

**Ethicists' and Non-Ethicists' Responsiveness to Student Emails:
Relationships among Expressed Normative Attitude, Self-Described Behavior, and
Experimentally Observed Behavior**

Joshua Rust
Department of Philosophy, Unit 8250
Stetson University
421 North Woodland Boulevard
DeLand, FL 32723
jrust at domain- Stetson.edu

Eric Schwitzgebel
Department of Philosophy
University of California at Riverside
Riverside, CA 92521-0201
eschwitz at domain- ucr.edu

January 10, 2011

**Ethicists' and Non-Ethicists' Responsiveness to Student Emails:
Relationships among Expressed Normative Attitude, Self-Described Behavior, and
Experimentally Observed Behavior**

Abstract:

Do professional ethicists behave any morally better than do other professors? Do they show any greater consistency between their norms and their behavior? In response to a survey question, a large majority of professors (83% of ethicists, 83% of non-ethicist philosophers, and 85% of non-philosophers) expressed the view that “not consistently responding to student emails” is morally bad. A similarly large majority of professors (>80% of all groups) claimed to respond to at least 95% of student emails. We sent these professors, and others, three emails designed to look like queries from students: one concerning office hours, one about declaring a major, and a third about a future course of the professor’s drawn from posted schedules of classes. All three emails were tested against spam filters, and we had direct confirmation that almost all target email addresses were actively used. Professors responded to about 60% of the emails. Ethicists’ email response rates were within statistical chance of the other two groups’. Expressed normative view correlated with self-estimated rate of email responsiveness, especially among the ethicists. However, for all groups of professors, measured email responsiveness was virtually unrelated to either expressed normative view or self-estimated email responsiveness.

Keywords: ethics, morality, moral psychology, experimental philosophy, attitude-behavior consistency, social psychology

Word Count (excluding title and abstract): 7,072, plus one figure and five tables

**Ethicists' and Non-Ethicists' Responsiveness to Student Emails:
Relationships among Expressed Normative Attitude, Self-Described Behavior, and
Experimentally Observed Behavior**

1. Introduction.

Many philosophers – though probably not a majority – appear to believe that professional ethicists behave, on average, at least a little morally better than do socially comparable non-ethicists. (For survey results on this question, see Schwitzgebel and Rust 2009.) Until recently, however, no direct empirical research on this issue was available. Real-world moral behavior is somewhat difficult to study, and professional ethicists are a sophisticated and thinly distributed group.

Two studies from our laboratory suggest that ethicists do *not* behave better than do socially comparable non-ethicists. The first study examined the rates at which ethics books were missing from leading academic libraries compared to similar non-ethics books in philosophy. We found that ethics books were actually more likely to be missing (Schwitzgebel 2009). The second study started from the background assumption that voting in public elections is a civic duty and examined public records of voter participation in five U.S. states. We found that ethicists and political philosophers voted at the same rate as did non-ethicist philosophers and professors in departments other than philosophy (although political science professors did vote about 10-15% more often; Schwitzgebel and Rust 2010). These two studies are, of course, insufficient by themselves to be anything more than suggestive about the overall moral behavior of ethicists.

In this article, we present the results of a third study. We examine the rates at which ethicists and two professorial comparison groups respond to emails from students, on the assumption (evidently shared by most professors: see below) that failing to respond to emails from students is morally bad. This study extends our research into a new domain of behavior, one that is directly personal – between a professor and a student – rather than public and impersonal (returning library books, voting in public elections).

We also measured participants' normative attitudes about responding to student emails, and we asked them to estimate the proportion of student emails they responded to. Thus, in contrast with our previous research, the present study allows us to compare directly measured behavior with expressed normative attitudes and self-reported behavior. Of particular interest, perhaps, is the question of whether ethicists show more or less consistency than do other professors between their expressed normative attitudes and their directly measured behavior. At least on the issue of email responsiveness, are ethicists more or less likely than other groups to act on the norms they espouse? Arguably, norm-behavior consistency is both an intellectual and a moral virtue. And arguably it is just the sort of virtue that practical moral reasoning, if effective, should support.

Why should we care whether ethicists behave any morally better, or any more consistently with their espoused norms, than do socially similar non-ethicists? One reason concerns the efficacy of ethics education. If courses in ethics tend to improve moral behavior or at least to increase norm-behavior consistency, one might expect that professional ethicists, who are much exercised in the kind of study at the center of such courses, would tend to show excellent behavior or at least greater norm-behavior consistency. If professional ethicists are no

better in either respect, that creates a *prima facie* (though perhaps resolvable) empirical challenge for those who would advocate ethics instruction for its effects on behavior.

A second reason – and our main reason for conducting this research – is that examining the moral behavior of ethicists, and their norm-behavior consistency, can help shed empirical light on the general question of the relationship between moral reflection, especially philosophical moral reflection, and moral behavior. If ethicists behave no better, and show no greater norm-behavior consistency, than do non-ethicists, that creates a *prima facie* (though again perhaps resolvable) challenge for psychologists and philosophers who regard philosophical moral reflection as, in general, or at least on average, morally improving. This includes the advocates and inheritors of Socratic and Enlightenment ideals that place intellectual styles of moral reflection near the center of moral development and proper moral motivation (e.g., Plato 4th c. BCE/1961; Kant 1785/1998; Kohlberg 1984). Negative results – and our results are, as you will see, negative – would appear to cohere better with approaches that either see intellectual or philosophical moral reflection as largely post-hoc rationalization of norms arrived at by non-intellectual means (e.g., Nietzsche 1887/1998; Posner 1999; Haidt 2001; Knobe and Leiter 2007) or see moral intellectualism as enervating, less motivationally powerful than the spontaneous moral reactions that intellectualism undercuts (e.g., Baier 1985; Williams 1985). On the latter view, one might expect even less moral behavior and less norm-behavior consistency among ethicists than among non-ethicists.

