

Do Ethics Classes Influence Student Behavior? Case Study: Teaching the Ethics of Eating Meat

Eric Schwitzgebel¹, Brad Cokelet² & Peter Singer³

¹University of California at Riverside

² University of Kansas

³ Princeton University

Author Note

Correspondence concerning this article should be addressed to Eric Schwitzgebel, Department of Philosophy, University of California at Riverside, Riverside, California, 92521-0201, USA. Email: eschwitz at ucr.edu.

Abstract

Do university ethics classes influence students' real-world moral choices? We aimed to conduct the first controlled study of the effects of ordinary philosophical ethics classes on real-world moral choices, using non-self-report, non-laboratory behavior as the dependent measure. We assigned 1332 students in four large philosophy classes to either an experimental group on the ethics of eating meat or a control group on the ethics of charitable giving. Students in each group read a philosophy article on their assigned topic and optionally viewed a related video, then met with teaching assistants for 50-minute group discussion sections. They expressed their opinions about meat ethics and charitable giving in a follow-up questionnaire (1032 respondents after exclusions). We obtained 13,642 food purchase receipts from campus restaurants for 495 of the students, before and after the intervention. Purchase of meat products declined in the experimental group (52% of purchases of at least \$4.99 contained meat before the intervention, compared to 45% after) but remained the same in the control group (52% both before and after). Ethical opinion also differed, with 43% of students in the experimental group agreeing that eating the meat of factory farmed animals is unethical compared to 29% in the control group. We also attempted to measure food choice using vouchers, but voucher redemption rates were low and no effect was statistically detectable. It remains unclear what aspect of instruction influenced behavior.

Keywords: consumer choice; ethics instruction; experimental philosophy; moral psychology; moral reasoning; vegetarianism

Do Ethics Classes Influence Student Behavior? Case Study: Teaching the Ethics of Eating Meat

Do university ethics classes influence students' real-world moral choices? The question is important, but difficult to answer.

The question is important partly for curricular reasons relevant to the humanistic mission of the university and the practical value of teaching ethics courses (Abend, 2014; Nussbaum, 1997). It is also important because of what the answer would reveal about the relationship between moral behavior and moral cognition of the sort that is fostered in ethics classes. A long philosophical tradition holds that we can become morally better people by reflecting philosophically in the right sort of way – or at least that good philosophical reasoning can reduce the likelihood of acting immorally on the basis of tempting but bad arguments (Hadot, 1995; Ivanhoe, 2000; Kant, 1785/1996; Mill, 1859/2003). Also, most empirical psychological models that address the relationship between moral reasoning and moral behavior allow that reasoning can have an important influence on moral behavior (Cushman, Young, & Greene, 2010; Ellemers, van der Toorn, Paunov, & van Leeuwen, 2019), including presumably reasoning of the sort encouraged in philosophy classes.

On the other hand, there are grounds for pessimism about whether exposure to academic philosophy influences moral behavior. Often, philosophical moral reasoning appears to be recruited only after the fact to justify ethical conclusions one would have reached anyway (Haidt, 2012; Schwitzgebel & Ellis, 2017). Also, professional ethicists, who presumably engage in philosophical moral reflection more often than and at least as well as people who do not specialize in ethics, appear to behave no morally better than do others of similar social background, across a wide variety of measures (Schönegger & Wagner, 2019; Schwitzgebel & Rust, 2016).

Such claims about the behavioral effects of philosophical reflection are at least in principle empirically testable (Haidt, 2012; Kohlberg 1981; May, 2018; Narvaez, 2014; Schwitzgebel, 2007). They are, however, difficult to test directly. Real-world moral behavior is in general difficult to accurately and systematically measure, and few philosophers are able to run controlled experiments with large numbers of students randomly assigned to one sort of ethics instruction or another. If any controlled experiments exist that measure the effects of university-level philosophy ethics classes on real-world, non-laboratory, non-self-reported moral behavior, we are unaware of them. The empirical literature on the effectiveness of business ethics and medical ethics instruction, for example, is generally of poor quality, shows weak results, and uses unconvincing self-report or laboratory measures (Schwitzgebel, 2013). A recent meta-analytic review of ethics instruction in the sciences (Watts, Medeiros, Mulhearn, Steele, Connelly, & Mumford, 2017) does not even attempt to look at behavioral measures as a criterion of instructional success, noting that almost all studies use measures of moral reasoning, moral attitude, knowledge of standards, or Rest's Defining Issues Test (Rest, Narvaez, Thoma, & Bebeau, 1999) as their main dependent variables. We aimed to run the first controlled experiment testing the real-world behavioral consequences of university-level philosophy instruction.

We chose the ethics of eating meat as our test case for three reasons. First, opinion on this question in our target population (university students in California) is highly variable and probably unstable, and thus it seemed a promising target for opinion change, compared to ethical issues on which opinion is more firmly fixed. Second, among philosophers who write on the issue there is widespread (though not perfect) consensus that it is generally morally better for the typical North American to eat less factory farmed meat (Adams 1990/2015; Camosy, 2013; DeGrazia, 1996; Huemer, 2019; Korsgaard, 2018; Pollan, 2006; Regan, 1983; Scruton, 2004; Singer 1975/2009). Thus, in exposing students to the philosophical arguments in favor of this ethical conclusion, we would be exposing them

to arguments about which the range of reasonable disagreement spreads mainly between the view that avoiding factory farmed meat is morally required to the view that avoiding factory farmed meat is morally good but not required. Good philosophical reflection on these issues should then tend to move students, if it moves them at all, toward lower factory farmed meat consumption, rather than in divergent directions as might be expected on more contentious issues. Third, meat consumption is directly measurable and its choice points are obvious. Success in acting on an ethical decision to “be kinder”, though perhaps measurable, is a much fuzzier business. “Avoid dishonesty”, though somewhat more concrete, often admits of excuses, and decisions about whether to be honest are not easily countable and not always salient. In contrast, a decision to avoid factory farmed meat when other options are available is in most cases straightforwardly implementable; there are a few consistent choice points per day (one’s meal purchases); and with access to receipts, meal purchase choices can be accurately measured without need of self-report. For students eating at a standard university cafeteria, avoiding factory farmed meat consumption is likely to require avoiding meat altogether.

