HBS Working Paper #06-007
Revised: 2/18/2007

Slippery Slopes and Misconduct: The Effect of Gradual Degradation on the Failure to Notice Others' Unethical Behavior

Francesca Gino
Carnegie Mellon Tepper School of Business

And

Max H. Bazerman,
Graduate School of Business Administration
Harvard University

Running head: SLIPPERY SLOPES

Slippery Slopes and Misconduct:

The Effect of Gradual Degradation on the Failure to Notice

Others' Unethical Behavior

FRANCESCA GINO
GSIA Room A-10A
Tepper School of Business/CMU
Pittsburgh PA 15213
Tel: 412-268-5810
e-mail: fgino@andrew.cmu.edu

MAX H. BAZERMAN
Baker Library 453
Harvard Business School
Boston MA 02163
Tel: 617-495-6429
e-mail: mbazerman@hbs.edu

Draft: February 18, 2007

The authors gratefully acknowledge the financial support of the Harvard Business School Division of Faculty Research and Development and the support of the staff and facilities of the Computer Laboratory for Experimental Research. We are especially grateful to Daylian Cain, Andrew McAfee, Don Moore, Lamar Pierce, Jan Rivkin and Sandra Sucher for their insightful feedback on earlier drafts. Please address correspondence to fgino@andrew.cmu.edu or mbazerman@hbs.edu.

Abstract

Four laboratory studies show that people are more likely to accept others' unethical behavior when ethical degradation occurs slowly rather than in one abrupt shift. In the studies, participants served in the role of watchdogs charged with catching cheating in a series of trials. The cheating they observed increased either gradually or abruptly; people were more likely to overlook cheating that increased gradually. Our studies also provide evidence as to why people accept cheating by others. Our results indicate that the effect is due at least in part to the failure to notice that unethical behavior is occurring when the change is gradual rather than abrupt.

Key words: unethical behavior, implicit biases, gradual changes, ethical decision-making

Slippery Slopes and Misconduct:

The Effect of Gradual Degradation on the Failure to Notice

Others' Unethical Behavior

Companies such as Enron, Tyco, Parmalat, and WorldCom are often cited as examples of business disasters resulting from unethical management behavior. More broadly, evidence suggests that ethical misconduct is a widespread phenomenon. Both the popular press and academic studies have noted that managers and leaders in modern organizations either lack strong ethical standards (Andrews, 1989; Longnecker, 1985; Molander, 1987; Pitt & Abratt, 1986) or are willing to abandon them in the face of economic incentives or competitive pressures (Gellerman, 1986; Hosmer, 1987). For instance, while accounting firms in this country are charged with providing independent financial monitoring and reporting client mismanagement, they do so at the risk of displeasing their clients and losing lucrative service contracts (Moore, Tetlock, Tanlu, & Bazerman, 2006). As a result, "independent" auditors often have been found to be complicit in their clients' unethical practices (Levitt & Dwyer, 2002).

While we do not disagree with these views suggesting intentionality in unethical conduct, this paper presents evidence consistent with a more recent view of unethical behavior that argues that many unethical acts occur without the conscious awareness of the actor who engages in the unethical conduct (Banaji, Bazerman, & Chugh, 2003; Bazerman & Banaji, 2004). Specifically, we focus on the act of watching others' behaviors in order to guard oneself against unethical conduct.

While numerous definitions of unethical behavior exist in the literature (e.g., Brass, Butterfield, & Skaggs, 1998; Trevino, Weaver, & Reynolds, 2006), for our current purposes, we rely on Jones' (1991) broad conceptualization of unethical behavior as reflecting any action that is "either illegal or morally unacceptable to the larger community" (p. 367). Examples include violations of ethical norms or standards (whether they are legal standards or not), stealing, cheating, lying, or other forms of dishonesty. We predict and show that people are more likely to accept others' unethical behavior when ethical degradation occurs slowly, over time, rather than in one abrupt shift. Our results also speak to the psychological process underlying this phenomenon; we find that much of the effect occurs because people simply fail to notice degradation when it occurs gradually.

Acting on others' unethical behavior

Under what conditions do people tend to overlook others' unethical behavior? This question can be answered in part by research on organizational whistle-blowing (e.g., King, 1997; Near & Miceli, 1986; Miceli, Dozier & Near, 1991). Whistle-blowing is defined as the disclosure by organizational members of wrongdoing within an organization to the public or to those in positions of authority inside the organization (Near & Miceli, 1985). The literature on whistle-blowing aims to explain individual reactions to acts of perceived unethical behavior (Gundlach, Douglas & Martinko, 2003) and has identified individual, organizational, and situational factors (Near & Miceli, 1996).

Religious and ideal values (Miceli & Near, 1992; Sims & Keenan, 1998), moral standards and judgment (Miceli et al., 1991), power and justice (Near, Dworkin & Miceli, 1993), and individual locus of control (Miceli & Near, 1992) have all been shown to affect the tendency to whistle-blow. Scholars have also investigated differences in personal characteristics between employees who do or would report colleagues' unethical behavior and inactive observers (e.g., Sims & Keenan, 1998). Whistleblowers tend to have higher levels of pay, seniority, and education than inactive observers (Miceli & Near, 1984). Furthermore, whistleblowers typically have more years of service and higher levels of professional status than inactive observers (Miceli & Near, 1988).

As for contextual characteristics affecting people's reactions to others' unethical behavior, previous research has identified various organizational (e.g., whistle-blowing culture) and situational (e.g., nature of the wrongdoing) factors. For example, prior work has found that people are more likely to report wrongdoing when the organizational climate supports whistle-blowing (Near & Miceli, 1995) or when the organization is perceived to be responsive to complaints (Miceli & Near, 1988). In addition, it has been suggested that whistle-blowing is less likely when wrongdoing is ambiguous or when group members must assist each other with the task in question (Greenberger, Miceli & Cohen, 1987).

One premise of the literature on whistle-blowing is that both whistle-blowers and inactive observers do in fact notice others' unethical behavior. If this is correct, inactive observers are themselves acting unethically (e.g., Miceli, Near, & Schwenk, 1991), as in the case of internal auditors who fail to report malfeasance. However, our research calls this premise into question. We build on research on implicit biases (Banaji et al., 2003; Bazerman & Banaji, 2004), which views the majority of unethical behaviors as *unintentional* and *ordinary*, affecting virtually everyone. According to this perspective, most people may believe that their behavior is ethical and unbiased; at the same time, most of us fall prey to unconscious biases in ethically relevant domains (Banaji et al., 2003; Bazerman & Banaji, 2004), succumbing to implicit forms of prejudice, in-group favoritism, conflicts of interest, and the tendency to over-claim credit (Banaji et al., 2003).

Providing further support for the “ordinary unethical behavior” perspective, we will argue that certain unethical behaviors occur when people unconsciously “lower the bar” over time through small changes in their acceptance of others’ ethicality. We investigate the relationship between abrupt versus gradual changes in unethical behavior and the propensity of third parties to acknowledge this unethical behavior. A significant portion of unethical behavior in organizations, we believe, results from the “boiling frog syndrome” (Senge, 1994). According to this old folk warning, if you throw a frog into boiling water, it will quickly jump out. But if you put a frog in a pan of cold water and raise the temperature ever so slowly, the gradual warming will make the frog doze happily. In fact, the frog will eventually cook to death due to its inability to sense the gradual increase in water temperature.

The message of this tale is that, because its environment changes so gradually, the frog is never stimulated to take bold action to save its life. Similarly, we suggest that third parties (e.g., accountants working for clients involved in business scandals) may not react to a gradual deterioration of their clients’ ethicality. Building on the literature on implicit biases, we argue that managers or third parties might, over time, grow accustomed to unethical behavior. In particular, we examine the possibility that a gradual deterioration of others’ behavior affects people’s propensity to overlook such behavior. We present the results of four experiments in which we vary the process of deterioration of others’ behavior – either gradually, through small changes, or in a single, abrupt shift. In the studies, we show that people are less likely to act on the unethical behavior of others if the behavior develops gradually over time (along a slippery slope) rather than occurring abruptly. Our results also indicate that the failure to act on unethical behavior when it occurs gradually stems simply from a failure to notice it.

BACKGROUND AND HYPOTHESES DEVELOPMENT

Our focus on the slippery slope view of ethical erosion is suggested by prior work that has manipulated situational factors to induce individuals' compliance or commitment to certain actions as a consequence of an initial small step (i.e., the process of "getting one's foot in the door"). As shown in the studies of Freedman and Fraser (1966), often a big request (such as allowing a large, ugly "Drive Carefully" sign to be installed in front of one's house) can be facilitated via foot-in-the-door techniques (e.g., first getting the target to sign a safe-driving petition). Indeed, the theory of cognitive dissonance (Festinger, 1957) would predict that acts of minimal compliance might motivate people to subsequently adopt a consonant behavior (Cialdini, 2001).

Foot-in-the-door techniques might also apply to unethical behavior. For instance, someone might be more willing to look the other way if a peer stole a dollar a day for a year than if the peer stole $365 all at once. Similarly, an accountant might be more willing to approve the financial statements of clients who year after year misreport their revenues by small amounts in a self-serving direction than the financial statements of clients who misreport abruptly and by large amounts. This reasoning leads us to our first hypothesis:

> *Hypothesis 1:* People are less likely to act on the unethical behavior of others if the unethical behavior develops gradually over time than if the unethical behavior occurs abruptly.