A note about the ethics of our own study: Our study involved deceiving hundreds of philosophers and other professors, since we posed as undergraduates seeking email responses from professors. Some scholars think that any sort of deception in human research is to be avoided (Erikson 1967; Baumrind 1985). We do not share that view: Deception can sometimes

be justified in scientific research when no significant harm is brought to the participant (see also the ethics code of the American Psychological Association). We took care to keep our emails brief so that reading and responding to those emails would consume no more of the recipient's time than would reading and considering an invitation to participate as a research subject in an ordinary psychological study. We also used a coding procedure, involving unique tracking numbers and the division of data among different computers in different locations, that prevented us from learning which individual email recipients responded and did not respond to the email messages that we sent, thus protecting the privacy of respondents and especially non-respondents. All aspects of our research were approved in advance by the Human Research Review Board at U.C. Riverside. You might still believe our deceptions unethical, and you might regard the above as merely post-hoc rationalization on our part. That would of course fit nicely with the type of pessimistic view about intellectual moral reflection that our research appears to support.

2. The First Email.

Recipients. Our first email was sent to 937 recipients: 334 ethicists, 317 non-ethicist philosophers, and 286 professors in departments other than philosophy ("non-philosophers"). Recipients were drawn from tenure-track faculty at university departments in the same five U.S. states in which we had conducted our voting-rate research: California, Florida, Minnesota, North Carolina, and Washington State. (The original pool was identical to the voting rate used pool in Schwitzgebel and Rust 2010; over time we adjusted some of the information, for example by updating email addresses, adding some newly hired professors, and excluding professors for whom we did not have good contact information.) We focused on these recipients because we

also planned to collect survey data from them on their normative attitudes on voting, responsiveness to undergraduate emails, and several other issues, which we could then compare to their self-reported behavior and their actual, directly measured behavior. (See Section 6 below and Schwitzgebel and Rust in preparation for further details about our survey.) Philosophers in the target departments were classified as “ethicists” if any of the following terms or their cognates appeared in the area of specialization information on their homepage: “ethics”, “moral”, “political”, “law”, “policy”, “race”, “feminism”, “women”, and “justice”. Philosophers who did not list areas of specialization were excluded from analysis; so also were philosophers listing “action” or “religion” among their specializations but no other ethics-related terms. The remaining philosophers were classified as non-ethicists. By this relatively broad measure of “ethicist”, about half of all philosophers in the target departments qualified as ethicists. We conducted a separate analysis of the data using a narrower classification of “ethicist” (“ethics”, “moral”, or “political” among the first two listed specializations) and the results were essentially the same. Non-philosophers were sampled proportionately from faculty directories at the same universities. Professors with no valid listed email address (either unlisted or returned to us undeliverable) are excluded from the numbers above. This included 20 ethicists, 45 non-ethicist philosophers, and 72 non-philosophers. We also noted the gender of recipients (using voter registration data when available or, when feasible, our best guess from the faculty member’s name and picture), and we noted whether recipients were at a “research-oriented” university (having a Ph.D. program in philosophy, plus CalTech) or instead a “teaching-oriented” university.

Message. The first email was sent in April-May 2008. Emails were sent in small batches and were checked against spam filters at Yahoo, Google, U.C. Riverside, and Stetson University.

The sender was "J.R. <hi5university@yahoo.com>" and the subject line was "Office hours?"

The body of the email read as follows:

Dear Prof. [last name]:

Could you please let me know your office hours this term?

Thanks!

Results. We treated as a reply any response other than an automated reply. The main results are presented in Table 1. Although we found a small trend toward higher response rate among ethicists, that trend did not approach statistical significance ($\chi^2 = 1.3$, $p = .51$). Professors at research and teaching institutions also did not detectably differ in response rate (59.8% vs. 55.8%, $\chi^2 = 1.5$, $p = .22$). There was a marginally significant tendency for men to respond at a higher rate than women (59.2% vs. 52.7%, $\chi^2 = 3.0$, $p = .08$); however this trend was not confirmed by subsequent data (presented in Sections 3 and 4 below).

TABLE 1: Response rates to Email 1

	Total number of emails sent	Total number of replies	Response rate
Ethicists	334	197	59.0%
Non-ethicist philosophers	317	184	58.0%
Non-philosophers	286	156	54.6%

Some professors replied by giving their office hours; others explained why they did not have office hours (e.g., on leave); still others responded with an inquiry about who was asking. Among respondents who stated their office hours, all three groups claimed approximately the

same number of office hours per week: 3.1 hours/week for ethicists vs. 3.3 for non-ethicist philosophers and 3.4 for non-philosophers (ANOVA, $F = 0.65$, $p = .52$). (As one might expect, professors at teaching-oriented institutions claimed more office hours than those at research institutions: t-test, mean 3.8 hours vs. 2.4 hours, $t = 8.81$, $p < .001$.) These data are of secondary interest to our hypothesis, to the extent professors can choose how many office hours to hold per week and choosing to hold more office hours might reflect a morally praiseworthy higher level of availability to students.