Anecdotally, it is common for people to claim that they have been moved by philosophical arguments to become practicing vegetarians. Primatologist Jane Goodall, for example, says she converted to vegetarianism after reading Peter Singer’s *Animal Liberation* (Goodall, McAvoy, & Hudson, 2005; Singer 1975/2009), as does Ingrid Newkirk, the president and co-founder of People for the Ethical Treatment of Animals (PETA, 2015). Recent empirical research based on self-reports also suggests that philosophical argumentation or presentation of information about the conditions of factory farmed animals can lead some people to reduce their meat consumption (Feltz & Feltz, 2019; Wright, in press). All of this provides some hope that we would find an effect.

However, other research is less encouraging: Philosophy professors, despite being much more likely than other professors to say that eating meat is morally bad, self-report eating meat at about the

same rate as other professors, suggesting that exposure to philosophical argumentation on this issue might shift their attitudes relative to other professors but not their behavioral choices (Schwitzgebel & Rust, 2014; though see discussion in Schöneegger & Wagner, 2019, and Schwitzgebel, 2019b). Also, self-report of vegetarianism is not as highly correlated as one might expect with the avoidance of eating meat, perhaps due to varying, flexible, or aspirational conceptions of “vegetarianism” (Haddad and Tanzman, 2003; Vinnari, Montonen, Härkänen, & Männistö, 2009). Thus, even if self-reports change after exposure to philosophical material, actual meat consumption might change much less. Given the inaccuracies of self-report of dietary change, The Humane League now recommends that future research on the effectiveness of animal advocacy programs focus to the extent possible on purchase data collected directly from retailers or other objective measures (Peacocke, 2018).

Some people who eliminate or reduce their consumption of factory farmed meat after reading philosophical work on the topic emphasize the logical rather than the emotional appeal of the work (e.g., Dhammika, 2019). However, philosophical writing and teaching contain elements other than argumentation in the narrow sense – such as appeals to emotion, practical advice, presentation of exemplars, and social influence from instructors and fellow students. Although we hope that reasoning and argumentation play a central role in philosophical reading and in the philosophy classroom, philosophical instruction as it occurs “in the wild”, so to speak – in actual philosophy classrooms, taught by actual philosophy instructors using their usual teaching methods – might influence students’ behavioral choices through a variety of mechanisms that are difficult to control, balance, and empirically separate. Although we hope in follow-up research to begin to isolate mechanisms by manipulating instructional techniques, this study aims only to explore the presence or absence of an effect of instruction without manipulation of teaching variables.

We divided students in four large introductory classes in a large public university in California into two groups: a group that focused on the ethics of eating meat and a group that focused on the ethics of charitable giving. Students in the meat ethics group were required to read a philosophical essay arguing that it is unethical to eat the meat of factory farmed animals, and they analyzed and critically assessed the argument in one fifty-minute group discussion. Students in the ethics of charity group were required to read a philosophical essay arguing that it is unethical to spend money on luxuries instead of donating it to worthy charitable organizations, and they analyzed and critically assessed the argument in one fifty-minute group discussion. In a subsequent anonymous questionnaire, all students were asked their opinions about several moral issues, including charitable giving and meat eating. Finally, we examined dining-card purchase data for both groups of students, before and after the group discussions.

All three authors hypothesized that students exposed to the material on the ethics of eating meat would be substantially more likely to disapprove of meat eating in the questionnaire. Two of the authors, Cokelet and Singer, based on their own experience teaching the ethics of eating meat, hypothesized that students in that condition would also purchase substantially less meat. Singer in particular is widely known around the world for his philosophical advocacy of vegetarianism, which he believes is behaviorally effective. The other author, Schwitzgebel, based on his previous research described above, hypothesized that any difference in meat-purchasing behavior would be statistically undetectable between the groups. If philosophy professors, despite being much more likely than comparison groups to describe meat eating as morally bad, do not detectably differ in their food choices, it seems reasonable to conjecture that students exposed to philosophical ethical arguments against eating meat would behave no differently from students not exposed to such arguments. Thus, on the central questions of the study, this was a partly adversarial collaboration.

Method

Participants and Instructors

Participants were 1332 students enrolled in four introductory Philosophy courses at University of California, Riverside: Philosophy 1 (Introduction to Philosophy) in Spring 2017 and Fall 2017, Philosophy 2 (Contemporary Moral Issues) in Spring 2017, and Philosophy 5 (Evil) in Fall 2017. The courses ranged in size from 263 to 367 students. Philosophy 5 was taught by one of the authors of this article (Schwitzgebel). The other classes were taught by other Philosophy Department faculty. Philosophy graduate student teaching assistants (TAs) – 15 in all, across the four courses – ran weekly group discussion sections, with enrollment capped at 25 students per section and three sections per TA. Four of the TAs served in both Spring and Fall.

Design

In a week selected by the instructor, half of the course discussion sections covered the ethics of charity and half covered the ethics of eating meat. Neither the ethics of eating meat nor the ethics of charity was discussed by the course instructor during the main course lectures.

