

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

Joshua Rust

Eric Schwitzgebel

Abstract:

Do professional ethicists behave any morally better than do other professors? Do they show any greater consistency between their normative attitudes 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 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. Ethicists’ email response rates were within chance of the other two groups’. Expressed normative view correlated with self-estimated rate of email responsiveness, especially among the ethicists. However, empirically measured email responsiveness was at best weakly correlated with self-estimated email responsiveness; and professors’ expressed normative attitude was not significantly correlated with empirically measured email responsiveness for any of the three groups.

Keywords:

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

## *1. Introduction.*

Approximately half of American ethicists 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.

A few studies from our laboratory suggest that ethicists do *not* behave better than do socially comparable non-ethicists. In one study, we 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). Another 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 than other professors; Schwitzgebel and Rust 2010). In another study we found ethicists to behave no more courteously than non-ethicists at conferences of the American Philosophical Association, as measured by the number of times audience members talked over a speaker or commentator, door slamming during a session, and amount of trash and cups left behind (Schwitzgebel et al. 2012). Still another study suggests that ethicist participants

at American Philosophical Association conferences are just as likely to avoid paying conference registration fees as are non-ethicists (Schwitzgebel forthcoming).

In this article, 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 personal – between a professor and a student – rather than public and impersonal.

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 regarding email responsiveness, are ethicists more or less likely than other groups to act on the norms they espouse? Arguably, attitude-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 attitude-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 attitude-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 and attitude-behavior consistency of ethicists can help shed empirical light on the psychological determinants of moral behavior. On what one might think of as the “booster” view of philosophical moral reflection, habits of philosophical moral reflection will tend to improve overall moral behavior. Boosterism of this sort has been historically influential in philosophy, and one can see various qualified versions of it in, for example, Plato (4th c. BCE/1961), Aristotle (4th c. BCE/1962), Kant (1785/1998), and Mill (1859/2003); and also in twentieth-century moral psychology, especially in the tradition of Kohlberg (1984). The booster view might be advocated at various levels of specificity, ranging from the most general, according to which abstract philosophical moral reflection improves moral behavior in general, to the very concrete on which philosophical moral reflection about some single, imminent behavioral choice will tend to improve moral behavior on that particular occasion. The possibility that the relationship between moral reflection and moral behavior plays out differently for different issues (e.g., voting vs. personal relationships to students) is part of what recommends the multi-pronged approach we have taken to this issue across our various studies.

Negative results – and our results are, as you will see, negative – create a *prima facie* empirical challenge for boosterism about the behavioral effectiveness of philosophical moral reflection: If philosophical moral reflection is morally beneficial, why don't professional ethicists seem to be improving morally as a result of their presumed expertise in it? While we think that this question admits of several potential responses consistent with boosterish views (e.g., ethicists are not in fact more expert, or ethicists tend to start out with deficient moral

intuitions and use intellectual tools to improve to average, or some forms of philosophical moral reflection are morally improving and some are morally worsening and they cancel out on average), on the face of it, negative results would seem to fit more naturally with various “scoffer” views about the relation between moral reflection, attitudes, and action. A “scoffer” might hold that philosophical moral reflection is primarily post-hoc rationalization of particular antecedently held moral beliefs (e.g., Posner 1999; Haidt 2001, 2012), or, alternatively, that it is primarily post-hoc rationalization of one’s behavioral inclinations (e.g., Nietzsche 1887/1998; Knobe and Leiter 2007), or that philosophical moral reflection is enervating and tends to undercut one’s spontaneous moral inclinations (e.g., Baier 1985; Williams 1985). For the purposes of our study, it’s worth noting that these three scoffer views seem to differ in their predictions about attitude-behavior consistency: The rationalization-to-behavior view might predict a higher correlation between attitude and behavior, or at least to one’s opinions about one’s behavior, than do the other two views. Perhaps, even, the view that intellectualism is enervating might predict that philosophical reflection would tend to *reduce* attitude-behavior correlation. We don’t pretend to be able to sort out the full implications of such models here; no doubt they can be developed in a variety of directions. Our point is only to gesture at the types of issues on which our study, if successful, should begin to cast some light.