One possible explanation for this hypothesis is that individuals have set thresholds of acceptability for others' unethical behavior, but fail to notice when these thresholds are crossed gradually. Abundant evidence shows that individuals are much less likely to perceive gradual changes to the status quo than they are to notice absolute levels. Specifically, research in psychophysics has shown that people easily adapt to the stimuli around them, such as smells, sounds, and ambient light levels (see, for instance, Gray, 1994; or Thorson & Biederman-Thorson, 1974). Research in visual perception has revealed the profound degree to which people are blind to gradual changes, and can only detect changes in visual stimuli when these changes occur abruptly (Levin, 2002). Similarly, research on prospect theory has highlighted the fact that people are much more sensitive to changes in their current wealth than they are to variation in their total wealth (Camerer, 2000; Kahneman & Tversky, 1979).

This insensitivity to gradual changes explains how frogs might allow themselves to be boiled. It may also explain why people might notice the ethical misconduct of others when these behaviors first occur, but fail to notice them if others' behavior becomes unethical gradually, over time. In addition, Tenbrunsel and Messick (2004) note that decision makers may be less likely to detect a series of small changes in ethicality and accuse the target of unethical behavior than if the changes occurred in one dramatic move. These arguments lead to our second hypothesis:

> *Hypothesis 2:* People's propensity to accept others' unethical behavior when ethical erosion occurs gradually stems from people's failure to notice small changes in others' unethical behavior.

Cain, Loewenstein, and Moore (2005) have used this "slippery slope" explanation for unethical behavior to describe how auditors' conduct becomes unethical. Specifically, the researchers argue that auditors may be blind to clients' internal changes in accounting practices as long as the changes are made at a slow pace – on a slippery slope. Our current experiments test a laboratory analog of this conjecture in a between-subjects design with two different conditions (the *abrupt-change* condition and the *slippery-slope* condition).

In our studies, participants are shown a series of pictures of jars filled with pennies, each picture accompanied by an estimate of the amount of money contained in the jar. Participants get the chance to play the role of "approvers," who must either approve or reject another subject's estimate. It is common knowledge that the other party—the "estimator"—will be paid more for higher estimates, provided they are approved. Approvers are also paid more for approving higher estimates than lower estimates, but risk a penalty for approving egregiously exaggerated estimates. We compare gradual with abrupt changes in the exaggeration of estimates. We measure participants' propensity to accept unethical behavior by looking at the acceptance rate of the estimates they receive.

STUDY 1

Method

Four sessions were conducted in the computer laboratory of a university in the Northeastern United States. The procedure was identical across each experimental session, and participants were randomly assigned to one of two conditions. Each experimental session lasted approximately one hour and was conducted on computers. Participants received a $10 show-up fee and had the opportunity to win up to an additional $25 during the experiment. All payoffs were expressed in U.S. dollars. At the end of the experiment, participants were asked to complete a questionnaire that asked their gender, current position, and age. They then were paid their earnings, debriefed, thanked, and dismissed.

Participants. Seventy-six individuals participated in the experiment. They were recruited using ads in which they were offered money to participate in an experiment on decision-making. During recruitment and again in the experiment's instructions, participants were told that their earnings were a function of their choices during the experimental session. Participants were not given any feedback on their decisions as estimators and approvers until after the experiment. Thirty males (39%) and 46 females (61%) participated; their average age was 28. Most participants were students from local universities (53 out of 76, 70%), either undergraduates (31 out of 53, 58%) or graduate students (22 out of 53, 42%). Thirty-eight individuals participated in the slippery-slope condition, and 38 individuals participated in the abrupt-change condition.

Procedure. Upon arrival at the computer laboratory, participants were registered and randomly assigned to terminals. The room was arranged so that participants could not see each other during the experiment. Participants received copies of the instructions, which told them how to play the experiment using the computer. The researcher also read the instructions aloud, and the participants were given an opportunity to ask questions. Thus, before the session began, participants knew what each phase entailed and how the payoff was computed.

The experiment consisted of three phases, preceded by a few practice rounds. Each phase was composed of 16 rounds on the computer (see Table 1).

Table 1
Experimental phases, conditions, and task

	Slippery-slope Condition	Abrupt-change Condition
Pre-phase	Trial Phase (three rounds) in which the participant sees pictures of jars containing pennies and is told the correct amount of money in the jar	Same as in the other condition
Phase 1	Participants play the role of estimators for 16 rounds	Same as in the other condition
Phase 2	Participants play the role of approvers for 16 rounds and are given estimates by the experimenter (estimates for **slippery-slope condition**)	Participants play the role of approvers for 16 rounds and are given estimates by the experimenter (estimates for **abrupt-change condition**)
Phase 3	Participants play the role of approvers for 16 rounds and are given estimates provided by another subject in Phase 1	Same as in the other condition

In each round, participants saw a picture of a jar containing American pennies (see Figure 1 for an example). Pictures differed from round to round. Also, pictures of jars containing the same amount of money were different. Thus, if in Round 3 of both Phase 2 and Phase 3 the jar contained $10.04, a different picture was used in each phase (with the jar containing the same amount of money) in order to avoid an experience confound. This was done by shaking the jar between photographs to change the position of the coins. To avoid biases in estimates due to the shape of the jar, the same jar was used in all rounds.

Figure 1. Example of picture used in the experiment.

Before Phase 1, participants entered a pre-phase of three practice rounds. In the practice rounds, participants did not play any role; they simply were shown a series of pictures of jars containing different amounts of pennies. The first jar contained $9.75, the second contained $10.50, and the third contained $11.25. In the practice rounds, participants were told the true amount of money in the jars. This allowed them to get a sense of the pictures they would see during the experiment. Prior work has shown that people tend to underestimate the amount of money contained in jars (see, for instance, Bazerman & Samuelson, 1983). The three practice rounds were designed to reduce this tendency toward underestimation.

In Phase 1, participants played the role of "estimator"; in both Phase 2 and Phase 3, they played the role of "approver." When playing estimators, they were asked to guess the amount of money contained in the jar pictured on their computer screen. For each picture viewed in Phase 1, they were asked to provide a dollar estimate within the range of $0-18. Participants could also specify the cents amount for their estimates. When playing approvers, they instead were asked to either accept or reject an estimate of the amount of money contained in the pictured jar; the estimate was shown below the picture. If they wanted to accept the estimate, they first had to sign a document stating that they honestly believed the estimate was accurate within a range of 10%. They could sign the document by pressing the "SIGN" button pictured on the computer screen; after pressing it, their signature would appear on the document. If they wanted to reject the estimate, they had to press the "NO" button pictured on the computer screen.

In Phase 2, the estimates came from the experimenter. In Phase 3, the estimates were those of another participant who had been shown the same picture; specifically, the estimate was randomly selected from among the estimates provided by participants in Phase 1 (during the same experimental session). When playing Phase 2 and 3, participants did not know whether the estimates came from the experimenter or from another subject.

As Table 2 shows, in Phase 2, in the abrupt-change condition, the estimate equaled the true value of the money in the jar, which was $10 plus random error for the first 10 rounds. At Round 11, while the value of the money in the jars remained at $10 plus random error, the estimates jumped to $14 plus random error; they stayed at that level for the remaining five rounds. By contrast, in the slippery-slope condition, the estimates from the abrupt-change condition were increased by 40 cents in Round 2, 80 cents in Round 3, $1.20 in Round 4, etc. Then, from Rounds 11-16, the estimates were identical to the abrupt-change condition. The true amount of money contained in the jar only varied as a function of the small random error. We added the random error because round estimates might have made participants suspicious or cognizant that the estimates were provided by the experimenter.

Table 2
Estimates provided by the experimenter in each condition and true values

Estimate	Slippery-slope condition ($)	Abrupt-change condition ($)	Randomly generated error ($)	True value ($)
1	10.01	10.01	+ 0.01	10.01
2	10.37	9.97	- 0.03	9.97
3	10.84	10.04	+ 0.04	10.04
4	11.21	10.01	+ 0.01	10.01
5	11.58	9.98	- 0.02	9.98
6	12.00	10.00	+ 0.00	10.00
7	12.44	10.04	+ 0.04	10.04
8	12.79	9.99	- 0.01	9.99
9	13.22	10.02	+ 0.02	10.02
10	13.59	9.99	- 0.01	9.99
11	14.02	14.02	+ 0.02	10.02
12	14.01	14.01	+ 0.01	10.01
13	13.98	13.98	- 0.02	9.98
14	13.99	13.99	- 0.01	9.99
15	14.02	14.02	+ 0.02	10.02
16	14.03	14.03	+ 0.03	10.03

Participants' Payoff. In both conditions, participants' payoffs were computed based on the roles they played. Each participant played 16 rounds in each phase; her payoff was calculated for each round as follows. When playing the role of estimator, each participant received 8 percent of her estimate every time the approver accepted it. Thus, the estimator had an incentive to provide a high estimate in Phase 1. Yet if the approver rejected the estimate, the estimator would receive no payment. When playing the role of approver, the payoff depended on whether the approver accepted or rejected the estimator's estimate. If she rejected it, she would receive no payment. If she approved it, she would get 4 percent of the estimate. However, 10 percent of the time (randomly determined), a third party double-checked the estimate.[1] In that case, the approver would incur a penalty of $5 if she approved an estimate that was not within 10 percent of the true value. All these payoffs were common knowledge.