Approximately half of the TAs were assigned to teach the meat ethics material, and approximately half of the TAs were assigned to teach the charity ethics material. The TAs taught either only meat ethics sections or only charity ethics sections. Teaching duties were divided to approximately balance time of day and day of week between the meat ethics sections and the charity ethics sections. Because some classes had an odd number of TAs and not all sections filled with students, it was not possible to fully balance the meat and charity sections. In Spring 2017, most of the sections were paired in the same time slots (e.g., one meat ethics section and one charity ethics section, both at 9:00 a.m. on the same day of the week) and did not differ in any features visible in the student enrollment system. During registration, we slowly raised the enrollment caps for these sections from 5 to 25, ensuring quasi-random sorting of students into paired sections. However, we were unable to implement this quasi-

random enrollment procedure in Fall 2017. Both ethical concerns and concerns about maintaining the good will of instructors, TAs, and other campus staff prevented us from exerting complete managerial control over enrollment, TA assignment, and section day and time balancing.

Students in all discussion sections received normal instruction except for one week in the second half of the term. During this week, the main course instructor announced that as part of a teaching experiment, students would have different reading assignments depending on which section they were enrolled in. The TAs emailed the relevant reading materials to their students, readings either on meat or charity ethics. The email also contained a link to an optional related video. Students were encouraged to complete the reading before attending section. After the group discussion meetings, instruction resumed as normal.

Since our aim was to study the effects of philosophy instruction “in the wild”, so to speak, we did not attempt to control the teaching methods of the TAs, instead encouraging them to lead discussion in their usual style. TAs, all of whom knew the design and hypotheses of the study, volunteered to teach either meat ethics or charity ethics, whichever was of more interest to them. As it happened, more than half of the TAs were practicing vegetarians (reflecting the high rate of vegetarianism among U.S. philosophy PhD students), and the meat ethics sections were all led by these vegetarian TAs. In accord with their relatively personal or impersonal teaching styles, some but not all of the TAs revealed their personal ethical choices to the students in their discussion sections: their eating choices in the sections on the ethics of eating meat or, in the ethics of charity sections, their personal habits of charitable giving or luxurious spending. In future studies we plan to explore whether instructor attitudes, and communication of those, influence student behavior. But in the present study we believe that the decision to allow TAs to use their familiar teaching styles and to choose topics of interest to them enhances the ecological validity of the design. Most philosophy instructors in the U.S., especially

individual instructors leading small classes, have considerable liberty to choose teaching topics of interest to them, and we believe that typical university ethics classes with sections on the ethics of eating meat are taught by instructors and with methods that are similar to the ones in our study. (Post-hoc we asked TAs about their teaching style and what they thought might have been effective or ineffective, but this was not formally coded or entered into any of our statistical models.)

Approximately one week after the discussion meetings, we emailed all students a link to an anonymous questionnaire which they could complete for a small amount of extra credit, followed a few days later by a reminder link. We also encouraged instructors and TAs to alert their students to this extra credit opportunity. Around the same time, students also received a \$10 voucher from U.C. Riverside's Campus Dining Services for use at one campus restaurant. This voucher enabled us to track actual purchases, with names replaced by unique identifiers for privacy. Subsequently, although it was not part of the original design, we were able to obtain from the Oracle corporation all students' campus meal card purchases during Spring and Fall terms in 2017, again tagged to unique identifiers for privacy.

Our analyses rely on two-tailed statistical tests with an alpha level of .05. The main hypotheses, overall method, and main statistical results for the questionnaire and voucher portions in Spring 2017 were pre-registered at the Open Science Foundation [link to be provided after embargo]. Due to low voucher redemption rates in Spring 2017, we did not achieve our target N of 200 per group for that portion of the study and repeated the study in Fall 2017, as discussed below. We were unable to preregister our analyses of the food purchase data from Oracle, since we did not know in advance what the data would look like or what sorts of analyses would be possible. Because of risks to privacy, and also due to the proprietary nature of Oracle's detailed food purchase data, we have made public only the summary food purchase data. All stimulus materials and data are available at [link to be provided after embargo].

Teaching Materials

Students in the meat ethics condition were required to read James Rachels' article "The Basic Argument for Vegetarianism" (Rachels, 2004), an introductory-level ten-page philosophy article arguing that it is unethical to eat meat. Students were also encouraged to view on their own outside of class an optional eleven-minute vegetarianism advocacy video "What Came Before" (<http://whatcamebefore.com>). Students were informed that the video contains factory farm footage that some people might find upsetting, that it is optional, and that they would not be tested on its contents. Due to the potentially upsetting content of the video, it was important to emphasize its optionality, and this is standard practice among philosophy instructors who use videos of this sort in their teaching. In accord with our own usual teaching procedures for this material, we also thought it would be ethically preferable to use an opt-in rather than an opt-out structure even though this meant that we couldn't control which students viewed the video. We chose "What Came Before" in part because of preliminary research suggesting that it was more effective in encouraging online participants to request further information about vegetarianism than were some other videos with which it was compared (Cooney, 2013).

Students in the ethics of charity condition were required to read Peter Singer's article "The Singer Solution to World Poverty" (Singer, 1999/2007), a six-page opinion article that had originally appeared in the *New York Times Magazine*. In this article, Singer argues that it is unethical to spend hundreds of dollars on luxuries instead of giving that money to an effective charity that could save the life of someone in poverty. Students were also encouraged to view on their own outside of class an optional nineteen-minute video advocating donation to hunger-relief initiatives, "Ending Hunger Now" (https://www.ted.com/talks/josette_sheeran_ending_hunger_now). Students were informed that the

video contains footage of suffering children that some people might find upsetting, that it is optional, and that they would not be tested on its contents.