Discussion. Although we hoped that the message would be interpreted as from a student, and many respondents, to judge from their replies, did so interpret it, discussion with some recipients of this message suggests that not all were convinced: The lack of a name other than “J.R.” and the “spammy”-sounding email address “hi5university@yahoo.com” understandably raised some suspicions. Also, in conversation, some recipients who subsequently learned about our study suggested that there is no obligation to reply to emails about office hours, since students in one’s courses can find one’s office hours on the syllabus and people who are not students in one’s course have no right of access to one’s office hours; consequently, they suggested, declining to reply does not reflect any less well upon recipients than does replying. We are not entirely persuaded by this reasoning, but we attempt to address these concerns in the second and especially third email messages. One of the reasons we chose the topic and vague sender information we did was to avoid explicit deception: The person who sent the email was indeed “J.R.”, and simply requesting office hours is not tantamount to any statement about the identity of the person doing the requesting. In the second and third emails we resorted to more straightforward deception.

3. The Second Email.

Message. The second email message was sent almost one year later, in March and April 2009, to a mostly overlapping group of recipients (224 ethicists, 227 non-ethicist philosophers, and 218 non-philosophers, excluding bounce-backs). The sender was “Ryan Harrison” (<ryharrison89@gmail.com> or <raharrison89@gmail.com>), the subject line was “declaring a major”, and the body of the text was:

Dear Prof. [last name]:

I'm thinking about declaring a major in [major]. Do you know who the department's undergrad advisor is?

Thanks so much!

Ryan Harrison

The “major” field was completed with the name of the department with which the professor was affiliated (philosophy or otherwise), which we hoped would give the email a less “spammy” feel. Again, any non-automated response, regardless of content, was coded as a reply.

Table 2 displays the main results. Again, there is no statistically detectable difference between the groups ($\chi^2 = 1.0$, $p = .60$).

TABLE 2: Response rates to Email 2

	Total number of emails sent	Total number of replies	Response rate
Ethicists	224	120	53.6%
Non-ethicist philosophers	227	113	49.8%

Non-philosophers	218	118	54.1%
------------------	-----	-----	-------

Neither institution type nor gender were predictive (research vs. teaching, 51.5% vs. 53.2%, $\chi^2 = .27$, $p = .60$; male vs. female 53.0% vs. 53.6%, $\chi^2 = .02$, $p = .90$). Overall, recipients were marginally less likely to respond to Email 2 than to Email 1 (52.5% vs. 57.3%, $\chi^2 = 3.71$, $p = .054$).

Results. The results for Email 2 appear to confirm our findings on Email 1: The groups did not differ, overall, in their responsiveness. However, in discussion a number of recipients indicated that they felt no obligation to respond since information about the major advisor is widely available from other sources. Perhaps this explains the slightly lower response rate, despite the somewhat less spammy sender and content. Again, we the authors are not entirely convinced that it isn't somewhat morally better – somewhat kinder or more generous of one's time – to respond to emails of this sort than to ignore them; but in light of such concerns we decided to create a more personal-seeming third email. Also, we decided to expand the pool of recipients for the third email, giving us more power to detect small differences. And finally, since part of the explanation for the mediocre response rates to the first two emails may have been that some were sent to professors no longer actively involved in teaching at the university (though still with valid email addresses), we collected information about which professors were scheduled to teach undergraduate courses in the upcoming term.

4. The Third Email.

Recipients. We examined course catalog information for professors at the target universities in the original five-state pool and divided recipients into two groups: Group 3A included only recipients scheduled to teach an undergraduate level course in the upcoming term

(Fall 2009); Group 3B were recipients either not scheduled to teach an undergraduate course in the upcoming term or for whom teaching information was unavailable. From the original five-state pool, 392 recipients fell into group 3A and 517 fell into group 3B. We then expanded Group 3A by adding 601 recipients from other states, selected and sorted in the same manner as for Email 1. A similar proportion of ethicists, non-ethicist philosophers, and non-philosophers belonged to Groups 3A and 3B (approximately two-thirds in Group A for all three professor types: 67.8% of ethicists vs. 68.7% of non-ethicist philosophers vs. 64.6% of non-philosophers, $\chi^2 = 2.30$, $p = .32$). Due to a programming error, some professors received an email containing incorrect name information; these professors are excluded from the above numbers.

Message. The third email was sent several months after Email 2, in July-August, 2009, always before the beginning of the recipient's school term. The sender was "Kati Sanchez" and the email address was "katisanchez11@gmail.com" or a similar address with a different two-digit number.

Recipients in group 3A – those scheduled to teach in the fall term – received a different email from those in group 3B. For recipients in group 3A the subject line was "question regarding your fall course" and the text of the email was:

Dear Professor [last name] –

I was planning to take your fall course, [class]. However, due to some unavoidable family obligations, I will not be able to attend the first two or possibly three class meetings. But the course sounds interesting, I'm a good student with a 3.6 GPA, and I'm willing to put in a little extra work to get up to

speed. Do you think it would possible to succeed in your class, despite those missed days? Or should I register for a different class?

Thanks for your time responding to this!

Kati Sanchez

The [class] field was completed with the name of an undergraduate course the professor was scheduled to teach in the coming term. In general, we filled [class] with the lowest level class the professor was scheduled to teach, avoiding graduate level courses.

For group 3B the subject line was “Teaching next term?” and the text was

Dear Professor [last name]:

Will you be teaching any lower division classes next term or in the next year? A friend of mine took a class of yours a while back and highly recommended you as a teacher, but I didn't see any introductory level classes of yours in next term's course schedule. I'm a new transfer student into [university], starting in the fall.