Note that the approver's payoff schedule includes a conflict of interest. Approvers were paid more when they allowed inflated estimates, but at the risk of a potential penalty. This payoff scheme was designed to parallel the payoffs faced by regulators or watchdogs, such as auditors charged with verifying that organizations report truthful information to the public. These watchdogs typically have a financial incentive to please those they are supposed to be monitoring, which might mean turning a blind eye to small infractions. However, watchdogs face a substantial penalty if they are caught allowing a great deal of cheating.

Participants were told that their payoffs at the end of the experiment were computed as the sum of their performance in Phase 1 and their performance in either Phase 2 or Phase 3. Whether they were paid based on either Phase 2 or Phase 3 was predetermined by the experimenter but unknown to participants.

Dependent measure. Our hypotheses were tested on the data collected in Phase 2, in which our manipulation occurred. Phases 1 and 3 were used to avoid deceiving participants, while allowing us to manipulate the estimates provided in Phase 2. Indeed, we wanted participants to believe that they might be judging estimates made by other participants. Thus, the experiment might have consisted only of Phase 2 had we told participants that the estimates they were asked to approve came from other participants when shown the same pictured jar. However, this would have involved deception, which we wanted to avoid.

---

[1] This check was automatically made by the computer.

Our dependent measure is the number of approvals (i.e., “yes”es) of the estimator’s estimate over Rounds 11-16. As mentioned above, in the abrupt-change treatment, the abrupt change happened in Round 11. In subsequent rounds, participants in both conditions were asked to review the same estimate (see Table 2, estimate 11 through estimate 16). To test our first hypothesis, we compared the cumulative number of approvals in Rounds 11-16 in the two conditions. Based on Hypothesis 1, we expected to find a statistically significant difference for such measures, with a significantly higher number of approvals in the slippery-slope condition than in the abrupt-change condition.

Ethicality of the estimates: manipulation check. During Rounds 11-16 in Phase 2 of the experiment, participants were asked to review inaccurate estimates. Did they perceive inaccurate or exaggerated estimates as unethical? To verify that they did not have difficulty determining the accuracy and ethicality of these estimates, a manipulation check was needed. We did not include a manipulation check in the main study for fear that it would be too intrusive and arouse suspicion about the experiment’s purpose. Instead, we ran a manipulation check on a separate sample of participants. We recruited 36 individuals to complete a 10-minute study for which they were paid $5. The pilot study consisted of a short paper survey that included the same pictures we used in Round 1 and Rounds 11-16 of the main study. Specifically, participants were asked to look at seven different pictures of jars with pennies in them. Each jar was presented on a different page. The first page after the instruction included an example of a pictured jar with an estimate of the true amount of money contained in the jar. This was done so that participants had a sense of the accuracy of the estimates. On each page of the survey after the example, participants saw a picture and, under the picture, an estimate of the money contained in the jar. Participants were told that the estimates were provided by individuals who participated in another study in which they were asked to guess the amount of money contained in the jar and were paid based on the accuracy of their estimates. On each page, we also asked participants to indicate how appropriate, acceptable, and ethical the estimate was along a 7-point scale (ratings range from 1 = not acceptable [appropriate, ethical] at all, to 7 = very acceptable [appropriate, ethical]). Finally, we asked them to indicate how much they thought the estimate differed from accuracy; possible answers were: within 75% of accuracy, 50%, 25%, 10%, 5%, 3% and 1%. (The answer “within 75%” refers to an estimate that is not very accurate, while the answer “within 1%” refers to a very accurate estimate.) We found that these four scales were highly and significantly correlated (see Table 3).

Table 3
Partial correlations, controlling for jars and subjects (pilot for Study 1). Significance (2-tailed) is reported in parentheses

	Appropriate	Acceptable	Ethical	Accurate
Appropriate	1.000 .	.863 (.000)	.715 (.000)	.511 (.000)
Acceptable		1.000 .	.729 (.000)	.515 (.000)
Ethical			1.000 .	.434 (.000)
Accurate				1.000 .

For our purposes, the most important correlations are probably between accuracy and ethicality and between acceptability and ethicality: the higher the perceived inaccuracy of the estimate, the more unethical participants considered it to be (.43, $p < .001$); and the higher the perceived acceptability of the estimate, the more ethical they considered it to be (.73, $p < .001$). Given these correlations, we created just one measure for the perceived ethicality of the estimates. The scale reliability (Cronbach's alpha) for such measure was .87. A factor analysis confirmed that its items represented a single dimension. Overall, we found that the use of exaggerated estimates was a successful manipulation for unethical behavior, as participants perceived inaccurate estimates to be unethical.

Results

Phase 1 estimates. We examined the estimates that participants provided in Phase 1. Because the experimental manipulation occurred in Phase 2, we expected to find no significant difference in the estimates between the two conditions. The Phase 1 estimates were subjected to an analysis of variance in which treatment (slippery-slope vs. abrupt-change) served as the between-subjects factor and rounds (i.e., repeated measure) served as the within-subjects factor. As expected, the main effect of treatment was not significant, F (1, 74) = .03, p = .87. Thus, we combined the estimates in the two conditions and computed the mean estimate per round. Most estimates (13 out of 16) that participants provided were higher than the true value. On average, over all rounds, participants' estimates were significantly higher (M = 10.78, SD = 0.80) than the true amount of money contained in the jar (M = 10.01, SD = 0.02), t (30) = 3.84, p = .001, d = 1.36. This overestimation is probably due to the payoff scheme used in the experiment: as explained earlier, estimators had an incentive to provide high estimates in Phase 1.

Rate of approval. We computed both the absolute number of approvals and the percentage of participants approving the estimator's estimates in Phase 2 within each condition. The results of this analysis are reported in Table 4. As shown, the percentage of "yes"es was higher in the slippery-slope condition than in the abrupt-change condition in each round following the abrupt change, i.e., Rounds 11-16 (see columns labeled "% of 'yes'es"). The table also reports the number of times each subject approved the estimators' estimates over Rounds 11-16 in each condition (see columns labeled "Number of 'yes'es"). In each round, the difference in the number of approvals between conditions is statistically significant (all p-values are lower than .05). These results are consistent with our first hypothesis, that people are less likely to act on the unethical behavior of others if the unethical behavior develops gradually over time than if it occurs abruptly.

Table 4
Number of approvals and rate of approval in Rounds 11-16 for each condition, Study 1

	Slippery-slope condition		Abrupt-change condition		
Round	Number of "yes"es	% of "yes"es	Number of "yes"es	% of "yes"es	p-value
11	10	26 %	1	3 %	.003
12	23	61 %	9	24 %	.001
13	20	53 %	11	29 %	.036
14	23	61 %	12	32 %	.011
15	19	50 %	8	21 %	.008
16	23	61 %	14	37 %	.039

At the aggregate level (over Rounds 11-16), the mean rate of approval in the slippery-slope condition was 52% (SD = 0.13). In the abrupt-change condition, the rate of approval was only 24% (SD = 0.12). As for the number of approvals, on average, in the slippery-slope condition, participants approved the estimate 3.11 times over the last six rounds of Phase 2 (SD = 2.18). In the abrupt-change condition, the mean number of approvals for each participant over Rounds 11-16 was only 1.45 (SD = 1.64). This difference is statistically significant ($t(74) = 3.749$, $p < .001$, $d = 0.86$), thus supporting our first hypothesis.

The impact of Phase 1 estimates on Phase 2 approvals. To examine whether approval rate in Phase 2 was affected by the estimates participants provided in Phase 1, we used the following hierarchical model:

$$PH2\_APPROVAL_{ij} = \alpha_0 + \beta_1\ TREAT_i + \beta_2\ PH1\_EST_{ij} + subjects'\ RANDOM\ EFFECTS + \varepsilon$$

where index i refers to participants and index j refers to rounds. We only considered Rounds 11 through 16. The dependent variable was approval (yes=1, no=0) for each participant and for each considered round in Phase 2. Explanatory variables were: (i) a dummy variable indicating the experimental condition ($TREAT_i$), and (ii) the estimates participants provided in Phase 1 ($PH1\_EST_{ij}$). The parameter β_1 measured the effect of our manipulation on the propensity of participants to approve estimates in Phase 2, while the parameter β_2 measured the effect of the estimates participants provided in Phase 1 on approval in Phase 2.

This analysis reveals that participants were more likely to approve an estimate in Phase 2 when they overestimated the amount of money contained in the jar in Phase 1 ($\beta_2 = 0.04$, $t = 4.66$, $p < .0001$). Yet, even after controlling for this effect, the effect of our manipulation is still in the predicted direction and statistically significant ($\beta_1 = 0.27$, $t = 6.43$, $p < .001$): participants were more likely to approve an estimate in Phase 2 when playing in the slippery-slope condition than when playing in the abrupt-change condition.

Discussion

The results of Study 1 are consistent with Hypothesis 1: People are significantly less likely to act on the unethical behavior of others if the unethical behavior develops gradually, over time, than if it occurs abruptly. This study demonstrated how the slippery-slope phenomenon affects the acceptance of the unethical behavior of others. Our second study begins to uncover the mechanism behind this effect.