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 excuses 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 five U.S. states: California, Florida, Minnesota, North Carolina, and Washington State. 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 academic homepage: "ethics", "moral", "political", "law", "policy", "race", "feminism", "women", and "justice". We excluded from analysis philosophers who did not list areas of specialization, philosophers listing "action" or "religion" among their specializations but no other ethics-related terms, and professors with unlisted or invalid email addresses. Non-philosophers were sampled proportionately from

faculty directories at the same universities. We also noted gender and 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 a higher response rate among ethicists, that trend did not approach statistical significance.<sup>1</sup> Professors at research and teaching institutions also did not detectably differ in response rate (59.8% vs. 55.8%).<sup>2</sup> There was a marginally significant tendency for men to respond at a higher rate than women (59.2% vs. 52.7%).<sup>3</sup> However, this trend was not confirmed by subsequent data (presented in Sections 3 and 4 below).

TABLE 1: Response rates to Email 1.

---

<sup>1</sup>  $\chi^2(3) = 1.3, p = .51.$

<sup>2</sup>  $\chi^2(2) = 1.5, p = .22.$

<sup>3</sup>  $\chi^2(2) = 3.0, p = .08.$

	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.<sup>4</sup> (As one might expect, professors at teaching-oriented institutions claimed more office hours than those at research institutions: mean 3.8 hours vs. 2.4 hours.<sup>5</sup>) 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, the replies from some recipients of this message suggest 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, some recipients who subsequently learned about our study

---

<sup>4</sup> ANOVA,  $F(2, 294) = 0.65, p = .52$ .

<sup>5</sup>  $t(494) = 8.81, p < .001$ .



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 professors than does replying. We are not entirely persuaded by this reasoning but we attempt to address these concerns in the messages of the second and especially third email studies. 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 among the groups.<sup>6</sup>

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 significantly predictive (research vs. teaching, 51.5% vs. 53.2%; male vs. female 53.0% vs. 53.6%).<sup>7</sup> Overall, recipients were marginally less likely to respond to Email 2 than to Email 1 (52.5% vs. 57.3%).<sup>8</sup>

*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

---

<sup>6</sup>  $\chi^2(2) = 1.0, p = .60.$

<sup>7</sup>  $\chi^2(1) = .27, p = .60; , \chi^2(1) = .02, p = .90.$

<sup>8</sup>  $\chi^2(1) = 3.71, p = .054.$

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 in response rate. 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).<sup>9</sup>

---

<sup>9</sup>  $\chi^2(1) = 2.30, p = .32.$

Due to a programming error, some professors received an email containing incorrect name information; these professors have been 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.

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 response rates than non-ethicists.<sup>10</sup> 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%

---

<sup>10</sup> Version 3A:  $\chi^2(2) = 3.50$ ,  $p = .17$ ; Version 3B:  $\chi^2(2) = 1.54$ ,  $p = .46$ ; combined:  $\chi^2(2) = 2.56$ ,  $p = .29$ .

	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.<sup>11</sup> Recipients were a bit more likely to respond to Version 3A than to Version 3B: 66.4% vs. 60.7% overall.<sup>12</sup> 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 than to Email 1.<sup>13</sup> 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

---

<sup>11</sup> Research vs. teaching, 63.7% vs. 66.5%,  $\chi^2(1) = 1.28$ ,  $p = .26$ ; male vs. female, 65.0% vs. 63.1%,  $\chi^2(1) = .25$ ,  $p = .62$ .

<sup>12</sup>  $\chi^2(1) = 4.78$ ,  $p = .03$ .

<sup>13</sup>  $\chi^2$ , all  $p$ 's < .01 except Email 3B vs. Email 1,  $\chi^2(1) = 1.61$ ,  $p = .21$ .