Sincerely,

Kati Sanchez

Results. Table 3 displays the main results. As with Email 1, but not Email 2, ethicists showed a statistically non-significant trend toward higher responses rates than non-ethicists (Version 3A: $\chi^2 = 3.50$, $p = .17$; Version 3B: $\chi^2 = 1.54$, $p = .46$; combined: $\chi^2 = 2.56$, $p = .29$). In

Section 5, we will consider whether the aggregate results of all three emails suggest a higher response rate for ethicists.

TABLE 3: Response rates to Email 3

		Total number of emails sent	Total number of replies	Response rate
Version 3A	Ethicists	307	212	69.1%
	Non-ethicist philosophers	363	241	66.4%
	Non-philosophers	316	196	62.0%
Version 3B	Ethicists	156	101	64.7%
	Non-ethicist philosophers	176	103	58.5%
	Non-philosophers	185	110	59.5%

Neither institution type nor gender were predictive (research vs. teaching, 63.7% vs. 66.5%, $\chi^2 = 1.28$, $p = .26$; male vs. female 65.0% vs. 63.1%, $\chi^2 = .25$, $p = .62$). Recipients were a bit more likely to respond to Version 3A than to Version 3B: 66.4% vs. 60.7% overall ($\chi^2 = 4.78$, $p = .03$). Response rates were significantly higher to Email 3A than to Emails 1 and 2, and response rates to Email 3B were significantly higher than to Email 2 but not to Email 1 (χ^2 , all p 's < .01 except Email 3B vs. Email 1, $\chi^2 = 1.61$, $p = .21$). The somewhat different response rates for the various email messages might reflect differences in the plausibility of the messages and the extent to which they seem to call for reply, and for Email 3A they might also reflect a difference in the recipient pool (since it was limited to professors scheduled to teach in the

upcoming term). Despite these differences in message and recipient pool, however, variability in response rates between versions is only moderate, from a minimum of 52.5% for Email 2 to a maximum of 66.4% for Email 3A.

5. Combined Measures.

The simplest combined measure, but least statistically conservative (because it treats all trials as statistically independent), compares the total percentage of replies to all emails from the three groups. This measure yields total response rates of 61.7% for the ethicists, 59.2% for the non-ethicist philosophers, and 57.7% for the non-philosophers – a trend that does not reach statistical significance despite 3,109 total trials ($\chi^2 = 3.44$, $p = .18$). An alternative statistical approach examines the mean response among those professors who received all three emails (182 ethicists, 178 non-ethicist philosophers, and 165 non-philosophers). The mean number of responses was 1.78 for the ethicists compared to 1.67 for the non-ethicist philosophers and 1.78 for the non-philosophers – not detectably different (ANOVA, $F = 0.81$, $p = .45$). We conclude that ethicists are not detectably more likely to reply to email messages of the sort we sent, though we cannot rule out a small, statistically undetected difference of up to about 5%.

One potential source of concern is the possibility that some of our emails were going to inert email addresses – technically valid addresses (so that we received no bounce-back message) but de facto unchecked by the respondents. The combined data of the three email studies, plus the survey study to be described in Section 6 below, speak against this interpretation of the mediocre response rates to our email messages. Among the 525 professors who received all three of our emails, 488 (93.0%) either responded to one of our email messages or took the

electronic version of our survey by following a link sent to them by email. We therefore conclude that almost all our email addresses were actively checked by our email recipients.

As mentioned in Sections 2 and 3 above, recipients of the first two email messages, when they learned about our study, sometimes asserted that such messages do not call for response. Few professors said this about Email 3, and we hope that most readers who agree that professors have a general obligation to respond to most emails from students will agree that Email 3A is among those that call for reply. If the mediocre response rates to Emails 1 and 2 were primarily due to the fact that those messages did not call for reply, then we should expect that many professors who did not reply to either of those two messages would still have replied to Email 3. However, among the 262 professors who received Email 1, Email 2, and Email 3A, only 35 (13.4%) showed this pattern of response – little different from the 12.5% one would expect were professors evenly distributed among the eight possible reply patterns. We agree with those who would say that Email 3A compels response more than do Emails 1 and 2, but it does not appear that professors' response rates show a particularly high level of sensitivity to such differences in message content.

One might think that many professors would be either consistent responders or consistent non-responders. We did not find this pattern in the data. Among the 525 professors who received all three emails, only 52 (9.9%) replied to none of them and only 115 (21.9%) replied to all three. Most professors showed intermediate patterns of responsiveness, with 145 (27.6%) replying to one message and 214 (40.6%) replying to two messages. Ethicists were similarly intermediate in their responses (9.9%, 25.3%, 41.8%, and 23.1% for 0 to 3 responses respectively; $3 \times 4 \chi^2 = 2.07$, $p = .91$). Professors who replied to Email 1 were more likely to reply to Email 2 (56.5% vs. 45.5%, $\chi^2 = 7.09$, $p = .008$), and professors who replied to Email 2

were more likely to reply to Email 3 (68.5% vs. 59.0%, $\chi^2 = 5.71$, $p = .02$). However, we would describe those effect sizes as moderate. If our messages are representative, most professors, including most ethicists, show a pattern of inconsistent, mediocre response to emails from students.