STUDY 2

While Study 1 showed that the effect of the gradual deterioration of others' ethical behavior on observers' acceptance of unethical conduct is strong, it did not provide evidence that people accept others' unethical behavior because they *fail to notice* it, as our second hypothesis predicts. It is possible that participants realized that the estimates they were asked to approve were not accurate but decided to accept them to gain the possibility of increasing their payoffs. Study 2 examines whether people notice gradual changes in others' unethical behavior.

The two experiments differ in the number of rounds participants play in the role of approvers. While participants in Study 1 played approvers in all of the rounds relevant to the purpose of this paper (i.e., Phase 2, as we will explain in detail below), Study 2 participants are first exposed to the process of degradation of the other party's estimates for a few rounds and then are asked to estimate the amount of money contained in the jar for another few rounds. This estimation task tests whether people notice the change in the other party's ethical behavior and do not react to it or whether, as our second hypothesis predicts, they fail to notice such change if it occurs gradually. Study 2 allows us to test whether or not participants' judgment is distorted by the process they go through in the first 12 rounds (either abrupt change or slippery slope). We expect the estimates of participants in the slippery-slope condition to be significantly higher than those of participants in the abrupt-change condition.

Method

Participants. As in Study 1, we recruited participants through ads offering money to participate in an experiment on decision making. Seventy-four individuals agreed to participate (42% male and 58% female). The average age of participants was 26 years (SD = 9.86). Most participants were students from local universities (78% of them), primarily undergraduates (43 out of 58 students, 74%).

Procedure. Study 2 followed the same procedure used in Study 1 with only one difference: While in Study 1 participants played the role of approver in all 16 rounds of Phase 2, in Study 2 they played the role of approver in the first 12 rounds of Phase 2 and the role of estimator in Rounds 13, 14, 15, and 16. With this new design, we can compare the estimates of participants in Rounds 13-16 of Phase 2 between conditions.

Payment. As in Study 1, participants received $10 as a show-up fee and, in addition, had the opportunity to win up to $25 during the experiment. The payoff structure was the same as in Study 1 except for Rounds 13-16 of Phase 2. Because participants played the role of estimators during those rounds, they received 8% of their estimate each time it was within 10% of the true amount of money in the pictured jar. The average overall payoff in Study 2 was $23.21 ($SD$ = 5.37).

Dependent measure. As in Study 1, we tested our hypothesis on the data collected in Phase 2, in which our manipulation occurred. In Study 2, we are interested in two dependent measures. The first is the number of approvals of the estimator's estimates in Rounds 11 and 12. Indeed, as in Study 1, the abrupt change happened in Round 11; the estimates that approvers received in Rounds 11 and 12 were the same across conditions. Thus, to find further support for the hypothesis that people are less likely to act unethically by accepting the incorrect assessments of others when the behavior of others erodes slowly rather than in one abrupt shift, we compared the cumulative number of approvals in Rounds 11 and 12 in the two conditions. Similar to our prediction in Study 1, we expected to find a statistically significant difference for such measures, with a higher number of approvals in the slippery-slope condition than in the abrupt-change condition. Our second dependent measure is given by the estimates participants provided over Rounds 13-16 while playing the role of estimators. To test the hypothesis that people accept others' unethical behavior without acting against it because they fail to notice it, we compared the average estimate in Rounds 13-16 in the two conditions. We expected to find higher estimates in the slippery-slope condition than in the abrupt-change condition.

Results

Phase 1 estimates. The estimates that participants provided in Phase 1 were subjected to an analysis of variance in which treatment (slippery slope vs. abrupt change) served as a between-subjects factor and rounds (i.e., repeated measure) served as a within-subjects factor. Replicating the results of Study 1, this analysis reveals no significant main effect of treatment, F (1, 72) = .01, p = .94. After combining the estimates from the two conditions, we computed the mean estimate per round. On average, over all the rounds, the estimates that participants provided were significantly higher (M = 10.60, SD = 0.74) than the true amount of money contained in the jar (M = 10.01, SD = 0.02), t (30) = 3.17, p = .004, d = 1.13, suggesting that participants tended to overestimate the amount of money contained in the jars. This result is consistent with the findings of Study 1.

Rate of approval. We first replicated the analysis conducted on the data from Study 1 by computing both the absolute number and rate of approval in Phase 2 within each condition. We computed such values for Rounds 11 and 12; the results of this analysis are reported in Table 5. As shown, the percentage of "yes"es was higher in the slippery-slope condition than in the abrupt-change condition in each round following the abrupt change, i.e., Rounds 11-12 (see columns labeled "% of 'yes'es"). In each round, the difference in the number of approvals between conditions is statistically significant (both p-values are lower than .05). These results show further support for our hypothesis that people are less likely to act against the unethical behavior of others if the unethical behavior develops gradually, over time, than if it occurs abruptly.

At the aggregate level (over Rounds 11 and 12), the mean rate of approval is higher in the slippery-slope condition (M = 43%, SD = 0.23) than in the abrupt-change condition (M = 12%, *SD* = 0.06), χ^2 (1, *N* = 200) = 24.10, *p* < .001, Cramer's V = .35. This result replicates the findings of Study 1.

Table 5
Absolute number and rate of approval in Rounds 11-12 for each condition, Study 2

	Slippery-slope condition		Abrupt-change condition		
Round	Number of "yes"es	% of "yes"es	Number of "yes"es	% of "yes"es	p-value
11	10	27 %	3	8 %	.032
12	22	59 %	6	16 %	< .001

Failure to notice. At the aggregate level, over Rounds 13-16, the mean estimate per participant was significantly higher in the slippery-slope condition (*M* = 11.14, *SD* = 1.32) than in the abrupt-change condition, (*M* = 10.23, *SD* = 0.74), *t* (72) = 3.65, *p* = .001, *d* = .85. This result provides at least partial evidence in support of the hypothesis that participants were less likely to notice others' unethical behavior when it gradually deteriorated over time than when it occurred abruptly. The value for the estimates provided by participants in each round over Rounds 13-16 are shown in Table 6. Differences between conditions are statistically significant in each round (all p-values lower than .05).

Table 6
Mean of estimates (in $) across participants in Rounds 13-16, Study 2. Standard deviation is reported in parentheses

Round	Slippery-slope condition	Abrupt-change condition	p-value
13	11. 38 (1.70)	10.19 (0.73)	< .001
14	11.26 (1.77)	10.28 (0.58)	.002
15	10.95 (1.31)	10.23 (1.00)	.026
16	11.06 (1.48)	10.21 (1.20)	.008

The impact of Phase 1 estimates on Phase 2 approvals. To examine whether approval rate in Phase 2 was affected by the estimates that participants provided in Phase 1, we conducted the same mixed-model analysis we ran for Study 1. We estimated the following hierarchical model:

$PH2\_APPROVAL_{ij} = \alpha_0 + \beta_1\, TREAT_i + \beta_2\, PH1\_EST_{ij} + subjects'\ RANDOM\ EFFECTS + \varepsilon$

where index *i* refers to participants and index *j* refers to rounds. We only considered Rounds 11 and 12, as Study 2 participants played the role of estimators over Rounds 13-16 of Phase 2. The effect of the estimates that participants provided in Phase 1 is positive: participants were more likely to approve an estimate in Phase 2 when they overestimated the amount of money contained in the jar in Phase 1 ($\beta_2 = 0.04$, $t = 2.60$, $p = .01$). Yet, even after controlling for this effect, the effect of our manipulation is still in the expected direction and statistically significant ($\beta_1 = 0.28$, $t = 4.19$, $p < .001$): participants were more likely to approve an estimate in Phase 2 when playing in the slippery-slope condition than when playing in the abrupt-change condition.

Discussion

The results of Study 2 are consistent with both of our hypotheses and also further support the findings from Study 1. Namely, people are more likely to act unethically by accepting the incorrect assessments of others when the behavior of others erodes slowly rather than in one abrupt shift, and this effect is at least partially a consequence of people's failure to notice the change in others' unethical behavior over time. We had been concerned that participants would correctly perceive that the estimates provided by others were not accurate but nevertheless decide to approve these estimates. Yet Study 2 provides evidence that this is not the case. The process that participants went through in the first 10 rounds (reviewing estimates that changed in a slippery-slope manner versus a single abrupt change) significantly affected their estimates of the amount of money contained in pictured jars in subsequent rounds. This suggests that people fail to notice the gradual deterioration of others' unethical behavior.

One might be concerned that the methodology used in this study could be the source of an alternative explanation for our findings. We suggested that people are more likely to accept the unethical behavior of others when others' behavior changes gradually rather than abruptly. An alternative explanation might be that the number of exposures to unethical estimates increased participants' willingness to accept those estimates. In both studies, participants in the slippery-slope condition were exposed to a higher number of unethical overestimates than those in the abrupt-change condition. Thus, it is possible that, as the number of unethical estimates increases, so does acceptance of those estimates. Yet, had people in the slippery-slope condition perceived that the received estimates were not accurate, we likely would have observed a lower rate of approval in the slippery-slope condition than in the abrupt-change condition. Indeed, while participants in the abrupt-change condition in the first 10 rounds of Phase 2 were asked to either approve or reject estimates that were within the 10% range of accuracy, participants in the slippery-slope condition had to either approve or accept estimates, seven of which were outside the 10% range of accuracy. Thus, these participants had more opportunities to react to others' unethical behavior. In essence, the effect of the exposure to a higher number of unethical estimates in the slippery-slope condition goes in the opposite direction of our prediction and should, if anything, work against the strength of our findings.