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.<sup>14</sup> 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 – again not detectably different.<sup>15</sup> 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 unmonitored 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

---

<sup>14</sup>  $\chi^2(2) = 3.44, p = .18.$

<sup>15</sup> ANOVA,  $F(2, 522) = 0.81, p = .45.$

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. No recipient has so far expressed this view 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 3A. However, among the 262 professors who received Email 1, Email 2, and Email 3A, only 35 (13.4%) showed that N-N-Y 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 professors' response rates don't appear to show much 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).<sup>16</sup> Professors who replied to Email 1 were more likely to reply to Email 2 (56.5%

---

<sup>16</sup>  $\chi^2(6) = 2.07, p = .91$ .



vs. 45.5%), and professors who replied to Email 2 were more likely to reply to Email 3 (68.5% vs. 59.0%).<sup>17</sup> 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 forthcoming). Among the 25 to 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

---

<sup>17</sup>  $\chi^2(1) = 7.09, p = .008; \chi^2(1) = 5.71, p = .02.$

one's mother, charitable donation, and (as the eighth prompt) "not consistently responding to student emails". Each prompt was 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.

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 (between Email 1 and Email 2). Recipients received up to five communications, four by email and 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%).<sup>18</sup> 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.<sup>19</sup>

---

<sup>18</sup>  $\chi^2(2) = 6.68, p = .04.$

<sup>19</sup>  $\chi^2(1) = 4.96, p = .03.$

*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 the 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.<sup>20</sup> 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<sup>21</sup>), 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 forthcoming.) Seven 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 advanced students who say that such estimates match their own experience in attempting to

---

<sup>20</sup>  $\chi^2(2) = 0.37, p = .83.$

<sup>21</sup> ANOVA,  $F(2, 563) = 5.19, p = .006.$

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. The trend toward higher self-reported response rates among non-philosophers did not approach statistical significance.<sup>22</sup>

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-reported at least 98% responsiveness.<sup>23</sup> As displayed in

<sup>22</sup>  $\chi^2(2) = 2.01, p = .37; \chi^2(2) = 1.73, p = .42.$

<sup>23</sup>  $\chi^2(1) = 27.72, p < .001.$

Table 5, ethicists showed a much stronger relationship between their expressed normative attitudes and their self-described behavior than did the two other groups.<sup>24</sup> If the behavioral self-reports are to be trusted, ethicists show much higher attitude-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	differential
Ethicists	74.2%	25.9%	49.3%
Non-ethicist philosophers	67.9%	46.7%	21.2%
Non-philosophers	73.7%	60.0%	13.7%

---

<sup>24</sup> Permutation test, reshuffling group labels: 0.6% of 10,000 samples met or exceeded ethicists’ 49.3% differential; overall correlation between expressed normative attitude on the 1-9 scale and self-reported response rate (-log(101-x)-transformed) was .38 for ethicists, .24 for non-ethicist philosophers, and .13 for non-philosophers; a significant between-groups difference using Fisher’s r-to-z conversion.

*Relationship of self-reported behavior and measured behavior.* However, the behavioral self-reports are not to be trusted. Self-reported email responsiveness was only modestly related to measured email responsiveness. The correlation between individuals' (log(101-x)-transformed) self-reported percentage responsiveness and their observed percentage responsiveness was  $r = .14$  ( $p < .001$ ) – generally considered a “low” correlation in social science research. For ethicists, the correlation was so low as to be only marginally significant with this sample size ( $r = .13$ ,  $p = .08$ ). Both among recipients as a whole and for ethicists in particular, self-reported email responsiveness predicted 2% of the variance in measured responsiveness. (The untransformed correlations are similar.) Another way of looking at the issue is this: Professors who claimed to reply to at least 98% of student emails 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.<sup>25</sup> By this measure, the only subgroup showing a statistically significant relationship between self-described behavior and measured behavior was the non-philosophers: 65.8% vs. 53.0%, a 12.8% differential.<sup>26</sup> For ethicists the corresponding percentages were 66.7% vs. 58.7%, an 8.0% differential).<sup>27</sup>

*Relationship of expressed normative attitude and measured behavior.* Using a suite of correlations,  $\chi^2$  tests, 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 or whether we treated it as a continuous variable

---

<sup>25</sup>  $t(339) = 2.26$ ,  $p = .02$ .