6. Self-Reported Normative Attitude and Self-Reported Email Responsiveness.

Survey questions. Most of the professors in the original five-state voting data pool (980 total: 337 ethicists, 329 non-ethicist philosophers, 314 non-philosophers) also received a survey questionnaire concerning “professors’ moral attitudes and behavior”. The survey methodology and results are presented in detail in Schwitzgebel and Rust (in preparation). Among the 25-28 survey questions, two concerned email responsiveness. Part One of the survey began as follows:

Please indicate the degree to which the action described is morally good or morally bad by checking one circle on each scale.

We recognize that it may be difficult to rate moral goodness and badness on a numerical scale, that different moral goods may be incommensurable, and that the goodness or badness of an action can vary with context. We encourage you to set aside such concerns as best you are able, interpreting the questions below as straightforwardly as possible. You are also welcome to clarify your answers, raise objections to the wording of the questions, etc., in the margins.

There followed nine prompts concerning theft, professional society membership, blood and organ donation, vegetarianism, voting in public elections, not staying in regular contact with one’s mother, and charitable donation, with each prompt followed by a nine-point scale with

“very morally bad” at one endpoint (coded as 1) and “very morally good” at the other endpoint (coded as 9). The 3, 5, and 7 points were labeled as “somewhat morally bad”, “morally neutral”, and “somewhat morally good” respectively. The eighth prompt was “not consistently responding to student emails”.

Part Two of the survey asked respondents to self-report their behavior on the same issues covered in Part One (except theft). Question 22 of the survey asked: “About what percentage of student emails do you respond to?” This question was followed by “enter a percentage” and a blank field.

The survey was sent in February to March 2009 (thus between Email 1 and Email 2). Recipients received up to five communications, four by email, one by traditional post, until they either completed the survey or opted out of further communications. Overall response rates to the survey were very good by social science standards, with non-philosophers moderately less likely to respond (perhaps unsurprisingly, given the philosophical nature of the survey): ethicists’ response rate 58.8%, non-ethicist philosophers’ 63.2%, non-philosophers’ 53.2% ($\chi^2 = 6.68$, $p = .04$). Professors who responded to our “student” email communications were more likely to complete the survey (also unsurprisingly, given that four of the five survey communications were by email), though the difference was again only moderate: 62.0% of those who replied to Email 2 (the email temporally closest to the time of the survey) completed the survey, compared to 53.3% of those who did not reply to Email 2 ($\chi^2 = 4.96$, $p = .03$).

Expressed normative attitude. The large majority of respondents rated “not consistently responding to student emails” on the morally bad side of the scale (i.e., 1-4 on the 9-point scale). The size of majority was similar among the three groups: 82.6% of ethicists rated it as morally bad, compared to 83.4% of non-ethicist philosophers and 84.9% of non-philosophers ($\chi^2 = 0.37$,

$p = .83$). The mean response to the normative question was 3.29, near the “somewhat morally bad” label of the scale. Non-philosophers rated not responding a little morally worse on average than did philosophers (3.05 vs. 3.36 for ethicists and 3.42 for non-ethicist philosophers; ANOVA, $F = 5.19$, $p = .006$), but we are unsure whether this difference was due to scaling differences among the groups (philosophers possibly reserving the endpoints for more extreme moral and immoral actions than non-philosophers), non-response bias (since a smaller percentage of non-philosophers than philosophers responded), or a genuine difference in group opinion. (For further discussion of these issues see Schwitzgebel and Rust in preparation.) Only 7 respondents out of 566 (1.2%) rated not consistently responding to student emails on the morally good side of the scale. The remainder – 15.2% of respondents – rated it as morally neutral.

Self-reported behavior. All groups of respondents self-reported very high rates of email responsiveness. 50.5% of respondents claimed to reply to 100% of student emails, 67.4% claimed to reply to at least 98% of student emails, and 83.6% of respondents claimed to respond to at least 95% of student emails. Even independently of the response rates to our three email messages, which may be unrepresentative as measures of overall response rate to student emails, such high numbers seem implausible. In informal polling of student audiences, we have found few students who say that such estimates match their own experience in attempting to communicate with professors. Although survey non-response bias may explain a part of these results – as described above, there was a moderate positive relationship between survey responsiveness and responsiveness to our emails – that modest relationship cannot go very far in explaining such extreme responses. As displayed in Table 4, the three groups differed little in their self-reported email responsiveness ($\chi^2 = 2.01$, $p = .37$; $\chi^2 = 1.73$, $p = .42$).

TABLE 4: Self-reported rates of responsiveness to student emails

	Percentage claiming 100% responsiveness	Percentage claiming at least 95% responsiveness
Ethicists	49.5%	81.5%
Non-ethicist philosophers	47.7%	83.1%
Non-philosophers	55.1%	86.7%

Relationship of expressed normative attitude and self-reported behavior. Among professors who rated “not consistently responding to student emails” on the morally bad side of the scale, 67.5% claimed to reply to at least 98% of student emails. In contrast, among those who did not rate “not consistently responding to student emails” on the morally bad side of the scale, only 43.0% of professors self-reporting at least 98% responsiveness ($\chi^2 = 27.72, p < .001$). As displayed in Table 5, ethicists showed much higher levels of consistency between their expressed normative attitudes and their self-described behavior than did the two other groups (permutation test, reshuffling group labels: 0.6% of 10,000 samples met or exceeded ethicists’ 49.3% differential). If the behavioral self-reports are to be trusted, ethicists show much higher norm-behavior consistency than the other groups.

TABLE 5: Relationship of expressed normative attitude toward “not consistently responding to student emails” and self-described behavior.

