

## **Do Ethics Classes Influence Student Behavior?**

Eric Schwitzgebel

Department of Philosophy

University of California at Riverside

Riverside, CA 92521-0201

August 15, 2013

## **Do Ethics Classes Influence Student Behavior?**

### Abstract:

Does university-level ethics instruction influence students' practical behavior? Little direct empirical work has been done on this question. I briefly review the existing data: two studies on honesty in laboratory settings and two re-analyzable economics studies on donation to student charities. Although these data suggest no overall effect of ethics instruction on moral behavior, the research is so limited that it is difficult to draw any conclusion. I then turn to the larger literature that examines the influence of ethics classes on students' self-reported moral attitudes. This literature is so flawed that, again, it is difficult to draw any conclusion, but the most reasonable interpretation appears to be that ethics instruction has at most a small influence on student attitudes. On the assumption that the influence of ethics instruction on students' practical behavior would be substantially less than the influence on students' verbally espoused attitudes, I conclude that ethics courses probably have at most a tiny influence on behavior. Finally, I discuss reasons to worry that whatever tiny influence ethics courses do have might as easily be negative as positive.

## **Do Ethics Classes Influence Student Behavior?**

### *1. Introduction.*

Real-world moral behavior is hard to study. Maybe this partly explains the otherwise surprising absence of scientific research on the question that forms the title of this essay.

Drawing mostly on indirect evidence, I will conclude that university-level ethics instruction probably has little if any effect on student behavior, and that whatever effects there are might be negative as well as positive.

### *2. Two Laboratory Dishonesty Studies.*

I have found only two studies that aim to directly examine the effects of university-level ethics instruction on students' moral behavior: Bloodgood, Turnley, and Mudrack 2008; and Mayhew and Murphy 2009. These studies are not, I think, very compelling evidence. But let's look at them.

Bloodgood, Turnley, and Mudrack (2008) asked 230 business students to complete four timed word-search tasks. Students discarded their word-search sheets in a provided trashcan and then told the experimenters how many words they had found in each task. The students were told that they were being randomly assigned to groups of six or seven, that the person reporting the highest score in each group would receive \$10, and that every member of the group reporting the highest total score would receive \$5. Bloodgood and colleagues then retrieved the discarded word-searches from the trashcan and compared the number of words students said they had found to the number of actually circled words. In all, only 16% of participants reported having found more words than they had circled. Bloodgood and colleagues did not find that students

overall were less likely to cheat if they had taken a course in business ethics. However, they did find statistical interaction effects suggesting the possibility of less cheating among high ACT-score students and low-religiosity students if they had taken business ethics classes than if they had not. Bloodgood and colleagues interpreted this pattern as suggesting that business ethics classes do reduce student cheating, but only among less religious students who have “more room to grow” and among intelligent students who are “better at learning and applying the ethical lessons” (p. 565-566).

Mayhew and Murphy (2009) used a similar design. They gave business students an opportunity to anonymously report how well they had done on a trivia quiz. Students were paid depending on how the number of correct and incorrect answers they said they had, although they were also told that under certain conditions their answers would be checked and their payment reduced if they were found to have overreported the number correct. Mayhew and Murphy found no difference between students who had taken a business ethics course and those who had not. However, they did find that students who had taken business ethics were less likely to inflate their scores in a *non*-anonymous condition, a result Mayhew and Murphy interpreted as suggesting that ethics instruction might have “established expectations of ethical behavior” which were not internalized (p. 407).

These studies have serious limitations, if our aim is to assess the effects of ethics instruction on student behavior. In neither study were students randomly assigned into ethics vs. non-ethics courses, which opens up the possibility of uncontrolled confounding factors such as maturity, proximity to graduation, and ethics instruction from other sources. The students knew they were being observed by professors at their university, in a laboratory setting, and so might have behaved very differently from how they would have behaved in a more typical practical

setting. And it's not clear that either study accurately measures immoral behavior: The Mayhew and Murphy has a beat-the-experimenter game-like character that might make strategic misreporting ethically neutral; and it's possible that a non-negligible minority of the participants in Bloodgood et al., knowing their word-find sheets were being discarded, might not have bothered to circle every word they found, which if true could largely account for the 16% of students found "cheating". Even taking such misreporting at face value as a measure of dishonesty, it's not clear how closely such dishonesty resembles or correlates with the more serious types of misconduct that are presumably the instructional focus of business ethics classes.

Thus, I find it difficult to conclude much from these studies.