STUDY 3

In Study 2, we began to uncover the psychological process behind the apparent willingness of people to accept others' unethical behavior when such behavior deteriorates gradually. With Study 3, we provide additional support for the mechanism behind the demonstrated effect by introducing a no-incentive condition. Based on our second hypothesis, we predict that the main effect between gradual and abrupt changes in estimate will be maintained even in the no-incentive condition. This contrasts with the possibility that our entire effect is driven by the intentional unethical behavior of self-interested actors to disregard overestimates.

In addition to the incentives manipulation, Study 3 participants are asked to complete an additional task at the end of Phase 2, i.e., a 12-item word-completion task. Word-fragment completion tests have been previously used in the psychology literature and have been shown to be extremely sensitive to the activation of either recently encountered constructs (Tulving, Schacter, & Stark, 1982) or self-generated ones (Bassili & Smith, 1986). The word completion task allows us to test whether or not participants' choice of words is influenced by the process they go through in reviewing estimates in Phase 2 of Study 3 (either abrupt-change or slippery-slope). We expect participants in the slippery-slope condition to use significantly fewer words related to unethical behavior than participants in the abrupt-change condition.

Method

Design and procedure. Study 3 employs a 2 (change in estimates: slippery-slope vs. abrupt-change) x 2 (incentives: incentives vs. no incentives) design in which both change in estimates and incentives are between-subject factors. In each condition, the procedure used in Study 3 followed the one used in Study 1, with two differences. First, while participants in Study 1 played all three phases sequentially with no interruption, in Study 3 we added an additional task at the end of Phase 2. Specifically, participants were asked to complete a 12-item word-completion task in which each word stem (e.g., F R _ _ _) could be completed with an unethical behavior-related word (e.g., fraud) or with other words unrelated to ethicality (e.g., frame, fresh, frogs, fries, fruit, front, etc.). This task is designed to test whether being in the slippery-slope condition is more likely to activate thoughts relevant to relaxed ethical standards than being in the abrupt-change condition. For instance, when people are asked to complete the word "C H _ _ _," if they are thinking about cheating, they will be more likely to fill in "cheat" than alternatives such as "chime" or "chimp."

Study 3 also differed from Study 1 by including a no-incentive condition. The no-incentive condition differed from the incentive condition only in its payoff structure. As specified below, while the payoff in the incentive condition varied depending on participants' choices during the experiment, participants in the no-incentive condition were paid a fixed fee for their participation.

Participants. One hundred forty-eight individuals (49% male and 51% female) participated in Study 3. Participants were recruited using ads offering money to participate in an experiment on decision making. The average age of participants was 29 years (*SD* = 12). Most participants were students from local universities (74% of them), primarily undergraduates (87 out of 109 students, 80%). At the beginning of the experiment, participants were randomly assigned to one of the four conditions; Table 7 shows the number of participants in each condition.

Table 7
Number of participants by condition, Study 3

		Change in estimates	
		Abrupt change	Slippery slope
Incentives	No incentives	38	37
	Incentives	37	36

Payment. Participants in the incentive conditions received $10 as a show-up fee and, in addition, they had the opportunity to win up to $25 during the experiment. The payoff structure was the same as that used in Study 1. Participants in the no-incentive conditions received $10 for their participation in addition to the \$10 show-up fee. The average overall payoff in Study 3 was \$17 (*SD* = 4.88).

Results

Phase 1 estimates. As in Study 1 and 2, we examined the estimates that participants provided in Phase 1. These estimates were subjected to an analysis of variance in which change in estimates (slippery slope vs. abrupt change) and incentives (incentives vs. no incentives) served as the between-subjects factor and rounds (i.e., repeated measure) served as the within-subjects factor. None of these factors was significant.[2]

Thus, we combined the estimates in the four conditions and computed the mean estimate per round. On average, over all of the rounds, the estimates that participants provided were significantly higher ($M = 10.69$, $SD = 0.77$) than the true amount of money contained in the jar ($M = 10.01$, $SD = 0.02$), t (30) = -3.53, $p = .001$, $d = 1.25$. Note that this overestimation is significantly different from zero in both the incentives conditions ($M = 0.77$, $SD = 0.87$; t (30) = 3.54, $p = .001$) and the no-incentives conditions ($M = 0.60$, $SD = 0.69$; t (30) = 3.49, $p = .002$) and is larger in the former, as we might expect given the payoff structure (but not significantly: t (30) = -0.63, $p = .53$).

Rate of approval. To compare the results of Study 3 with those of Studies 1 and 2, we computed the absolute and percentage number of approvals in Phase 2 within each condition. At the aggregate level (over Rounds 11-16), the mean rate of approval was higher in the slippery-slope/no-incentives ($M = 31\%$, $SD = 0.11$) condition than in the abrupt-change/no-incentives

[2] Main effect for change in estimates: F (1, 144) = .64, $p = .43$, $\eta^2 = .004$; main effect for incentives: F (1, 144) = .31, $p = .58$, $\eta^2 = .002$; interaction effect: F (1, 144) =.02, $p = .89$, η^2 = .000.

condition (*M* = 16%, *SD* = 0.10), χ^2 (1, *N* = 450) = 16.66, *p* < .001, Cramer's V = .19. As for the two conditions with incentives, the rate of approval was 56% (*SD* = 0.12) in the slippery-slope/incentives condition and only 32% (*SD* = 0.11) in the abrupt-change/incentives condition, χ^2 (1, *N* = 438) = 30.75, *p* < .001, Cramer's V = .27. These results replicate the findings of Study 1 and 2 and provide additional support for Hypothesis 1.

Similar results are obtained by comparing the rate of approval across conditions in each round over Rounds 11-16. Table 8 reports the number and percentage of approvals by round, pooled by condition. In the incentives conditions, differences in rate of approval between the slippery-slope and abrupt-change treatments are statistically significant in each round (all p-values are lower than .05). However, in the no-incentives conditions, such differences are statistically significant only in a few rounds.

Table 8
Number of approvals in Rounds 11-16 for each condition, Study 3. Rates of approval are reported in parentheses

	No Incentives			Incentives		
Round	Slippery slope	Abrupt change	p-value	Slippery slope	Abrupt change	p-value
11	4 (11 %)	1 (3 %)	.16	13 (36 %)	5 (14 %)	.025
12	12 (32 %)	8 (21 %)	.27	23 (64 %)	10 (27 %)	.002
13	11 (30 %)	6 (16 %)	.15	23 (64 %)	15 (41 %)	.046
14	12 (32 %)	8 (21 %)	.27	24 (67 %)	16 (43 %)	.044
15	13 (35 %)	2 (5 %)	.001	19 (53 %)	10 (27 %)	.025
16	17 (46 %)	9 (24 %)	.043	23 (64 %)	14 (38 %)	.026

The average number of approvals over Rounds 11-16 was used as the dependent variable in a univariate ANOVA with change in estimates (slippery slope vs. abrupt change) and incentives (incentives vs. no incentives) as between-subjects factors. This analysis reveals a significant main effect for both factors. On average, participants approved a higher number of estimates in the incentives condition (*M* = 2.67, *SD* = 2.36) than in the no-incentives condition (*M* = 1.37, *SD* = 1.72), $F(1,144) = 16.23$, $p < .001$, $\eta^2 = .101$. Furthermore, participants approved a higher number of estimates in the slippery-slope condition (*M* = 2.66, *SD* = 2.24) than in the abrupt-change condition (*M* = 1.39, *SD* = 1.87), $F(1,144) = 15.57$, $p < .001$, $\eta^2 = .098$. The interaction term was not significant ($F(1,144) = .89$, $p = .35$, $\eta^2 = .006$).

Given the results of this analysis, we compared the rate of approval within both incentive conditions and both change-in-estimates conditions. As Table 9 shows, the percentage of "yes"es was higher in the slippery-slope condition than in the abrupt-change condition in each round following the abrupt change, i.e., Rounds 11-16. The table also reports the number of times each subject approved the estimator's estimate over Rounds 11-16 in each condition. In each round, the difference in the number of approvals between the two conditions is statistically significant (all p-values are lower than .05). These results are consistent with our first hypothesis, that people are less likely to act against the unethical behavior of others if the unethical behavior develops gradually, over time, than if the unethical behavior occurs abruptly.

Table 9
Number of approvals and rate of approval in Rounds 11-16 for the slippery-slope and abrupt-change conditions, Study 3

	Slippery-slope condition		Abrupt-change condition		
Round	Number of "yes"es	% of "yes"es	Number of "yes"es	% of "yes"es	p-value
11	17	23 %	6	8 %	.010
12	35	48 %	18	24 %	.002
13	34	47 %	21	28 %	.019
14	36	49 %	24	32 %	.032
15	32	44 %	12	16 %	<.001
16	40	55 %	23	31 %	.003

Finally, Table 10 reports the results for the incentives vs. no-incentives comparisons. In each round, the difference in the number of approvals between the two conditions is statistically significant (all p-values are lower than .05), suggesting that the presence of incentives exacerbates the demonstrated effect of failing to act against others' gradual ethical erosion.

Overall, these results suggest that the main effect between gradual and abrupt change in estimates found in Study 1 and 2 is maintained even when incentives are not present. This provides stronger evidence that the effect is due in part to the failure of people to notice gradual erosion in others' ethical behavior, as predicted in Hypothesis 2.