	Percentage reporting at least 98% responsiveness to
--	---

	student emails		
	If said bad not to respond	If said not bad not to respond	<i>differential</i>
Ethicists	74.2%	25.9%	49.3%
Non-ethicist philosophers	67.9%	46.7%	21.2%
Non-philosophers	73.7%	60.0%	13.7%

Relationship of self-reported behavior and actual behavior. However, the behavioral self-reports are not to be trusted. Self-reported email responsiveness was only modestly related to measured email responsiveness. Only for Email 3 was there a statistically detectable difference between the reply rates of professors who self-reported responding to at least 98% of student emails and professors who self-reported responding at lower rates, and the difference was only moderate: 68.3% vs. 58.3% ($\chi^2 = 4.79, p = .03$). (For the other emails the rates were 63.5% vs. 60.1% and 60.3% vs. 52.2% respectively, $\chi^2 = 0.53, 2.07, p = .49, .15$.) Aggregating across the emails, the professors who claimed to reply to at least 98% of student email replied to an average of 64.3% of the emails we sent, compared to a 57.4% average response rate for professors who claimed to reply to fewer than 98% of student emails ($t = 2.26, p = .02$). The only subgroup showing a statistically significant relationship between self-described behavior and measured behavior was the non-philosophers: Among the non-philosophers, the self-described 98%+ responders replied to an average of 65.8% of the emails we sent, compared to an average 53.0% response rate among those not claiming 98%+ responsiveness, a 12.8% differential ($t = 2.58, p = .01$). For ethicists the corresponding percentages were 66.7% vs.

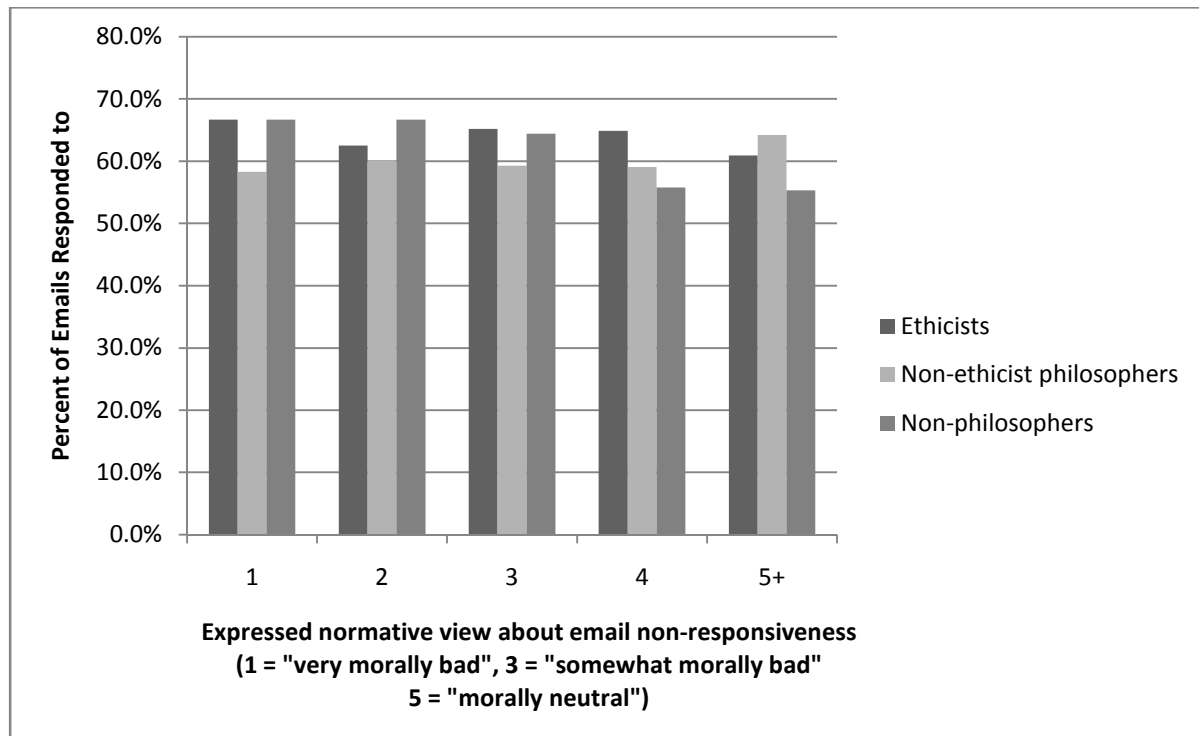
58.7% (8.0% differential, $t = 1.46$, $p = .15$); for non-ethicist philosophers 60.6% vs. 59.2% (1.4% differential, $t = 0.28$, $p = .78$).

Here is another way to capture the modest relationship between self-reported email responsiveness and responsiveness to our three email messages: The correlation between individuals' self-reported percentage responsiveness and their observed percentage responsiveness was $r = .16$ ($p < .001$) – generally considered a “low” correlation in social science research. For ethicists, the correlation was so low as not to be statistically significant with this sample size ($r = .12$, $p = .11$). Among recipients as a whole, self-reported email responsiveness predicted 3% of the variance in measured responsiveness; among ethicist recipients, self-reported responsiveness predicted 1% of the variance.