### *3. Re-Analyses of Two Economics Studies.*

A small literature, launched by Carter and Irons (1991) and Frank, Gilovich, and Regan (1993), examines whether studying economics makes people more selfish. This literature is potentially interesting as indirect evidence: If studying economics makes people more selfish, that seems to imply that university instruction can influence people's moral behavior for the worse – and if in some cases for the worse, maybe in other cases for the better?

Unfortunately the results of this economics literature are inconclusive. The evidence does seem to suggest that economics students tend to make more selfish, money-maximizing choices in settings that are explicitly structured as economic games (e.g., both offering and accepting small amounts in "ultimatum games"). But it's not clear that such selfishness generalizes to "real-world" (i.e., non-laboratory, non-"game") settings, nor that studying economics causes whatever differences do exist. None of the three existing studies in this literature that use real-world, non-self-report measures finds an effect. Laband and Beil (1999)

find economists no less likely than other social scientists to underreport income when paying income-based membership dues to academic societies. Frey and Meier (2003) and Bauman and Rose (2011), examining data from two universities in which all students were given the option of donating small fixed amounts of money to selected charities when registering for classes, found that although students majoring in business or economics were less charitable on average than students with other majors, over the course of their education their rates of charitable giving did not decrease relative to students of other majors. From this, both groups of researchers concluded that teaching economics did not reduce students' charitable giving.

Unfortunately, neither Frey and Meier nor Bauman and Rose collected data on ethics instruction in particular. However, both groups did collect data on student major. Thus, I was able to re-analyze the original raw data to see if philosophy majors, who normally take several ethics classes, increased their rates of charitable giving over the course of their studies. What I found was the mirror-image of the data on business and economics majors: Philosophy students were more likely than other students to donate to charity, but their rates of charitable giving did not detectably change over time or relative to other students. Thus, exposure to the philosophy ethics curriculum did not appear to influence philosophy majors' rates of giving to these charities. For more detailed analysis, see the Appendix.

Of course, the most sophisticated charitable givers might choose to save their charitable dollars for charities other than those selected by the Universities of Zurich and Washington in this period. And since random assignment is not possible, the data are likely to be confounded, for example with socio-economic status, which tends to be higher for humanities and social science majors than for other majors, at least in the U.S. (Leppel, Williams, and Waldauer 2001; Goyette and Mullen 2006). Furthermore, philosophy students might join and leave the major and

the university at different times than do other students and for different reasons, introducing year-by-year subject-pool differences.

To summarize Sections 2 and 3: Direct empirical evidence of the behavioral effects of university ethics education are meager and problematic. However, for whatever they are worth, the data do not point toward an overall effect.

#### *4. The Influence of Ethics Instruction on Self-Reported Attitudes.*

There are a few dozen published studies on the influence of university ethics instruction on students' self-reported moral attitudes. Conveniently, two meta-analyses from 2009 address the bulk of the research.

Waples, Antes, Murphy, Connelly, and Mumford (2009) analyze the literature on business ethics instruction, looking at a variety of attitudinal and self-report variables across 25 studies (some studies of university students, but also some studies of ethics instruction for professional businesspeople). They find an overall weighted average effect size of  $d = 0.29$ ; for studies focusing on university students,  $d = 0.28$ . (Cohen's  $d$  is the difference in means divided by the standard deviation;  $d = 0.20$  is usually considered a "small" effect size and  $0.50$  is considered a "medium" effect size.) Waples and colleagues conclude that "business ethics instruction, as reported in the literature, is at best minimally effective in enhancing ethics among students and business people" (p. 146). Similarly, Antes et al. (2009) analyze 20 studies of ethics instruction in the sciences (e.g., in biomedical ethics and in the responsible conduct of research), finding a weighted mean effect size of  $d = 0.42$  and concluding that "ethics instruction is at best moderately effective as currently conducted" (p. 397).

It is, however, I think, practically impossible to interpret such analyses without a sense of the methods and results of the studies included. Examining the target studies in detail, the results are, in my view, more disappointing than suggested even by the modest language of Waples, Antes, and collaborators quoted above. To give readers a feel for the literature, I will briefly discuss three sample studies. I selected the three studies by starting with the two studies with the most citations since 2009 in Google Scholar as of July 31, 2013. Since both were studies of business ethics, I added a third study from the other meta-analysis – the study most-cited since 2009, excluding one study that examined the medical school curriculum in general without any direct measure of ethics instruction in particular. In my view, these three studies are representative of the literature.