Table 10
Number of approvals and rate of approval in Rounds 11-16 for the incentives condition, Study 3

	No-incentives condition		Incentives condition		
Round	Number of "yes"es	% of "yes"es	Number of "yes"es	% of "yes"es	p-value
11	5	7 %	18	25 %	.002
12	20	27 %	33	45 %	.019
13	17	23 %	38	52 %	<.001
14	20	27 %	40	55 %	<.001
15	15	20 %	29	40 %	.008
16	26	35 %	37	51 %	.049

Failure to notice. To compare the results of Study 3 with those of Studies 1 and 2, we also examine whether approval rate in Phase 2 was affected by the estimates that participants provided in Phase 1. We estimated the following hierarchical model:

$$PH2\_APPROVAL_{ij} = \alpha_0 + \beta_1 \, TREAT_i + \beta_2 \, INCENTIVES_i + \beta_3 \, PH1\_EST_{ij} + subjects' \; RANDOM \; EFFECTS + \varepsilon$$

where index *i* refers to participants and index *j* refers to rounds. Explanatory variables were: (i) a dummy variable indicating the experimental condition for change in estimates (slippery slope vs. abrupt change) ($TREAT_i$), (ii) a dummy variable indicating the experimental condition for incentives (*$INCENTIVES_i$*), and (iii) the estimates participants provided in Phase 1 ($PH1\_EST_{ij}$). The effect of the estimates provided in Phase 1 is positive: participants were more likely to approve an estimate in Phase 2 when they overestimated the amount of money contained in the jar in Phase 1 ($\beta_3 = 0.06$, $t = 10.01$, $p < .001$). Yet, even after controlling for this effect, the effect of our manipulations is still in the expected direction and statistically significant: participants were more likely to approve an estimate in Phase 2 when playing in the slippery-slope condition than when playing in the abrupt-change condition ($\beta_1 = 0.22$, $t = 7.83$, $p < .001$), and they were also more likely to approve an estimate in Phase 2 when playing in the incentives condition than when playing in the no-incentives condition ($\beta_2 = 0.20$, $t = 7.09$, $p < .001$).

Word-completion task. We conducted an ANOVA with the number of unethical-related words as the dependent variable and incentives (incentives vs. no incentives) and change in estimates (slippery slope vs. abrupt change) as the between-subjects factors. The mean values by condition are reported in Figure 2.

Figure 2. Mean number of words with unethical-related meaning by condition, Study 3. Error bars represent standard errors.

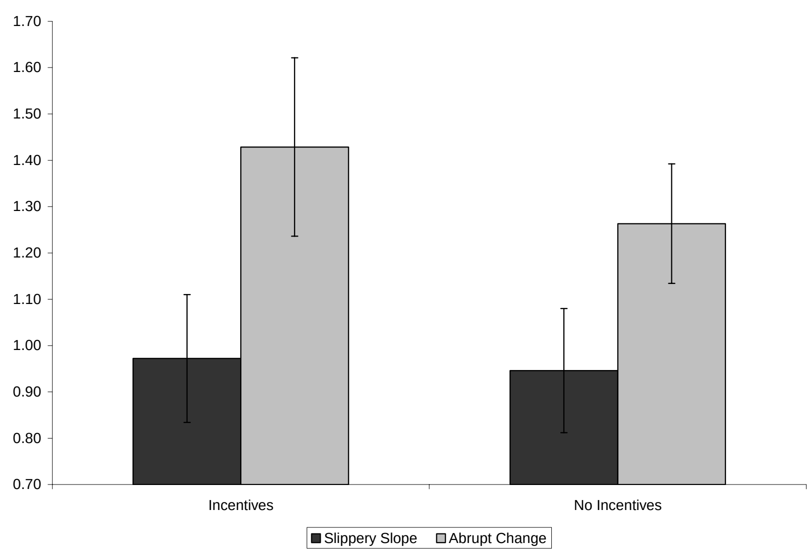

This analysis reveals a significant main effect for change in estimates: as we expected, participants in the abrupt-change condition used a significantly higher number of words related to unethical behavior (*M* = 1.34, *SD* = 0.98) than did participants in the slippery-slope condition (*M* = 0.96, *SD* = 0.81), $F(1,142) = 6.75$, $p = .01$, $\eta^2 = .05$. No other significant effect was found ($p = .52$ for the main effect of incentives and $p = .64$ for the interaction effect).

Discussion

The results of Study 3 replicate those of Study 1 and 2: People are more likely to act on the unethical behavior of others when changes in others' unethical behavior occur abruptly than when they develop gradually over time, thus supporting Hypothesis 1. Furthermore, the results of Study 3 show that people's failure to act upon others' unethical behavior can be explained by their failure to notice the changes in unethical behavior when these changes occur gradually, thus supporting Hypothesis 2.

In Study 3, we added a no-incentive manipulation to test whether incentives fuel the implicit process through which people fail to notice gradual changes in the unethical behavior of others and, as a result, fail to act upon it. The main effect between gradual and abrupt change in estimates was maintained even when incentives were not present, providing stronger support for the argument that people do not notice gradual changes in others' unethical behavior.

STUDY 4

The three studies presented thus far were designed on the assumption that providing exaggerated estimates is an example of unethical behavior. Do people perceive this behavior as unethical? The manipulation check described in Study 1 suggests that they do. Like this pilot study, Study 4 includes measures of the extent to which participants perceive incorrect estimates of the amount of money contained in the pictured jars to be unethical. In addition, Study 4 examines whether people realize that given estimates are inaccurate and investigates how ethicality ratings may change when exposure to unethical acts increases.

Method

Participants. As in the other three studies, we recruited participants through ads offering money to participate in an experiment on decision making. Thirty-two individuals agreed to participate (44% male and 56% female). The average age of participants was 27 years (*SD* = 11.24). Most participants (69%) were students from local universities.

Procedure. At the beginning of the experiment, participants were randomly assigned to one of two experimental conditions, with 16 participants in the slippery-slope condition and 16 in the abrupt-change condition. As in the pilot conducted for Study 1, participants had to complete a short paper survey that included the same pictures we used in Rounds 1-16 of the other experiments described above. Specifically, participants were asked to look at sixteen different pictures of jars with pennies in them. Each jar was presented on a different page. As in the pilot, the first page after the instruction included an example of a pictured jar with an estimate of the true amount of money contained in the jar.

On each page of the survey after the example, participants were shown a color picture of a jar and, under the picture, an estimate for the amount of money contained in the jar. Participants were told that the estimates had been provided by another participant from a previous session who was paid based on the accuracy of her estimates. In reality, the estimates and the pictures matched those from Phase 2 of Study 1 (as reported in Table 2).

On each page, we also asked participants to indicate how appropriate, acceptable, and ethical the estimate was. Specifically, participants were asked to answer the following five questions by circling the most appropriate answer.

1. *On a 1-7 scale please indicate how appropriate you think the estimate is (1 = not appropriate at all; 7 = very appropriate)*
2. *On a 1-7 scale please indicate how acceptable you think the estimate is (1 = not acceptable at all; 7 = very acceptable)*
3. *On a 1-7 scale please indicate how ethical you think the estimate is (1 = not ethical at all; 7 = very ethical)*
4. *Please indicate how much you think the estimate differs from accuracy. Possible answers were: within 1%, within 5%, within 10%, within 25%, within 50%, and within 75%.*
5. *Please indicate how much money you think is contained in Jar #1 (in $).*

After reviewing the 16 jars, participants were asked a few demographic questions (age, gender, nationality, and status), then thanked, debriefed, paid $20 each for their participation, and dismissed.

Results and Discussion

Ethicality of estimates. As in the pilot study, we found that the four scales used in the paper survey were highly and significantly correlated (see Table 11, Columns 1 through 4). For instance, the higher the perceived inaccuracy of the estimate, the more unethical participants considered it to be (.41, $p < .001$); and the higher the perceived acceptability of the estimate, the more ethical participants considered it to be (.58, $p < .001$). Given these correlations, we created just one measure for the perceived ethicality of the estimates. The scale reliability (Cronbach's alpha) for such measure was .86. A factor analysis confirmed that its items represented a single dimension. In sum, we found that participants perceived inaccurate estimates as inappropriate, unacceptable, and unethical, thus confirming the results of our pilot study.

Additional analyses reveal a significant negative correlation between the ethicality scale and the absolute difference between given and reported estimates. Specifically, the closer participants' reported estimates were to the given estimates, the higher the ethical ranking the given estimate received (see Table 11, Column 5). Note that, here, we use the term "given estimates" to refer to the estimates participants saw under the pictured jar and the term "reported estimates" to refer to the estimates participants provided for the amount of money contained in the pictured jar.

Table 11
Partial correlations, controlling for jars, subjects, and condition. Significance (2-tailed) is reported in parentheses

	Appropriate (1)	Acceptable (2)	Ethical (3)	Accurate (4)	Abs(Given-Reported) (5)
Appropriate	1.000	.909	.546	.484	-0.516
	.	(.000)	(.000)	(.000)	(.000)
Acceptable		1.000	.579	.471	-0.522
		.	(.000)	(.000)	(.000)
Ethical			1.000	.407	-0.412
			.	(.000)	(.000)
Accurate				1.000	-0.660
				.	(.000)
Abs(Given-Reported)					1.000
					.