Relationship of expressed normative attitude and measured behavior. Using a suite of χ^2 and t-tests, we found no statistically significant relationships between expressed normative attitude and measured behavior, whether we analyzed the emails individually or whether we looked at professors' individual percent responsiveness, whether we considered attitude as a dichotomous variable (attempting the split in several ways) or whether we treated it as a continuous variable (looking at overall rating on the 1-9 scale), and whether we considered all professors as a group or broke the responses down by subgroup. Despite the substantial relationship between expressed normative attitude and *self-reported* behavior among ethicists, the relationship between expressed normative attitude and measured behavior was statistically insignificant for all three groups. For example, the average reply rate to our emails for professors who rated not consistently responding to student emails on the bad side of the scale was 62.5%; for those who rated it as neutral (or good) the response rate was 60.6%, only a 2.1% differential, with a 95% confidence interval of -5.2% to 9.0% ($t = .53$, $p = .60$; for ethicists

64.8% vs. 60.8%; $t = 0.62$, $p = .54$, CI -9.0%-17.0%). Figure 1 displays another way of looking at the data.

FIGURE 1: Relationship between expressed normative view about “not consistently responding to student emails” and responsiveness to Emails 1-3, by group



Given the size of the confidence intervals, we cannot rule out the possibility of undetected moderate differences: For the groups combined, there may have been up to a 9% undetected population responsiveness differential between those rating non-responsiveness on the bad side of the scale vs. those not doing so. Indeed, some at least weak relationship between attitude and behavior seems both a priori likely and to be supported by the direction of the trends. But the absence of a statistically detectable effect is nonetheless striking. In social psychology, correlations are a standard measure of attitude-behavior consistency. In an influential meta-

analysis, Kraus (1995) reports an average attitude-behavior correlation across studies of $r = .32$. Had a correlation of that size been present in our data, we would have had ample power to detect it. In contrast, the correlation we find between professors' measured email reply rates and their expressed normative attitude on the nine-point scale is only $r = .05$ ($p = .24$; converting negative to positive to indicate consistency; for ethicists $r = .07$, $p = .36$.). Less than 1% of the variance in measured reply rate is explained by our respondents' expressed normative attitudes.

7. Conclusion.

Overall, we found no statistically significant difference between ethicists' and non-ethicists' responsiveness to three emails that were designed to look as though they were sent by students. On the assumption that it is morally better for professors to respond to student emails, and on the assumption that at least one of our three email messages was such that it was morally better to respond to it than to ignore it – e.g., kinder, more generous, more dutiful – this finding appears to fit with our previous research suggesting that ethicists, on average, behave no morally better than do other professors (Schwitzgebel 2009; Schwitzgebel and Rust 2010). In support of this working assumption we note that the large majority – about 83% – of survey respondents from all three groups of professors (ethicists, non-ethicist philosophers, and non-philosophers) rated “not consistently reporting to student emails” on the morally bad side of our nine-point scale.

We recognize, of course, the limited scope of the behaviors so far studied: email responsiveness, returning library books, and voting in public elections. We now have research in progress exploring further issues like vegetarianism and charitable donation (Schwitzgebel and Rust in preparation). We also recognize that any observational study, like the present one, that

does not involve random assignment into experimental groups risks being tangled with confounding factors. For example, perhaps ethicists are more likely to teach large classes than are other professors, and professors teaching large classes respond to a lower percentage of student emails because they are inundated and overwhelmed. However, we see no reason a priori to think such confounding factors would overall tend to disadvantage ethicists and thus mask an underlying disposition among ethicists to be more responsive to student emails. (For example, in our experience, Introduction to Philosophy, typically taught by non-ethicists, tends to enroll at least as many students as do introductory ethics classes.) Furthermore, as evidence of the moral non-superiority of ethicists accumulates across a variety of different measures, it becomes increasingly implausible to suppose that each non-effect is explained by undiscovered confounds that mask ethicists' real moral superiority.

Ethicists, non-ethicist philosophers, and professors in departments other than philosophy all reported very high rates of responsiveness to student emails, with a majority of respondents reporting that they respond to 100% of student emails and 84% of respondents reporting that they respond to at least 95% of student emails. For various reasons, such response percentages would seem, antecedently, to be unrealistic. The actual measured responsiveness to our three email messages was about 60% – much lower than 100% or even 95%. It could plausibly be argued that response rates to our three messages substantially underestimate the actual email responsiveness of professors, since all three of our messages came from unknown senders. On the other hand, however, it could be argued that response rates to our three messages, especially the first two, could be expected to overestimate responsiveness, since they admit of very quick reply and thus are less likely to be set aside for later reply and subsequently forgotten. We also emphasize that professors, especially those teaching large courses, do not always know the

names of all their students, much less the names of all their potential and upcoming students, so if they aim to adhere to the norm of consistently replying to student emails they should be willing to examine emails with plausible subject lines like “office hours?”, “declaring a major”, or “question regarding your fall course” even if the name is unfamiliar.

We see no particular reason to regard respondents’ very high estimates as deliberately deceptive. The more plausible interpretation, we think, is that most professors lack self-knowledge about their rates of email responsiveness. Neglected emails, we suspect, are typically forgotten and because forgotten unlikely to figure in one’s estimates of responsiveness. Generally speaking, people rarely notice their thoughtless rudenesses unless someone is bold enough to call them out. This epistemic failing is perhaps also a moral failing: If professors have an obligation to respond to emails from students, then arguably they also have a further obligation to track whether or not they are meeting the first obligation, so that if they are not meeting the first obligation they can take corrective measures. If this is correct, then the present study offers not just one measure of morality, email responsiveness, but two: email responsiveness and meeting one’s moral obligation not to be deluded about one’s level of email responsiveness. Professors remain far short of ideal by either measure, ethicists no less so than the others.