Ritter (2006) examined students from two versions of an Organizational Theory and Behavior class, a class required of all Management majors at a mid-sized university in the southern U.S. Students in one class received additional instruction in an ethics curriculum that students in the other class did not receive. Ritter used two measures of the effectiveness of instruction. One was a “moral awareness” measure adapted from Smith and Oakley (1997), which asked students to indicate the extent to which they found various situations, like padding one’s expense account or bribing foreign officials, “ethically acceptable”. Looking at 77 students total, Ritter found no statistically significant difference on this measure between students exposed the ethics curriculum and those not exposed.<sup>1</sup> A second measure was a “moral reasoning” measure adapted from Clarkeburn (2002) which presented students with ethical vignettes, asked how they would behave, and then asked them to describe five issues they considered in making their decision. Responses were then coded qualitatively on a 0-3 scale

---

<sup>1</sup> Dividing the measure into two factors, she reports  $t = 0.95$ ,  $p = .35$ , and  $t = -0.04$ ,  $p = .97$ .

from non-ethical to the highest level of ethicality. Students in both versions of the class were given this moral reasoning measure both at the beginning at the end of the semester. Ritter presents raw data rather than inferential statistics for this measure, but compiling the data it looks as though overall the control group had 37/56 (66%) low scores (0 or 1 on a scale of 0-3) both at the beginning and at the end of the term, while the treatment group went from 35/58 (60%) low scores to 29/58 (50%) – a change in the predicted direction, but not approaching statistical significance.<sup>2</sup> In a post-hoc analysis, Ritter reports that women appeared to have shifted more than men on both measures, including at an uncorrected p value of  $\leq .01$  on one version of the first measure. She concludes that “the positive effects of an ethics training program were witnessed only in women” (p. 161).

Eynon, Hill, and Stephen (1997) mailed a survey to a sample of 1092 certified public accountants from across the United States, achieving a 16% response rate. The survey asked recipients to report some demographic information including whether they had completed a college ethics course, to answer some opinion questions about ethics instruction, and to complete the three-scenario version of the Defining Issues Test. The Defining Issues Test (Rest 1979) presents respondents with paragraph-long moral dilemmas, such as whether to steal a highly-priced drug to save one’s spouse, and then asks respondents to rate the importance of different possible factors one might consider in deciding how to act. The DIT designers regard some of the factors as characteristic of mature “postconventional” moral thinking (e.g., “Whether the law in this case is getting in the way of the most basic claim of any member of society”) and others as characteristic of less mature thinking (e.g., “Whether the druggist deserves to be robbed for being so greedy and cruel”). In a multiple regression analysis, Eynon and colleagues found that

---

<sup>2</sup>  $Z = 1.1$ ,  $p = .26$ , 95% CI for diff -7.7% to 28.3%

CPAs who reported having taken a college ethics course tended to score about 8 points higher on the DIT, on a scale from 0 to 90. The authors interpret this result as “support for the effectiveness of ethics interventions” (p. 1304).

Myyry and Helkama (2002) report a study of 50 social psychology students who participated in a 20-hour ethics education curriculum in Finland, compared to a control group of six education students who participated in a qualitative research methods course. Participants read a 500-word story both at the beginning and at the end of the instruction. The story was a complex scenario involving child abuse, and participants were asked to “single out the elements that should be considered” in deciding whether the child should be placed in foster care. Answers were scored by the investigator, bearing in mind the extent to which the answers revealed “sensitivity to the special characteristics of the patient” and “awareness of what actions serve the rights and welfare of others”. Possible scores ranged from 0 to 34. Students who had completed the ethics curriculum showed a statistically marginal increase in score, from 8.9 to 10.0.<sup>3</sup> However, since the small control group’s mean score declined, the results of a statistical analysis comparing the two groups showed a greater positive shift for the ethics students than for the controls.<sup>4</sup>

The reader with a critical eye will find many causes for concern about these studies, including small sample sizes or low response rates, lack of proper controls, possible experimenter bias in coding, dubious use of statistics, participants possibly being influenced by guesses about the nature of the study, experimenters “teaching to the test”, and the validity of the measures. Not every study in the literature is subject to all of these concerns, but the three studies reported above are typical of the state of the literature in – if I may put it frankly – the

---

<sup>3</sup>  $t = 1.85, p = .07.$

<sup>4</sup> ANOVA,  $F = 7.7, p < .01.$

extent to which they invite considerable skepticism. More recent studies and older studies not included in the Waples and Antes meta-analyses are approximately similar in methods and quality (e.g., Jones 2009; Antes et al. 2010; Lau 2010; Warnell 2010; Lawrence, Reed, and Locander 2011; Burns 2012; Harkrider et al. 2012; Johl, Jackling, and Wong 2012; May and Luth 2013).