Overestimation. For each jar and for each participant, we computed the difference between the participant's reported estimate for the amount of money contained in the jar and the true value. A positive value for such difference indicates that the participant overestimated the correct amount, while a negative value suggests underestimation. To test whether our experimental manipulation affected people's estimates for the amount of money contained in the pictured jars, we conducted a repeated-measure ANOVA with change in estimates (abrupt change vs. slippery slope) as the between-subjects factor and jars as the within-subjects factor (repeated measure). This analysis reveals a significant main effect for change in estimates:

participants in the abrupt-change condition tended to underestimate the amount of money contained in the pictured jar (mean difference between the reported estimate and the true value = -0.12, *SD* = 2.12), while participants in the slippery-slope condition tended to overestimate it, (*M* = 1.10, *SD* = 1.58), F (1, 30) = 4.99, $p < .05$, $\eta^2 = .14$. This suggests that participants' reported estimates were affected by the given estimates – which changed, picture by picture, based on the experimental condition in which they participated.

Failure to notice. For each jar and for each participant, we computed the absolute difference between the reported estimate for the amount of money contained in the jar and the given estimate. Figure 3 shows the mean values for such difference across participants in each round, by experimental condition. While the absolute difference follows a linear trendline in the slippery-slope condition ($y = 0.1406x + 0.5691$, $R^2 = 0.64$), it follows an exponential trendline in the abrupt-change condition ($y = 0.5395e^{0.1272}x$, $R^2 = 0.70$). This suggests that participants were less likely to notice the inaccuracy of the estimates when the estimates gradually deteriorated over time than when they changed abruptly, as participants reported estimates closer to the given estimates in the slippery-slope condition (resulting in smaller values for the absolute difference between given and reported estimates in Rounds 11-16, as shown in Figure 3) than in the abrupt-change condition (resulting in larger values for the absolute difference in Rounds 11-16). If repeated exposure to unethical behavior explains the results of Studies 1, 2, and 3 supporting Hypothesis 1, then we would not observe the linear trendline in the slippery-slope condition with a positive coefficient for x.

As shown in Figure 3, before the abrupt shift occurring in Round 11, the absolute difference between the given estimate and the reported estimate in most rounds was higher in the slippery-slope condition than in the abrupt-change condition; the opposite relationship is observed after Round 11. At the aggregate level, over Rounds 11-16, the mean difference was significantly larger in the abrupt-change condition (*M* = 4.21, *SD* = 0.43) than in the slippery-slope condition, (*M* = 2.50, *SD* = 0.46), $t(10) = -6.66$, $p < .001$, $d = 3.84$. This suggests that participants failed to notice the inaccuracy of the estimates when the estimates eroded gradually than when they changed abruptly.

Figure 3. Mean values for the absolute difference between given estimate and reported estimate by condition (Study 4). Error bars represent standard errors.

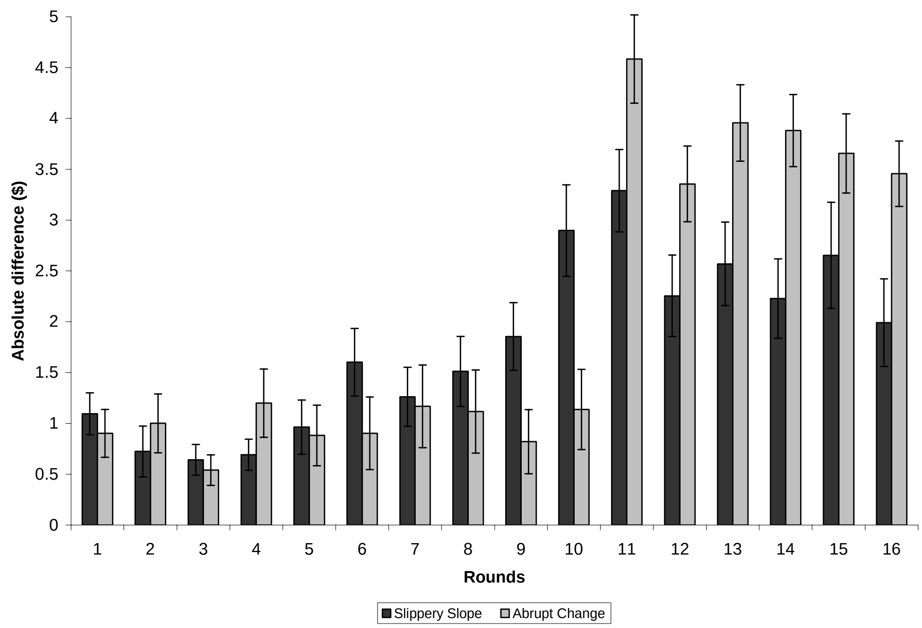

Table 12
Mean of absolute difference between given estimate and reported estimate across participants, for jars 11-16 (Study 4). Standard deviations are reported in parentheses

Jar	Slippery-slope condition	Abrupt-change condition	p-value
11	3.29 (1.62)	5.01 (2.34)	.02
12	2.26 (1.60)	3.77 (2.17)	.03
13	2.57 (1.64)	4.27 (1.88)	.01
14	2.23 (1.56)	4.28 (1.91)	.003
15	2.65 (2.09)	3.99 (1.98)	.07
16	1.99 (1.72)	3.99 (2.45)	.01

Table 12 shows the value of the estimates provided by participants on each page for jars 11-16. As shown, differences between conditions are statistically significant in each round. The nature and significance of the results do not change if the analysis is conducted with participants' reported estimates instead of the difference between given estimates and participants' reported estimates.

General Discussion and Conclusions

Where is the line between not accepting any exception to ethicality and behaving unethically or overlooking another person's unethical behavior? What does it take for ordinary people to slide across this line? We like to believe that only a few bad apples cross to the other side. In fact, under certain conditions, most of us can be expected to engage in unethical behavior. Here, our goal was to explore one condition that can lead ordinary people to cross that line. In particular, we hypothesized that people are less likely to act against the unethical behavior of others if the unethical behavior develops gradually, over time, than if the unethical behavior occurs abruptly. We also predicted that such failure to act is due to the failure to notice small changes in ethical erosion.

Imagine that an accountant employed by a well-known accounting firm is in charge of the audit of a large corporation with a strong reputation. The accounting firm and the client have an excellent relationship, and the accounting firm receives tens of millions of dollars in fees from the client each year. For three years, the accountant in charge has viewed and approved the client's high-quality, extremely ethical financial statements. Suddenly, the corporation begins stretching the limits of the law, and even breaking it in certain areas. Would the accountant notice these transgressions, and would she sign a statement certifying that the financial statements were acceptable according to government regulations? Our findings suggest that most accountants would notice the transgressions and would refuse to sign the statement.

Now suppose that the accountant saw and approved of high-quality, highly ethical financial statements for one year, after which the corporation begins stretching the law in a few areas, without appearing to break the law. In the third year, the firm stretched the ethicality of its returns a bit more. Some of the company's accounting decisions now possibly violated federal accounting standards. By the fourth year, the corporation stretched the limits of the law in more areas and occasionally broke the law. Based on the results of our study, we believe that auditors are much less likely to notice and refuse to sign the statements in the second version than in the first version, even if the unethical behavior is identical in year four of both stories.

Our studies show that when the unethical behavior of others develops gradually, over time, instead of occurring abruptly, people are more likely to overlook such behavior. Other situational factors might facilitate the crossing of the line into unethical behavior; these will be the focus of further studies. For instance, does creating an environment in which people feel pressured to justify their actions or decisions to others lead people to behave unethically? Several strategies can help people cope with such pressure (Tetlock, 1992). One of them is the acceptability heuristic, according to which people "adopt positions likely to gain the favor of those to whom they feel accountable" (Tetlock, 1992: 340). As demonstrated in prior work, at times people also use the acceptability heuristic when coping with ethical dilemmas. For instance, Pennington and Schlenker (1999) show that the audience to whom one is accountable can shape one's views and behaviors regarding an ethical decision. Related studies by Brief, Dukerich, and Doran (1991) and by Smith-Crowe, Umphress, Brief, Tenbrunsel, and Gee (2002) support such "conformity effects": when making decisions, people care less about the content of decisions (i.e., whether or not the decisions are ethical) than about the potential acceptance of such decisions by those to whom they are accountable. Similar effects of compliance were earlier shown by Milgram's studies (1974) on obedience to authority.

Another situational factor worthy of exploration in future research is diffusion of responsibility. When multiple people are responsible for a decision, rather than just one person, they are less likely to take responsibility for their unethical acts, thereby leading to an increase in unethical behavior (Bandura, 1999). Bandura (1999) notes that division of labor, or the subdivision of job responsibilities, is one mechanism that can lead to diffusion of responsibility. The task implemented in our studies may be an example of how diffusion of responsibility operates. Our task was subdivided: one person estimates, another accepts or rejects those estimates (participants), and another checks those estimates 10% of the time. It is possible that, in both the slippery-slope and the abrupt-change conditions, participants were more willing to accept incorrect estimates because they felt they may not be the "final say" as to whether or not the estimates were correct. Specifically, participants may have been more likely to agree with incorrect estimates because they knew that, 10% of the time, someone else would double-check the estimate. Also, participants may have felt even less responsibility because they themselves did not make the unethical estimates. Therefore, the segmentation of the task could have led to an increase in acceptance of unethical estimates.