<sup>26</sup>  $t(100) = 2.58$ ,  $p = .01$ .


<sup>27</sup>  $t(104) = 1.46$ ,  $p = .15$ .

(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% (for ethicists 64.8% vs. 60.8%, CI -9.0%-17.0%).<sup>28</sup> 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. Declining bar height from left to right reflects correlation between behavior and expressed normative view.

---

<sup>28</sup>  $t(134) = .53, p = .60; t(46) = 0.62, p = .54.$



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 weak relationship between attitude and behavior seems both a priori likely and 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. Several recent meta-analyses suggest average attitude-behavior correlations, across studies, of about  $r = .30$  to  $.50$  (Kraus 1995; Cooke and Sheeran 2004; Wallace, Paulson, Lord, and Bond 2005; Glasman and Albarracín 2006). 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, forthcoming; Schwitzgebel and Rust 2010; Schwitzgebel et al. 2012; Schwitzgebel and Rust forthcoming). In support of the background moral assumptions of this study, 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, voting in public elections, courtesy and registration free-riding at professional conferences. 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 perhaps professors teaching large classes respond to a lower percentage of student emails because they are inundated and overwhelmed. However, we see no reason 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 who use non-institutionally affiliated email addresses (e.g. “gmail.com”). 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 sender’s name is unfamiliar. Also, in our experience, it is not uncommon for students to use non-edu addresses in corresponding with their professors. Finally, and perhaps most importantly, even if professors are legitimately somewhat less responsive, across the board, to messages from unknown senders, regardless of topic, that would seem to have no direct bearing on the main issue of our research, which is the *relative* email response rates of ethicists compared to other professors,

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, in cases of implicit bias, or when perceived social pressure is high (see Kraus 1995; Wallace et al. 2005). Such factors do not appear to explain the non-relationship in the present case.

We think that even many cynics about the value of philosophical moral reflection ought to expect people to show a certain amount of attitude-behavior consistency: That's what post-hoc *rationalization* is all about, right? On the booster view of the attitude-behavior relationship, people rationally endorse norms and shape their behavior to fit them; according to the scofferist rationalization-to-behavior view, people know their behavior and then rationalize their way into self-congratulatory normative attitudes. In either case, attitudes 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. 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. On the face of it, this might seem to support a rational enervation model, on which moral reflection replaces motivationally powerful intuitions with motivationally inert general

principles, over either boosterism or a behavioral rationalization model, as suggested in the introduction.

It is possible, of course, 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 our non-finding involves combining the post-hoc behavioral rationalization model of attitude-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 attitudes and behavior unless the rationalizer knows what her behavior is. Otherwise, it will align normative attitudes 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 normative attitudes 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 no statistically significant relationship between self-reported behavior and actual behavior. This interpretation of our data would also harmonize with recent work suggesting that the order of presentation of hypothetical moral scenarios has a larger influence on philosophers', especially ethicists', judgments about related moral principles than on non-philosophers' judgments – a finding that suggests that philosophers may be more likely than other professors to

opportunistically recruit general moral principles in support of post-hoc rationalizations (in this case, post-hoc rationalizations of experimentally manipulated scenario judgments; Schwitzgebel and Cushman 2012). Expertise in ethics might be, to a substantial extent, expertise in post-hoc rationalization of opinions arrived at by largely unwelcome psychological mechanisms.

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.