Several important sources of distortion will tend to favor finding an effect where none exists. Experimenters generally want positive results, perhaps especially if they are examining the effectiveness of curricula they support. Participants tend to want to please experimenters by confirming their suspected hypotheses (Orne 1962; Rosnow and Rosenthal 1997). And positively-valenced measures – presumably including various types of educational attainment and various measures of moral sensitivity – often correlate for a variety of reasons other than a direct causal relationship among the correlates (Meehl 1990). Under such conditions, standard quantitative meta-analyses like those of Waples and Antes might tend to reveal a positive relationship even if none exists. In my judgment, the most striking feature of this literature is not that when one combines the results in meta-analysis a modest relationship shows up. Rather, it's that that despite the likely existence of several sources of bias toward positive results, studies do not report larger and more consistent relationships between what students are taught and what they later say, often to those same teachers. Ritter and Myyry and Helkama seem almost to beg their students to confirm their hypotheses, but they still find only modest and inconsistent results. Tentatively, I conclude the following: The effect of ethics instruction on students' moral attitudes is at most small.

What can we conclude from this about student moral *behavior*? It seems plausible that ethics classes would have less influence on students' practical behavior than on their self-

reported attitudes. Most university ethics instruction engages only indirectly with students' behavior outside the classroom, via engaging directly with students' attitudes. It's one thing to teach a student to verbalize a motto such as "the interests of all stakeholders should be given weight in corporate decision-making" or "you should not pad expense accounts", and it's quite another to lead a student to act differently, in light of such mottoes, than she would otherwise have acted. Thus, I suggest that we should not expect ethics instruction to have a medium-sized influence on practical behavior unless it has a large influence on attitudes, nor more than a tiny influence on practical behavior unless it has at least medium-sized influence on attitudes. I recommend a similar approach to evaluating the duration of effect: no medium-sized long-term influence without a large short-term influence, and no more than a tiny long-term influence without at least a medium-sized short-term influence.

Pulling it together: The existing literature on ethics instruction almost exclusively examines short-term effects on attitudes. (The Eynon et al. 1997, discussed above, is an exception, but very difficult to interpret without convergent evidence from other methods.) Those effects appear to be small overall. Thus, I recommend that we tentatively conclude that both the short-term behavioral effects and the long-term attitude effects are at most tiny. The long-term behavioral effects, if any, would then be still tinier.

##### *5. Plausibility Considerations and Other Types of Indirect Evidence.*

Sometimes I teach ethics courses. I hope that my students learn the types of things from my ethics courses that they learn from other courses in the humanities, including critical thinking, factual knowledge of the range of positions and arguments, and an appreciation of the

history of ideas. However, I also hope that my instruction has a positive effect on their moral behavior, including long after they have graduated from university.

My tentative conclusions above needn't destroy that hope. Those conclusions might even be entirely unsurprising. Of course – I might think – any practical effect on student behavior would have to be tiny! Most students do not apply what they learn in the classroom to their everyday lives, even when the connection might seem obvious to a specialist in the field; and many or most students will almost entirely forget the content of most of their university courses. Still, some minority of students will presumably remember, long after graduation, pieces of their classroom discussions of Kant vs. Mill, or of shareholders vs. stakeholders, or of the importance of informed consent in medical practice. And maybe some portion of that minority will then make a more ethical choice than they otherwise would have. Such an effect might be statistically undetectable by standard social scientific methods, but even if it's tiny it might be hugely important, if ten years from now one doctor or businessperson or parent from my class of three hundred makes a more ethical choice than she otherwise would on a matter of major importance for others' lives. For me, this hope is part of what makes ethics worth teaching. Nothing I've said so far compels me to abandon the view that university ethics instruction can sometimes have tiny but important long-term behavioral effects.

However, university ethics instruction might also have similarly tiny long-term *negative* effects on moral behavior. One way it might do so is through slightly improving students' capacities at "toxic rationalization" (Schwitzgebel and Rust forthcoming). Suppose you are tempted toward some unethical act. You are on the cusp, but you lean against. Wow, would profits jump! – and you'd get that big bonus and promotion, but there's just a small risk that toxins would leak into the water table. But then your ethics education comes back to you.