In neither condition did we assign “con” readings, arguing that meat eating is permissible or that people have no obligation to donate to charity. TAs presented both the pro and con considerations in their discussion sections, encouraging debate among the students, as is typical instructional procedure in philosophy discussion sections in the U.S. It is not unusual in philosophy instruction to assign a reading on only one side of a debate, especially if the reading expresses a philosophical position to which a majority of students are likely to be opposed, and then open up discussion of the pros and cons orally.

Both the Rachels and the Singer articles are commonly used in ethics classes (for example, in the authors’ own non-experimental teaching), and it is not uncommon for philosophy instructors to link to optional online advocacy videos for students who are interested in exploring issues further. These teaching materials are not unusual in the context of university-level philosophy ethics instruction.

Questionnaire

Eric Schwitzgebel emailed students in all courses, introducing himself as a U.C. Riverside philosophy professor interested in students’ attitudes about ethical issues, contacting them with the permission of their instructor. (This language was modified for Philosophy 5, where Schwitzgebel was the instructor.) Students were told they could participate in a short questionnaire on four ethical issues for a small amount of extra credit in the course.

To help ensure confidence in anonymity and reduce demand, the first page of the questionnaire contained the following language: “This study is being conducted by Eric Schwitzgebel and his collaborators at other universities. Your TA *will not be told your answers to these questions*. All identifying information will be stripped from them by an independent coder before Professor Schwitzgebel views the answers, so that he will not know how any particular student has answered.

You will not be graded on your particular answers, and you should feel free to disagree with your professor and TA about the ethical issues at hand.” Recruitment emails contained similar assurances.

The main body of the questionnaire consisted of three questions each on four topics, always in the same order: sexual intercourse outside of a committed, loving relationship; eating the meat of factory-farmed animals; spending a large amount of money on luxuries; and downloading music in violation of copyright laws. On each of the four topics, students were asked, again always in the same order, whether the behavior was unethical, whether they planned to avoid it, and whether if they engaged in that behavior they should feel guilty. All responses to this part of the questionnaire were on a seven-point scale from “strongly agree” (+3) to “strongly disagree” (-3).

The three meat ethics questions were:

4. Eating the meat of factory farmed animals is unethical.
5. I plan to choose non-factory farmed or vegetarian foods when they are available.
6. If I eat factory farmed animals, I should feel guilty about that.

Each set of three questions appeared on a new page, without a back button. The survey concluded by asking students whether they had watched the optional video “What Came Before” (in the meat ethics condition) / “Ending Hunger Now” (in the ethics of charity condition).

Vouchers

Within a few days after the section meetings, UCR’s Dining Services mailed a letter to students’ physical campus addresses, telling them that ten dollars had been added to their Student ID card for use at one of the campus restaurants, “The Barn”. This letter encouraged students to try the new upgraded menu at The Barn and provide feedback about their dining experience. Students who provided feedback were also entered into a drawing for \$150. After students had received a hard copy of the voucher, the head nutritionist at Dining Services followed up with email messages encouraging students to use their

free \$10 at The Barn and informing them that even if they had misplaced the hard copy of the voucher they could still claim the credit by showing their student ID. In Spring 2017, these communications were not visibly connected to their philosophy course in any way. In the fall, in an attempt to increase redemption rates, the head nutritionist from Dining Services appeared in the lecture hall to announce the availability of the vouchers and distributed them by hand as well as following up by email.

We were able to trace students' meal purchases with their \$10. Students' names were replaced with unique identifiers for their privacy. In accord with our pre-registration, each transaction was originally coded as red meat (beef, pork, lamb: coded as 5), poultry (chicken, turkey: 4), seafood (fish, shrimp: 3), vegetarian but not vegan (2), or vegan (1), with the entire transaction coded based on the highest-number menu item (e.g., a bacon turkey burger, salad, and lemonade would be coded as 5 because of the bacon). Although it is not clear whether eating red meat is ethically worse than eating poultry or fish (Singer, 2019), we chose this scoring system due to our impression that students in California tend to think of avoiding red meat, or avoiding both red meat and poultry, as types of vegetarianism or steps along the path toward vegetarianism, an opinion often expressed in popular treatments of vegetarianism (e.g., Berkeley Wellness, 2017; Miller, 2019). As described in the preregistration, we planned to analyze the data both parametrically in terms of mean ranks (1-5) and categorically (vegan or vegetarian vs. seafood, poultry, or red meat). However, parametric analysis yielded virtually identical results to the categorical analysis without adding much new information. Since it is also more difficult to interpret, we have confined those analyses to the supplementary online materials.

The Barn had recently upgraded its menu to highlight attractive vegan choices, which were clearly marked on the menu and advertised on posters displayed throughout the restaurant – a separate

initiative by UCR's head nutritionist, unconnected with this study, but part of the basis for our choosing The Barn as our restaurant location.

