.

Tourangeau, R., & Yan, T. (2007). Sensitive questions in surveys. *Psychological Bulletin*, **133**(5), 859–883.

U.S. Census Bureau (2008). *Current population survey.* Retrieved June 5, 2008, from http://www.census.gov/cps/.

Vasek, M. E. (1986). Lying as a skill: The development of deception in children. In R. W. Mitchell & N. S. Thompson (Eds.), *Deception perspectives on human and nonhuman deceit* (pp. 271–292). Albany: State University of New York Press.