Milton Friedman (1970) argued that your real moral duty as a manager is to maximize profits, not to use your position to advance a moral agenda that the company's stockholders might reject. Your ethics professor had used Friedman as a foil, hoping to convince you that businesspeople should use their power for the good of the community; but Friedman is what comes to you, and Friedman provides the rationalization you are seeking, the rationalization that permits you to decide (wrongly, I'm assuming) that the gain in profit outweighs the small risk to the water table, a risk which should be addressed to the extent necessary by other social mechanisms. Or maybe utilitarian moral theory proves convenient to help you superficially rationalize sacrificing some patient's interest for what you regard as the greater good of your medical research – an immoral choice you would not have made had the rationalization not been so readily available. The more theories you have ready to hand, and the cleverer your inner lawyer, the more resources you can bring to the task of rationalizing attractive misconduct.

Bernard Williams (1985) and Annette Baier (1985) have argued that philosophical debates about moral issues can undercut our intuitive sense of what's right, leaving us cynical or at sea. Hans-Georg Moeller (2009) has argued that philosophical ethics can lead to arrogance and retributive violence. Arguably, detailed knowledge from case studies of the past unethical behavior of corporate managers, doctors, political leaders, and others – especially if, realistically, most bad actors seem to profit from their misconduct – might loosen one's sense of what is thinkable or morally acceptable, or at least the impression that moderate infractions are not as bad as what others regularly get away with. An ethics class might leave a student feeling that ethical standards are unattainable, or the student might hate the class and be turned off forever to intellectual conversations about ethics. A student might learn that "ethics" is about the rigid application of stupid rules, or the kind of thing that concerns badly dressed sixty-year-old men,

or leads only to interminable, unwinnable debates. A student might learn more from the teacher's dismissive, judgmental demeanor than from the teacher's verbally espoused philosophy. Our lives and conversations are permeated with morally-loaded discourse, in natural and authentic discussion with friends, family, and business associates; and what overall effect a sophomore ethics class might have on this mix, if any, is difficult to guess.

One reason to think that the negative influences of university ethics instruction might approximately cancel out any positive influences is this: Ethics *instructors*, at least in philosophy departments – who are more intimately familiar with the university ethics curriculum than is anyone else in the world – don't appear to behave any morally better than do professors not specializing in ethics (Rust and Schwitzgebel forthcoming; Schwitzgebel and Rust forthcoming). Maybe this is because whatever moral insight they gain from their studies is approximately counterbalanced by enhanced skill at toxic rationalization or by loss of contact with emotionally-grounded conventional norms.

We might also compare university ethics instruction to religious moral instruction. An extensive literature addresses the relationship between religiosity and moral behavior – a literature that is in some ways similar to the literature on the effects of ethics instruction on moral attitudes. That is to say, a literature full of problematic studies, confounds, lack of random assignment (of course!), dubious self-report measures, and very likely tainted by a general positive-results bias; and which overall shows weak and inconsistent relationships. Although reviews of this literature often conclude that religiosity has a small but real positive influence on moral behavior (e.g., Baier and Wright 2001; Regnerus 2003; McCullough and Willoughby 2009), a skeptic might equally well conclude that religious thinking and religious instruction have both positive and negative influences on moral behavior, through a variety of mechanisms,

and that these influences likely approximately cancel out overall. If you are inclined to agree, then a comparison to ethics instruction in secular universities invites itself. Well-meaning moral exhortation might not consistently produce behavioral change in the direction hoped for.

#### *6. Conclusion.*

I tentatively conclude that the typical university ethics class has at most a tiny effect on students' moral behavior. Even tiny positive effects on moral behavior can be highly important, but unfortunately such tiny positive effects, if they exist, might be entirely canceled or outweighed by negative effects. Given the lack of direct evidence, it is hard to feel much confidence, but the most reasonable guess, I suggest, is that the average ethics class has an average moral effect on student behavior very close to zero and approximately as likely to be slightly negative as slightly positive.

Maybe, given the mediocre moral qualities of ethics instructors themselves and the dubious theories we endorse, we and the world should be grateful we have so little influence – a comfortable thought, in a way! But I'm inclined to favor, instead, a less comfortable thought: that an ethics class that has no overall positive influence on students' behavior is, in most cases, a half-broken thing.