Dining Card Data

About one-third of UCR students use their student ID cards for on-campus purchases. After a lengthy approval process, the campus vendor Oracle provided us with every student ID card purchase from April 17 to June 7, 2017, and from October 6, to December 29, 2017, for the campus dining locations they serviced, including date, time to the second, dining location and register number, purchase amount, and for some locations text listing the purchased items. Separately, UCR's Dining Services provided data from Blackboard (a separate service that tracks student-ID related information) for Student ID card purchases, listing Student ID, date, time to the second, purchase location, and purchase amount. For privacy, Student IDs were converted to unique identifiers using the same encryption code as was used in the questionnaire portion of the study. Finally, the Student ID information from Blackboard needed to be matched to item purchase text from Oracle. Both datasets contained information on purchase date, time to the second, purchase amount, and purchase location, enabling a combined output of student identifier (from Blackboard), purchase date, time, and amount (from both data sets), and purchased items (from Oracle). Due to clock inaccuracies, false negatives were extensive for some locations when match to the exact second was required. These false negatives were then compensated for by expanding the time window for a purchase price match at those locations, but doing so created a risk of false positives. (Clock inaccuracies were not consistent, even within location, so a single consistent offset could not be used.) The programmer adjusted the time windows by location in an attempt to balance false negatives against false positives. Post-hoc checks based on receipt images and plausibility considerations (e.g., its being unlikely that a register would have two different same-priced transactions within five seconds of each other) suggest final false positive rates of less than 2%

and false negative rates of less than 5% – a small amount of noise that contributes a bit more to our unsystematic sampling of purchases and could potentially reduce the ability of the study to capture an effect, but which is unlikely to be related to the variables of interest. All of this was completed before hypothesis testing, and the programmer was not informed which transactions belonged to students in which conditions.

The relevant discussion sections were held during the weeks of May 15 (Philosophy 1, Spring), June 5 (Philosophy 2, Spring), October 30 (Philosophy 5, Fall), and November 27 (Philosophy 1, Fall). Because the Spring Philosophy 2 discussion sections were held close to the end of the available Spring purchase data, students in that class were not included in the dining card data analysis.

Restaurant venues that did not list the items purchased were excluded from analysis. Each transaction at the remaining venues that had been matched to a student by means of the unique identifier was then coded as red meat (3), poultry (2), seafood (1), or vegetarian (0), with the entire transaction coded based on the highest value menu item. Because we could not reliably distinguish vegan from vegetarian choices, we used a four category scale (0-3) instead of the five category (1-5) scale we used for the vouchers, setting the floor at 0 rather than 1 to reduce the false appearance of numerical comparability. Again, we planned both parametric and categorical analyses, but for the reasons described above in the Vouchers section, the parametric analyses are confined to the online supplement. Given the number of transactions, it was not possible to hand-code every transaction, so transactions were coded based on search terms such as “chicken”, “hamburger”, and “shrimp”. The entire list of purchases was alphabetized and skimmed to see what items commonly appeared and what abbreviations were commonly used (e.g., “ckn” for chicken). After a first-pass coding, we spot-checked for miscoded transactions (e.g., a vegetarian “garden burger” miscoded as 3 due to containing “burger”), then added new search terms until spot checking suggested that our coding accuracy was over 99%.

Each student's discussion meeting date was noted, and transactions on that date were coded as date = 0. Transactions from the day before and after were coded as date = -1 and +1 respectively, from two days before and after as date = -2 and +2, etc. For analysis, dates of -1 or lower were classified as "before" the discussion meeting, and dates of 0 or higher were classified as "after". (The day of the discussion meeting was classified as "after" rather than "before" due to the expectation that a typical student may have read or skimmed the assigned reading the night before the scheduled section meeting. As it happened, only 1.5% of purchases were on Day 0, and the results are essentially the same if a slightly earlier or later date is used as the cutoff or if the day of intervention is excluded from analysis.) Before receiving any data, based on examination of menu prices at campus locations, we had hypothesized that the most informative analysis would be limited to purchases of \$4.99 or more, eliminating small snack and drink purchases, which even among regular meat eaters are likely to be vegetarian and would thus potentially add noise to the analysis. Examination of price information after receipt of the data but before analysis seemed to confirm \$4.99 as a reasonable cutoff price (the price of a cheeseburger, chicken pita sandwich, or vegetarian sushi roll). However, for completeness and to help mitigate potential concerns about p-hacking, we planned to run our main analyses both on the full data set and on the set of purchases of \$4.99 or more.

Consent, Deception, and Privacy

This study raises ethical concerns regarding consent, deception, and privacy that we wish to acknowledge here. Students did not consent to be sorted into sections on the topics of the ethics of eating meat or donating to charity. However, both topics are often covered in philosophy instruction, and we used materials that are normally used in ordinary, non-experimental philosophy instruction. It is not unusual in philosophy courses to have one's ethics challenged and to read and view potentially upsetting materials. The most potentially upsetting material, the films, were flagged as such and marked

as optional, and students were told they would not be tested on the material. Having different sections taught with different pedagogical techniques is within normal philosophical teaching practice, and we emphasized to the TAs that their duty was just their ordinary pedagogical duty: They were to teach the assigned topics as they would teach any other assigned topics, in accord with their usual teaching standards and techniques, with student learning as their overriding goal.

The voucher portion of the study involved deception by omission, since we did not reveal that we had funded the vouchers and that students were receiving the vouchers as part of an experiment. Nor were students informed that their Dining Card purchases were subsequently examined. There was no deception by commission, since The Barn did have a new menu, did collect student feedback that they considered valuable, and did distribute \$150 from their own funds to the winner of the drawing. Students were not debriefed about the deception because the benefits of debriefing were judged not to outweigh the costs and risks. The costs would have included inability to conduct follow-up studies and the risks included the possibility of students falsely inferring violations of privacy. In all portions of the study, student privacy was protected by replacing student names with unique identifiers so that no individual's purchases or questionnaire responses could be known. We also note that meal purchases in restaurants don't normally come with a high expectation of privacy, since the food purchase is known to the cashier, visible to passersby, and recorded by the company that manages the card transactions. It is not unusual for companies to study consumer choice under various conditions without explicit consent.

All aspects of the design were pre-approved by UCR's Institutional Review Board (IRB-SB-17-010).