Although normative attitude as expressed by response on our scale was substantially related to *self-reported* rates of email responsiveness, especially among ethicists, normative attitude was virtually unrelated to objectively measured email responsiveness. For all three groups of professors, response to our normative question about the morality of responding to student emails was almost entirely unpredictable of actual response to our three email messages. We were surprised by this result. Psychologists have tended to find that expressed attitude is

typically at least moderately predictive of measured behavior – and when it is not, it is often either because the measure of behavior is one-shot and not closely connected to the measured attitude or in cases of implicit bias (see Kraus 1995). Such factors do not appear to explain the non-relationship in the present case.

We think that even cynics about the value of philosophical moral reflection ought to expect a certain amount of norm-behavior consistency: That's what post-hoc *rationalization* is all about, right? On an optimistic view of the norm-behavior relationship, people rationally endorse norms and shape their behavior to fit them; on a pessimistic view, people know their behavior and then rationalize their way into self-congratulatory norms. In either case, norms and behavior should be at least somewhat aligned. Professors either think about whether responding to student emails is important and shape their behavior accordingly, or they notice their patterns of responsiveness or non-responsiveness to student emails and adopt a matching normative view. In either case, responders should tend to condemn the non-responders as shirking their duties, while non-responders should tend to regard themselves as under no obligation to respond. Our data, however, suggest that this is not the case.

One possibility is that either our measure of normative attitude or our measure of actual behavior is poor. We acknowledge this possibility, yet we also note that psychologists tend to find moderate attitude-behavior consistency even with relatively crude measures on both sides, measures fairly similar in complexion to our own. A more interesting possible explanation of this non-finding involves combining the post-hoc rationalization model of norm-behavior consistency with our suggestion above that professors have little self-knowledge of their actual rates of responsiveness to student emails. Post-hoc rationalization will not succeed in aligning norms and behavior unless the rationalizer knows what her behavior is. Otherwise, it will align

norms with false opinions about one's behavior. And indeed we do see alignment between expressed normative attitude and self-reported behavior, especially for the ethics professors. In the present case then, perhaps, professors' normative views about email responsiveness are to a substantial extent rationalized post-hoc to fit their inaccurate and radically optimistic self-assessments of their own responsiveness. In other words, they have shaped their norms to match their self-flattering illusions. If this is indeed the case, it seems to be especially so for the ethicists: They showed by far the strongest relationship between expressed normative view and self-reported behavior, and (unlike non-philosophers) they showed only a weak, statistically insignificant relationship between self-reported behavior and actual behavior. This interpretation of our data would also harmonize with other work suggesting that the order in which individual moral scenarios are presented has a greater effect on ethicists' than on other professors' judgments about related abstract moral principles – a finding that suggests that ethicists are more likely than other professors to engage in post-hoc rationalization of experimentally-manipulated scenario judgments (Schwitzgebel and Cushman in preparation). Expertise in ethics might be, to a substantial extent, expertise in post-hoc rationalization of opinions arrived at by largely unwelcome psychological mechanisms – a view that Nietzsche would no doubt have endorsed (Nietzsche 1887/1998; Knobe and Leiter 2007).

It is far too soon, we think, to say that such a starkly negative view of philosophical ethics is compelled by the psychological data. In fact, we the authors would reject such a view in any strong form. Yet it would nicely explain our results so far.

References:

- Baier, Annette (1985). *Postures of the mind*. Minneapolis: University of Minnesota.
- Baumrind, Diana (1985). Research using intentional deception. *American Psychologist*, 40, 165-174.
- Erikson, Kai T. (1967). A comment on disguised observation in sociology. *Social Problems*, 14, 366-373.
- Haidt, Jonathan (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814-834.
- Kant, Immanuel (1785/1998). *Groundwork of the metaphysics of morals*. Trans. M. Gregor. Cambridge: Cambridge.
- Knobe, Joshua, and Brian Leiter (2007). The case for Nietzschean moral psychology. In B. Leiter and N. Sinhababu, eds., *Nietzsche and morality*. Oxford: Oxford.
- Kohlberg, Lawrence (1984). *The psychology of moral development*. Cambridge, MA: Harper & Row.
- Kraus, Stephen J. (1995). Attitudes and the prediction of behavior: A meta-analysis of the empirical literature. *Personality and Social Psychology Bulletin*, 21, 58-75.
- Nietzsche, Friedrich (1887/1998). *On the genealogy of morality*. Trans. M. Clark and A.J. Swenson. Indianapolis: Hackett.
- Plato (4th c. BCE/1961). *The collected dialogues*. Ed. E. Hamilton and H. Cairns. Princeton: Princeton.
- Posner, Richard A. (1999). *The problematics of moral and legal theory*. Cambridge, MA: Harvard.

Schwitzgebel, Eric (2009). Do ethicists steal more books? *Philosophical Psychology*, 22, 711-725.

Schwitzgebel, Eric, and Fiery Cushman (in preparation). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers.

Schwitzgebel, Eric, and Joshua Rust (2009). The moral behavior of ethicists: Peer opinion. *Mind*, 118, 1043-1059.

Schwitzgebel, Eric, and Joshua Rust (2010). Do ethicists and political philosophers vote more often than other professors? *Review of Philosophy and Psychology*, 1, 189-199.

Schwitzgebel, Eric, and Joshua Rust (in preparation). The self-reported moral behavior of ethics professors.

Williams, Bernard (1985). *Ethics and the limits of philosophy*. Cambridge, MA: Harvard.