References:

- Antes, Alison L., Xiaoqian Wang, Michael D. Mumford, Ryan P. Brown, Shane Connelly, and Lynn D. Devenport (2010). Evaluating the effects that existing instruction on responsible conduct of research has on ethical decision making. *Academic Medicine*, 85, 519-526.
- Antes, Alison L., Stephen T. Murphy, Ethan P. Waples, Michael D. Mumford, Ryan P. Brown, Shane Connelly, S., and Lynn D. Devenport (2009). A meta-analysis of ethics instruction effectiveness in the sciences. *Ethics & Behavior*, 19, 379-402.
- Baier, Annette (1985). *Postures of the mind*. Minneapolis: University of Minnesota.
- Baier, Colin J., and Bradley R.E. Wright (2001). "If you love me, keep my commandments": A meta-analysis of the effects of religion on crime. *Journal of Research in Crime and Delinquency*, 38, 3-21.
- Bauman, Yoram, and Elaina Rose (2011). Selection or indoctrination? Why do economics students donate less than the rest? *Journal of Economic Behavior and Organization*, 79, 318-327.
- Bloodgood, James M., William H. Turnley, and Peter Mudrack (2008). The influence of ethics instruction, religiosity, and intelligence on cheating behavior. *Journal of Business Ethics*, 82, 557-571.
- Burns, David J. (2012). Exploring the effects of using consumer culture as a unifying pedagogical framework on the ethical perceptions of MBA students. *Business Ethics*, 21, 1-14.
- Carter, John R., and Michael D. Irons (1991). Are economists different, and if so, why? *Journal of Economic Perspectives*, 5, 171-177.

- Clarkeburn, Henriikka (2002). A test for ethical sensitivity in science. *Journal of Moral Education, 31*, 439-453.
- Eynon, Gail, Nancy Thorley Hill, and Kevin T. Stevens (1997). Factors that influence the moral reasoning abilities of accountants: Implications for universities and the profession. *Journal of Business Ethics, 16*, 1297-1309.
- Frank, Robert H., Thomas Gilovich, and Dennis T. Regan (1993). Does studying economics inhibit cooperation? *Journal of Economic Perspectives, 7*, 159-171.
- Friedman, Milton (1970). The social responsibility of business is to increase profits. *New York Times Magazine* (Sep. 13), 32 ff.
- Frey, Bruno S., and Stephan Meier (2003). Are political economists selfish and indoctrinated? Evidence from a natural experiment. *Economic Inquiry, 41*, 448-462.
- Goyette, Kimberly A., and Ann L. Mullen (2006). Who studies the arts and social sciences? Social background and the choice and consequences of undergraduate field of study. *Journal of Higher Education, 77*, 497-538.
- Harkrider, Lauren N., Chase E. Thiel, Zhanna Bagdasarov, Michael D. Mumford, James F. Johnson, Shane Connelly, and Lynn D. Devenport (2012). Improving case-based ethics training with codes of conduct and forecasting content. *Ethics & Behavior, 22*, 258-280.
- Johl, Shireenjit, Beverly Jackling, and Grace Wong (2012). Ethical decision-making of accounting students: Evidence from an Australian setting. *Journal of Business Ethics Education, 9*, 51-78.
- Jones, David Allen (2009). A novel approach to business ethics training: Improving moral reasoning in just a few weeks. *Journal of Business Ethics, 88*, 367-379.

- Laband, David N., and Richard O. Beil (1999). Are economists more selfish than other 'social' scientists? *Public Choice*, 100, 85-101.
- Lau, Cubie L. L. (2010). A step forward: Ethics education matters! *Journal of Business Ethics*, 92, 565-584.
- Lawrence, Katherine E., Kendra L. Reed, and William Locander (2011). Experiencing and measuring the unteachable: Achieving AACSB learning assurance requirements in business ethics. *Journal of Education for Business*, 86, 92-99.
- Leppel, Karen, Mary L. Williams, and Charles Waldauer (2001). The impact of parent occupation and socioeconomic status on choice of college major. *Journal of Family and Economic Issues*, 22, 373-394.
- May, Douglas R., and Matthew T. Luth (2013). The effectiveness of ethics education: A quasi-experimental field study. *Science and Engineering Ethics*, 19, 545-568.
- Mayhew, Brian W., and Pamela R. Murphy (2009). Impact of ethics education on reporting behavior. *Journal of Business Ethics*, 86, 397-416.
- McCullough, Michael E., and Brian L.B. Willoughby (2009). Religion, self-regulation, and self-control: Associations, explanations, and implications. *Psychological Bulletin*, 135, 69-93.
- Meehl, Paul E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, 66, 195-244.
- Moeller, Hans-Georg (2009). *The moral fool*. New York: Columbia.
- Myry, Liisa, and Klaus Helkama (2002). The role of value priorities and professional ethics training in moral sensitivity. *Journal of Moral Education*, 31, 35-50.