Results

Attendance and Exclusions

Of the 1332 originally enrolled students, 189 were excluded from the study on grounds of having been absent on the day of the discussion section (174), having taken an earlier class in this same study (10), or because it was unclear which discussion section they had attended (5), leaving 1143 students for analysis: 325 in Spring Philosophy 1 (198 meat ethics), 240 in Spring Philosophy 2 (112 meat ethics), 285 in Fall Philosophy 1 (178 meat ethics), and 293 in Fall Philosophy 5 (114 meat ethics).

Questionnaire

1036 students completed the questionnaire, of whom four were excluded for completing in under sixty seconds, yielding an overall response rate of 1032/1143 (90% overall, with very similar participation rates in all four classes, final $N = 539$ meat ethics, 493 charitable giving).

As expected, the meat ethics and charity ethics groups did not detectably differ in their answers to the six questions concerning the ethics of sex and the ethics of copyright ($|t| \leq 1.3$, $p \geq .19$). Since the assigned reading on the ethics of charity condemned luxurious spending, we anticipated that students in those sections might be more condemnatory of luxurious spending. However, we did not find such an effect. Indeed, respondents in the charity ethics condition expressed significantly *less* agreement that they planned to avoid spending large amounts of money on luxuries ($M = -0.33$) than did respondents in the meat ethics condition ($M = 0.04$, pooled $SD = 1.48$, $t(1030) = -4.01$, $p < .001$; $d = -0.25$). It is possible that this is a “backfire” effect: Students in the ethics of charity condition might have reacted negatively to having their luxurious spending ethically challenged. Alternatively, it is possible that students in the ethics of charity condition had a more vivid appreciation of how much would be involved in sacrificing one’s luxuries or that they had a broader idea of what would count as a “luxury”, given Singer’s characterization of “luxury” as including things like spending \$200 a month on restaurant meals. The “unethical” and “guilty” luxury questions were non-significantly in the direction of a backfire effect (unethical: $M_{meat} = 0.09$, $M_{char} = -0.09$, pooled $SD = 1.55$, $t(1030) = -1.86$, $p = .063$, $d = -$

0.12; guilty: $M_{meat} = -0.64$, $M_{char} = -0.74$, pooled $SD = 1.71$, $t(1030) = -0.95$, $p = .35$, $d = -0.06$). The large SDs on this issue and the meat ethics issue reflect, as expected, substantial differences of opinion on these matters of ethical controversy in the target population.

As expected, students in the meat ethics condition were more condemnatory of meat eating and expressed more agreement that they would avoid eating meat than were students in the charity ethics control condition, as shown in Table 1.

Table 1

Mean agreement (+3 to -3 agree/disagree scale) and percentage agreement (“slightly agree” (+1) or higher) with three claims about meat ethics

Question	Control Condition			Meat Ethics Students			Test Statistic		
	<i>M</i>	<i>SD</i>	% agr	<i>M</i>	<i>SD</i>	% agr	<i>t</i> (<i>DF</i>)	<i>p</i>	<i>d</i>
“Eating the meat of factory farmed animals is unethical.”	-0.46	1.69	29%	+0.12	1.65	43%	5.57(1030)	< .001	0.35
“I plan to choose non-factory farmed or vegetarian foods when they are available.”	-0.25	1.80	37%	+0.03	1.82	42%	2.54(1030)	.011	0.16
“If I eat factory farmed animals, I should feel guilty about that.”	-0.74	1.71	25%	-0.24	1.71	37%	4.63(1030)	< .001	0.29

To test for differences by class and TA we created a composite meat ethics score by averaging the three meat ethics responses. The four classes did not differ statistically in their composite score, ranging from a mean of -0.37 to -0.16 (ANOVA, $F(3, 1028) = 0.97$, $p = .41$, $\eta^2 = .00$). However, the teaching assistant might have made a difference: Among TAs teaching the meat ethics material, the composite meat ethics score ran from a mean of -0.65 to a mean of +0.28 (ANOVA, $F(8,530) = 2.17$, $p = .03$; $\eta^2 = .03$). Among the students of the TA whose sections disagreed most with the conclusions of the assigned reading from Rachels, only 27% agreed that eating the meat of factory farmed animals is unethical (a similar rate to the rate in the control condition), compared to 59% agreement among students of the TA whose sections agreed most with the conclusions of the assigned reading ($\chi^2(8) = 24.6$, $p = .002$; $\phi = .21$).

In the meat ethics condition, 33% (179/539) of students reported having watched the optional video, and another 25% (134) reported that they started watching the video but did not finish it. In the ethics of charity control condition, somewhat fewer students watched the video: 25% reported watching the whole video and 20% reported starting to watch but not finishing. (The charity video was longer and perhaps less interesting.) Students who reported watching the entire meat ethics video were more likely to agree that eating the meat of factory farmed animals is unethical (54%) than were students who reported not watching the video (38%) or only watching part of it (40%) ($\chi^2(2) = 11.5, p = .003; \phi = .15$). It is unclear whether the videos influenced student opinion or whether students who already tended to condemn meat eating were more likely to watch the meat ethics video. Nevertheless, even among students who did not watch the videos, those in the meat ethics condition expressed more agreement that eating the meat of factory farms is unethical than those in the control condition (38% vs. 27%, two-proportion z test: $z = 2.60, p = .009, OR = 1.65; M_{meateth} = -0.07, M_{charity} = -0.52, \text{pooled } SD = 1.65, t(497) = 3.04, p = .003, d = 0.27$).