Another important topic for future research concerns the limitations and boundary conditions of the effects we documented. As Ashforth and Anand (2003) indicate, unethical acts "become an integral part of day-to-day activities to such an extent that individuals may be unable to see the inappropriateness of their behaviors" (p. 4). Similarly, Brief, Buttram, and Dukerick (2001) describe institutionalization of corruption as occurring, in part, because of "repetitive cycles of routinized activity" (p. 480). Thus, another aspect of the slippery-slope thesis is the gradual inability of people to view the ethicality of their actions. Future studies should include a measure of the extent to which participants perceived acceptance of the incorrect amounts to be unethical and should investigate how these ethicality ratings might change when exposure to the unethical acts increases. We suspect that participants who fail to see the implications of their actions will be more willing to accept the unethical estimates and that as participants see more and more unethical estimates, they are less likely to perceive the ethical implications of their actions.

Future research should investigate potential mediators of the effect we demonstrate. For instance, it would be fruitful to explore more closely the psychological process underlying the effect of gradual deterioration of others' unethical behavior on people's acceptance of that behavior. We found that gradual changes over time in others' unethical misconduct decrease the likelihood of peers noticing such behavior and thus taking actions to stop it. Future research should explore how small and how slow the erosion in others' ethicality has to be to remain unnoticed. Future research might also investigate our effect using different contexts. For instance, research might consider settings that affect people's well-being more directly, such as the gradual deterioration of the environment due to pollution or FDA approvals of drugs from companies whose quality-control processes have slowly worsened over time.

We close by stressing that we believe and support research that identifies and attempts to reduce intentional unethical behavior. However, we believe that the unintentional aspects of ethical misconduct are also worthy of study. Accordingly, we have offered an additional perspective for understanding the psychology of this ordinary unethical behavior.

References

Andrews, K.R. (1989). Ethics in practice. *Harvard Business Review*, September-October, 99-104.

Ashforth, B.E., & Anand, V. (2003). The normalization of corruption in organizations. *Research in Organizational Behavior,* 25, 1-52.

Banaji, M.R., Bazerman, M.H., & Chugh D. (2003). How (un)ethical are you? *Harvard Business Review.* December, 81 (12), 56-64.

Bandura, A. (1999). Moral disengagement in the perpetration of inhumanities. *Personality & Social Psychology Review*, 3, 193-209.

Bassili, J. N., & Smith, M. C. (1986). On the spontaneity of trait attribution: Converging evidence for the role of cognitive strategy. *Journal of Personality and Social Psychology*, 50, 239-245.

Bazerman, M.H., & Banaji, M.R. (2004). The social psychology of ordinary ethical failures. *Social Justice Research*, 17(2), 111-115.

Bazerman, M. H., & Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution*, 27, 618-634.

Brass, D. J.; Butterfield, K. D.; and Skaggs, B.C. (1998). Relationships and Unethical Behavior: A Social Network Perspective, *Academy of Management Review*, 23(1), 14-31.

Brief, A.P., Buttram, R.T., & Dukerich, J.M. (2001). Collective corruption in the corporate world: Toward a process model. In M.E. Turner (Ed.), *Groups at Work: Theory and Research. Applied Social Research* (pp. 471-499). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Brief, A.P., Dukerich, J.M., & Doran, L.I. (1991). Resolving ethical dilemmas in management: Experimental investigations of values, accountability and choice. *Journal of Applied Social Psychology*, 17, 68-79.

Cain, D.M., Loewenstein, G. & Moore, D.A. (2005). The dirt on coming clean: Perverse effects of disclosing conflicts of interest. *Journal of Legal Studies*, 34(1), 1-27.

Camerer, C. F. (2000). Prospect theory in the wild: Evidence from the field. In D. Kahneman & A. Tversky (Eds.), *Choices, values, and frames* (pp. 288-300). New York: Russell Sage Foundation.

Cialdini, R. B. (2001). *Influence: Science and Practice* (4th Ed). Boston, MA: Allyn & Bacon.

Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford.

Freedman, J.L. & Fraser, S.C. (1966). Compliance without pressure: The foot-in-the-door technique. *Journal of Personality and Social Psychology*, 4, 195-202.

Gellerman, S.W. (1986). Why 'good' managers make bad ethical choices. *Harvard Business Review*, July-August, 85-90.

Gray, P. (1994). *Psychology* (2nd Ed). New York, NY: Worth Publishers.

Greenberger, D. B., Miceli, M. P., & Cohen D. J. (1987). Oppositionists and group norms: The reciprocal influence of whistle-blowers and co-workers. *Journal of Business Ethics*, 6, 527-542.

Gundlach, M. J., Douglas, S. C. & Martinko, M. J. (2003). The decision to blow the whistle: A social information processing framework. *Academy of Management Review,* 28(1), 107-123.

Hosmer, L. T. (1987). The Institutionalization of Unethical Behavior. *Journal of Business Ethics*, 6(6), 439-447.

Jones, T.M. (1991). Ethical decision making by individuals in organizations: An issue contingent model. *Academy of Management Review*, 16, 366-395.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263-291.

King, G. (1997). The effects of interpersonal closeness and issue seriousness on blowing the whistle. *Journal of Business Communication,* 34, 419-436.

Levin, D. T. (2002). Change blindness as visual metacognition. *Journal of Consciousness Studies, 9*(5-6), 111-130.

Levitt, A., & Dwyer, P. (2002). *Take on the street: What Wall Street and corporate America don't want you to know, and what you can do to fight back*. New York: Pantheon.

Longnecker, J.G. (1985). Management Priorities and Management Ethics. *Journal of Business Ethics*, 4, 65-70.

Miceli, M. P., Dozier, J. B., & Near, I. P. (1991). Blowing the whistle on data fudging: A controlled field experiment. *Journal of Applied Social Psychology,* 21, 271-295.

Miceli, M. P., & Near, J. P. (1984). The relationship among beliefs, organizational position, and whistle-blowing status: A discriminant analysis. *Academy of Management Journal,* 27(4), 687-705.

Miceli, M. P., & Near, J. P. (1988). Individual and situational correlates of whistle-blowing. *Personnel Psychology*, 41, 267-281.

Miceli, M. P., & Near, J. P. (1992). *Blowing the Whistle*. Macmillan Publishing Company, Riverside, NJ.

Miceli, M. P., Near, J. P., & Schwenk, C. R. (1991). Who blows the whistle and why? *Industrial and Labor Relations Review*, 45(1), 113-130.

Milgram, S. (1974). *Obedience to Authority: An Experimental View*. New York, NY: Harper & Row.

Molander, E. (1987). A paradigm for design, promulgation and enforcement of ethical Codes. *Journal of Business Ethics*, 6, 619-631.

Moore, D.A., Tetlock, P.E., Tanlu, L. & Bazerman, M. H. (2006). Conflicts of interest and the case of auditor independence: moral seduction and strategic issue cycling. *Academy of Management Review,* 31(1), 10-29.

Near, J. P., Dworkin, T. M., & Miceli, M. P. (1993). Explaining the whistle-blowing process: Suggestions from power theory and justice theory. *Organization Science,* 4, 393-411.

Near, J. P., & Miceli, M. P. (1995). Whistle-blowing: myth and reality. *Journal of Management*, 22(3), 507-526.

Near, J. P., & Miceli, M. P. (1995). Effective whistleblowing. *Academy of Management Review,* 20, 679–708.

Near, J. P., & Miceli, M. P. (1985). Organizational dissidence: The case of whistle-blowing. *Journal of Business Ethics,* 4, 1-16.

Near, J. P., & Miceli, M. P. (1986). Retaliation against whistleblowers: Predictors and effects. *Journal of Applied Psychology,* 71, 137-145.

Pennington, J. & Schlenker, B.R. (1999). Accountability for consequencial decisions: Justifying ethical judgments to audiences. *Personality and Social Psychology Bulletin*, 25, 1067-1081.

Pitt, L. F. & Abratt, R. (1986). Corruption in business—Are management attitudes right? *Journal of Business Ethics*, 5, 39-44.

Senge, P. M. (1994). *The Fifth Discipline: The Art and Practice of the Learning Organization*. New York, NY: Currency Doubleday.

Sims, R. L. & Keenan, J. P. (1998). Predictors of External Whistleblowing: Organizational and Intrapersonal Variables. *Journal of Business Ethics,* 17(4), 411–421.

Smith-Crowe, K., Umphress, E. E., Brief, A., Tenbrunsel A., & Gee, J. (2002). Cooking the books: The interactive effect of accountability and social dominance orientation on fraudulent financial reporting. *Unpublished manuscript.*

Tenbrunsel, AE, & Messick, DM , (2004). Ethical fading: The role of self-deception in unethical behavior. *Social Justice Research,* 17 (2), 223-236.

Tetlock, P.E. (1992). The impact of accountability on judgment and choice: Toward a social contingency model. In M. Zanna (Ed.), *Advances in Experimental Social Psychology* (pp. 331-376). New York, NY: Academic Press.

Thorson, J. & Biederman-Thorson, M (1974). Distributed relaxation processes in sensory adaptation. *Science* 183, 161–72.

Trevino, L. K., Weaver, R., & Reynolds, S. J. (2006). Behavioral Ethics in Organizations: A Review, *Journal of Management*, 32(6), 951 - 990.

Tulving, E., Schacter, D. L., & Stark, H. A. (1982). Priming effects in word-fragment completion are independent of recognition memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 8, 336-342.