- Orne, Martin T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, 17, 776-783.
- Regnerus, Mark D. (2003). Religion and positive adolescent outcomes: A review of research and theory. *Review of Religious Research*, 44, 394-413.
- Rest, James R. (1979). *Development in judging moral issues*. Minneapolis: University of Minnesota.
- Ritter, Barbara A. (2006). Can business ethics be trained? A study of the ethical decision-making process in business students. *Journal of Business Ethics*, 68, 153-164.
- Rosnow, Ralph L., and Robert Rosenthal (1997). *People studying people*. New York: Freeman.
- Roundtree, Aimee Kendall (2010). Success and failures teaching visual ethics: A class study. *Journal of Effective Teaching*, 10, 55-61.
- Rust, Joshua, and Eric Schwitzgebel (forthcoming). The moral behavior of ethicists and the power of reason. In H. Sarkissian and J. Wright, eds., *Advances in Experimental Moral Psychology*. Continuum.
- Schwitzgebel, Eric, and Joshua Rust (forthcoming). The moral behavior of ethicists. In W. Buckwalter and J. Sytsma, eds., *Blackwell companion to experimental philosophy*. Blackwell.
- Smith, Patricia L., and Ellwood F. Oakley, III (1997). Gender-related differences in ethical and social values of business students: Implications for management. *Journal of Business Ethics*, 16, 37-45.

Waples, Ethan P., Alison L. Antes, Stephen T. Murphy, Shane Connelly, and Michael D.

Mumford (2009). A meta-analytic investigation of business ethics instruction. *Journal of Business Ethics*, 87, 133-151.

Warnell, Jessica McManus (2010). An undergraduate business ethics curriculum: Learning and moral development outcomes. *Journal of Business Ethics Education*, 7, 63-84.

Williams, Bernard (1985). *Ethics and the limits of philosophy*. London: Fontana.

Appendix: Charitable giving of philosophy majors at University of Zurich and University of Washington.

In the Frey and Meier data from University of Zurich 1999-2004, philosophy majors were statistically more likely than other students to donate to at least one charity in their first semester: 82% vs. 75%.<sup>5</sup> That difference continued throughout schooling. For example, by the eighth semester, it was 83% vs. 73%.<sup>6</sup> Given the huge dataset, the non-philosophy majors' 2% decline in donation rate from 1<sup>st</sup> to 8<sup>th</sup> semester was statistically significant. Although no such decline was detectable among philosophy majors, a 2% decline is well within the 95% confidence interval of the difference between philosophy majors' 1<sup>st</sup> and 8<sup>th</sup> semester charity rates.<sup>7</sup> This lack of a statistically detectable difference in rate of change is confirmed by a logistic regression predicting donation from number of semesters of schooling (excluding students with over 10 semesters), philosophy major, and the interaction of semesters\*philosophy.<sup>8</sup> (An interaction effect would indicate a different relationship between amount of schooling and donation rates among philosophy majors than among non-philosophy majors.) See Figure 1.

---

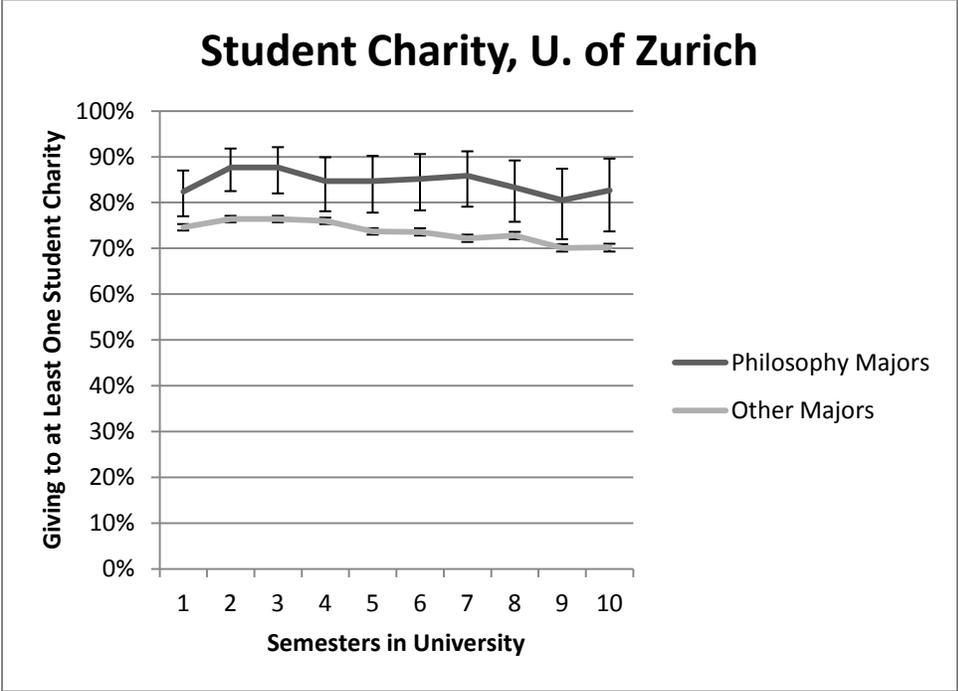
<sup>5</sup> 196/238 vs. 12594/16875,  $Z = 3.1$ ,  $p = .002$ , 95% CI for diff +2.8 to + 12.6%. This dataset includes seven additional semesters' worth of data not reported in Frey and Meier 2003.