Vouchers

Redemption rates for the vouchers were unexpectedly low. In the Spring term, only 143/565 (25%) of the vouchers were redeemed. The more assertive distribution technique of Fall term raised the redemption rate to 230/578 (40%), for a total overall redemption rate of 33% (373/1143), still slightly below our target N of 200 per group. One possible explanation is that The Barn was not as centrally located on campus or as familiar to students as some of the other dining options.

Despite the presence of well-advertised and seemingly attractive vegan options at The Barn, only 18/185 students (10%) in the meat ethics condition and 18/188 students (10%) in the charity ethics condition used the vouchers to purchase either a vegetarian or a vegan item (two-proportion z test: $z = 0.05, p = .96; 95\% \text{ CI of difference} = -5.8\% \text{ to } +6.1\%$).

Dining Card Data

By purchase. We matched 13,642 card purchases to 495 students who used their ID card for food purchase on campus at least once at one of the included locations: 4926 transactions by 162 students from Philosophy 1 Spring, 5124 transactions by 183 students from Philosophy 1 Fall, and 3592 transactions by 150 students in Philosophy 5 Fall (median 26 transactions per student, maximum 116). Of these, 5981 were transactions of at least \$4.99 (2042, 2334, and 1605 in the three classes, respectively). We included transactions from the entire range of data, spring as well as fall, for all students. Since the majority of students participated in discussion sessions between November 3 and December 1, we had more pre-intervention transactions (9759 total, 4326 of at least \$4.99) than post-intervention transactions (3883 total, 1655 of at least \$4.99). Also, some students might expend their allotted dining funds before the end of the term.

Table 2 shows the percentage meat purchases for both the control group who discussed the ethics of charity and for the experimental group who discussed the ethics of eating meat, before and after the section discussion meeting. Meat purchases declined in the meat ethics group but not in the ethics of charity group – a result that holds true whether all purchases are included or only purchases of \$4.99. As is evident from Table 2, before intervention, the groups were virtually identical in their percentages of meat purchases. For example, looking at purchases of \$4.99 or more, before the intervention, 52% of the purchases of both groups contained meat. The control group's rate of meat purchase remained at 52% after the intervention, while the meat ethics group's rate of meat purchase fell to 45%. This interaction effect is confirmed by a logistic regression on purchases of \$4.99 or more, predicting meat purchase (vegetarian = 0, meat = 1) from condition (control = 0, meat ethics = 1), time (before intervention = 0, after = 1), and the interaction variable condition*time (no other predictors were used).

In accord with the hypothesis that instruction influenced food choice, the interaction variable was the only statistically significant predictor: $b = -0.27$, $p = .022$).

Table 2

Percent meat purchases as measured from dining card receipts, before and after discussion of either meat ethics or ethics of charity

Condition	% Meat Purchases		Test Statistics		95% CI for diff	odds ratio
	before	after	<i>z</i>	<i>p</i>		
Control group: charity ethics						
all purchases (N = 6333)	27.9%	28.7%	-0.59	.55	-1.8% to +3.3%	1.04
\$4.99 or more (N = 2775)	51.6%	51.9%	-0.14	.89	-3.9% to +4.5%	1.01
Treatment group: meat ethics						
all purchases (N = 7309)	28.1%	24.8%	2.89	.004	-5.4% to -1.1%	0.85
\$4.99 or more (N = 3206)	51.6%	45.3%	3.27	.001	-10.2% to -2.6%	0.78

As is evident from Table 2, only about 50% of purchases over \$4.99 contained meat, compared to 90% of purchases in the Voucher portion of the experiment. Post-hoc examination reveals that these purchases are a mix of snack items, convenience store items, drinks, and vegetarian meals, not always clearly distinguishable from each other – for example, a random selection of seven vegetarian purchases over \$4.99 shows (1) a veggie delight sandwich with chips; (2) a banana, milk, and eggs; (3) several water bottles; (4) a chocolate bar, Cheetos, and mango slices; (5) a cheese quesadilla and fries; (6) several flavored milks; (7) a flavored café leche. Although the full dataset is proprietary information we have promised not to share, this is approximately representative. It is possible that meat is sometimes replaced by increased snacking rather than by full meals centered on a vegetarian entree.

By participant. We also examined the data participant-by-participant. This allowed us to associate behavior change with expressed attitude and also to assess whether the purchase-by-purchase results may have been driven by just a few participants. Of the 495 participants who had at least one recorded purchase, 464 had at least one purchase of at least \$4.99 before the day of the discussion section meeting, and 289 had at least one purchase of at least \$4.99 after the day of the discussion

section meeting. As noted above, the intervention was late in the date range for the majority of students. Overall, 277 participants (163 in meat ethics, 114 in ethics of charity) had purchases of at least \$4.99 both before and after the intervention.

Among the 114 participants in the ethics of charity control condition with recorded purchases of at least \$4.99 both before and after instruction, the average percentage of meat purchases before instruction was 53.3%, compared to 53.1% after. Among the 163 in the meat ethics condition, the average percentage was 55.9% before compared to 44.8% after. To examine this difference statistically, we employed a multilevel logistic regression including all participants, predicting whether the purchase contained meat (0 = vegetarian, 1 = meat) from a variable that was 1 if the purchase was made after instruction by a student who received the meat ethics instruction and 0 otherwise. Participant ID number was treated as a random effect and no other predictors were used. Purchases after instruction by meat ethics students were substantially less likely to contain meat than other purchases, with an estimated odds ratio of 0.69 (95% CI [0.57, 0.84], $b = -0.37$, $p < .001$). If purchases under \$4.99 are also included, the difference is smaller but still statistically significant, with an estimated odds ratio of 0.82 (95% CI [0.72, 0.94], $b = -0.20$, $p = .004$). Adding dummy predictors for condition (control = 0, meat ethics = 1) and time (before intervention = 0, after = 1), substantially expands the confidence intervals but does not much change the estimated effect sizes (\$4.99+: $OR = 0.72$, 95% CI [0.54, 0.97], $p = .029$; all purchases: $OR = 0.82$, 95% CI [0.66, 1.01], $p = .057$; dummy variable odds ratios 0.96-1.01).