Joshua Rust

Department of Philosophy

Stetson University

421 North Woodland Boulevard

DeLand, Florida 32723

[jrust@stetson.edu](mailto:jrust@stetson.edu)

Eric Schwitzgebel

Department of Philosophy

University of California at Riverside

Riverside, CA 92521-0201

[eric.schwitzgebel@ucr.edu](mailto:eric.schwitzgebel@ucr.edu)

## References

- Aristotle. 4th c. BCE/1962. *Nicomachean Ethics*. Trans. M. Ostwald. New York: Macmillan.
- Baier, Annette. 1985. *Postures of the Mind*. Minneapolis: University of Minnesota Press.
- Baumrind, Diana. 1985. "Research using Intentional Deception." *American Psychologist* 40: 165-74.
- Cooke, Richard and Paschal Sheeran. 2004. "Moderation of Cognition-intention and Cognition-behavior relations: A Meta-Analysis of Properties of Variables from the Theory of Planned Behaviour." *British Journal of Social Psychology* 43: 159-86.
- Erikson, Kai T. 1967. "A Comment on Disguised Observation in Sociology." *Social Problems* 14: 366-73.
- Glasman, Laura and Dolores Albarracin. 2006. "Forming Attitudes that Predict Future Behavior: A Meta-analysis of the Attitude-behavior Relation." *Psychology Bulletin* 132: 778-822.
- Haidt, Jonathan. 2001. "The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment." *Psychological Review* 108: 814-34.
- . 2012. *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. New York: Pantheon.
- Kant, Immanuel. 1785/1998. *Groundwork of the Metaphysics of Morals*. Trans. M. Gregor. Cambridge: Cambridge University Press.
- Knobe, Joshua, and Brian Leiter. 2007. "The Case for Nietzschean Moral Psychology." In *Nietzsche and Morality*, edited by Brian Leiter and Neil Sinhababu, 93-109. Oxford: Oxford University Press.

- 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.
- Mill, John S.. 1859/2003. "On Liberty". In *Utilitarianism and On Liberty*, edited by Mary Warnock. Malden, MA: Wiley-Blackwell.
- Nietzsche, Friedrich. 1887/1998. *On the Genealogy of Morality*, translated by Maudemarie Clark and Alan J. Swenson. Indianapolis: Hackett Publishing Co.
- Plato. 4<sup>th</sup> c. BCE/1961. *The Collected Dialogues*, edited by Edith Hamilton and Huntington Cairns. Princeton: Princeton University Press.
- Posner, Richard A. 1999. *The Problematics of Moral and Legal Theory*. Cambridge, MA: Harvard University Press.
- Schwitzgebel, Eric. 2009. "Do Ethicists Steal More Books?" *Philosophical Psychology* 22, no. 6: 711–25.
- Schwitzgebel, Eric. Forthcoming. "Are ethicists any more likely to pay their registration fees at professional meetings?" *Economics & Philosophy*.
- Schwitzgebel, Eric, and Joshua Rust. 2009. "The Moral Behaviour of Ethicists: Peer Opinion." *Mind* 118: 1043–59.
- . 2010. "Do Ethicists and Political Philosophers Vote More Often Than Other Professors?" *Review of Philosophy and Psychology* 1, no. 2: 189–99.
- . Forthcoming. "The Moral Behavior of Ethics Professors: Relationships among Self-Reported Behavior, Expressed Normative Attitude, and Directly Observed Behavior." *Philosophical Psychology*.



- Schwitzgebel, Eric, Joshua Rust, L. T. L. Huang, A. T. Moore, and J. Coates. 2012. "Ethicists' Courtesy at Philosophy Conferences." *Philosophical Psychology* 25, no. 3: 1-10.
- Schwitzgebel, Eric, and Fiery Cushman. 2012. "Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers." *Mind & Language* 27: 135-53.
- Wallace, David, Rene Paulson, Charles Lord and Charles Bond Jr.. 2005. "Which Behaviors Do Attitudes Predict? Meta-analyzing the Effects of Social Pressure and Perceived Difficulty." *Review of General Psychology* 9: 214-27.
- Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.