<sup>6</sup> 110/132 vs. 9065/12451,  $Z = 3.2$ ,  $p = .001$ , 95% CI for diff +4.1% to +17.0%.

<sup>7</sup>  $Z = -3.5$ ,  $p < .001$ , 95% for diff -0.8% to -2.8%;  $Z = 0.2$ ,  $p = .81$ , 95% CI for diff -7.0% to +9.0%.

<sup>8</sup> Number of semesters: OR = .97,  $Z = -16.1$ ,  $p < .001$ ; philosophy major: OR = 1.83,  $Z = 4.3$ ,  $p < .001$ ; interaction: OR = 1.01,  $Z = 0.5$ ,  $p = .64$ , 95% CI for interaction OR 0.96 to 1.06. Excluding students with more than 10 semesters, the correlation of semesters of schooling with a dummy of 1 for having donated is  $r = -.02$  ( $p = .36$ ) for philosophy majors and  $r = -.04$  ( $p < .001$ ) for other majors. (Students with many semesters of schooling are socially different from typical undergraduates and their time-charity relationship looks non-linear.)

Figure 1: Percent of students giving to at least one University of Zurich student charity, by semesters of university education, philosophy majors vs. all other majors. Error bars indicate one-proportion 95% confidence intervals.



The Bauman and Rose dataset, from University of Washington 1999-2002, confirms the Frey and Meier findings on a different student population. In their first year of university, philosophy majors were statistically more likely than other students to give to at least one charity: 38% vs. 21%.<sup>9</sup> And again, the difference continued through senior year, when the donation rate was 31% for the philosophy majors vs. 19% for the rest.<sup>10</sup> As in the Frey and Meier, the dataset was sufficiently large that the 2% decline among non-majors was statistically significant.<sup>11</sup> Also as in the Frey and Meyer, although there was no statistically detectable change in philosophy majors' donation rates, the 2% non-majors' decline was well within the 95% confidence interval of the difference in donation rates between philosophers' 1<sup>st</sup> and 4<sup>th</sup> years.<sup>12</sup> Finally, just like the Frey and Meyer, a logistic regression predicting donation to at least one charity from number of years in school, philosophy major, and the interaction of years\*philosophy finds no evidence of an interaction effect.<sup>13</sup> See Figure 2.

Thanks to both research groups for sharing their raw data with me.

---

<sup>9</sup> 43/114 vs. 3290/15497,  $Z = 3.6$ ,  $p < .001$ , 95% CI for diff +7.6% to +25.4%.

<sup>10</sup> 77/249 vs. 4682/24574,  $Z = 4.0$ ,  $p < .001$ , 95% CI + 6.1% to +17.6%.

<sup>11</sup>  $Z = 5.3$ ,  $p < .001$ , 95% CI -1.4% to -3.0%.

<sup>12</sup>  $Z = 1.3$ ,  $p = .28$ , 95% CI -17.4% to +2.8%. The correlation of years of schooling with a dummy of 1 for having donated is  $r = -.06$  ( $p = .11$ ) for philosophy majors and  $r = -.02$  ( $p < .001$ ) for other majors.

<sup>13</sup> Number of years: OR = .95,  $Z = -6.1$ ,  $p < .001$ ; philosophy major: OR = 2.43,  $Z = 4.1$ ,  $p < .001$ ; interaction: OR = .93,  $Z = -0.9$ ,  $p = .35$ , 95% CI for OR 0.81 to 1.08.

Figure 2: Percent of students giving to at least one University of Washington student charity, by years of university education, philosophy majors vs. all other majors. Error bars indicate one-proportion 95% confidence intervals.