We were curious whether we could find any evidence of vegetarianism among our participants. One hundred and thirty-three students had at least five purchases of \$4.99 or more after the day of the discussion meeting. Of these, 38 had fewer than 20% meat purchases after the discussion. We examined every transaction by these 38 students. Contrary to the expectation that we would find evidence of several practicing or newly converted vegetarians, almost all students either made

occasional meat purchases after the intervention or primarily purchased expensive drinks or snack items, which would not constitute clear evidence of vegetarianism (e.g., a habit of expensive frappuccinos). Only two students appeared to be unambiguous vegetarians, both in the meat ethics condition. One student had only vegetarian purchases throughout the recorded data. The other made several meat purchases before the day of the intervention, but afterwards made only vegetarian purchases, including some purchases that appear to have been full vegetarian meals. We infer, somewhat surprisingly to us, that the decline in meat purchases among the meat ethics group reflects a broad-based moderate reduction in meat purchases rather than the conversion of several students to vegetarianism.

Time course. Possibly, exposure to meat ethics has a brief effect (a day, a week, two weeks) which later disappears. We had limited statistical power to test this hypothesis, given the relatively thin rate of dining card purchases of at least \$4.99 at the target campus restaurants. Contrary to our expectations, we found no evidence that the effect decreased over time. For example, looking only at the data before intervention compared with the data in the two weeks after intervention slightly reduces, rather than increases, the estimated effect size (and the effect no longer crosses the threshold of statistical significance): Among students in the meat ethics condition, in the data before the intervention, 52% of purchases of at least \$4.99 were meat, compared to 47% in the two weeks after (1175/2276 vs. 228/487, two-proportion z test, $z = 1.93$, $p = .054$, $OR = 0.82$). (Recall from Table 2 that, if all data are included, 45% of these students' purchases after intervention contained meat.) Similarly, in a point biserial correlation analysis looking at all purchases of \$4.99 or more in the after-intervention meat ethics group, we found no relationship between time after intervention in days and choice of meat or non-meat options (meat = 1, vegetarian = 0, $r_{pb} = -.05$, $p = .14$, where a positive slope would indicate a tendency to increase meat purchases over time). Since the majority of students received the intervention

November 3 to December 1 and the data ended on December 29 (growing sparser in the second half of December), we estimate that the effect endures for at least several weeks.

Relationship with expressed attitudes. Unsurprisingly, expressed attitudes toward eating meat were correlated with purchase behavior. However, the relationship was smaller than one might expect. Agreement (-3 to +3) that eating factory farmed meat is unethical correlated at only $r = -.13$ ($p = .030$) with percentage of meat purchases after intervention (among purchases of at least \$4.99). The correlation was similar and only marginally significant for the guilt question $r = -.10$ ($p = .085$). Even agreement with “I plan to choose non-factory farmed or vegetarian foods when they are available” was only modestly correlated with observed purchase behavior after intervention: $r = -.22$ ($p < .001$).

Discussion

As far as we are aware, this is the first controlled study to show an effect of university-level ethics instruction, as conducted in ordinary philosophy classes, on non-laboratory behavior, using direct observational data rather than self-report. We found that after exposure to a philosophy article, a fifty-minute philosophy discussion section, and an optional online video concerning the ethics of eating factory farmed meat, students decreased their rates of meat purchasing from 52% to 45% of their food purchases of \$4.99 or more in campus dining locations for which receipts were available, compared to a constant rate of 52% among students in a control group exposed to similar materials on the ethics of charitable giving. This effect appears to have been a widespread moderate reduction of meat purchases among students rather than the conversion of several students to strict vegetarianism. Although we had only limited ability to detect the time course of the effect, we did not observe a decrease in effect size among students for whom we had several weeks or several months of data. The effect size is in our judgment striking given the brevity of the intervention and the fact that most university students are likely to have been previously exposed to arguments for and against vegetarianism.

Expressed moral opinion also changed substantially. In the ethics of charity control condition, 29% of participants agreed that “eating the meat of factory farmed animals is unethical”, compared to 43% in the ethics of eating meat condition. Again, the magnitude of the effect is in our view striking, given that most participants were probably already familiar with some arguments for vegetarianism. Although concerns about experimenter demand can be raised for any such study, we note that responses were rigorously anonymized and collected by an experimenter who was, for three of the four classes, not among the students’ instructors, and we saw no such demand effect in responses to the questions about the ethics of charitable giving in the control condition.

We note the following limitations:

First, although assignment to the treatment or control group was quasi-random for the majority of participants in the Spring term, assignment was more haphazard in the Fall term, and in no class were we able to completely balance the two conditions, due to both practical and ethical limitations on our ability to control student enrollment. Thus, although students could not have known in advance which sections would discuss which topic, it is possible that students in some of the discussion sections on meat ethics were systematically different from students in some of the discussion sections on the ethics of charity.

Second, we obtained purchase data for only the minority of students who used their student ID card for purchases, and only for the minority of campus locations that provided detailed receipts through the Oracle corporation, giving us a limited number of temporally scattered observations. Although we have no reason to think that use of Student ID or not, or food purchase at the target locations versus other locations, would be systematically related to any of the variables of interest, it remains possible that some unobserved correlate partly or entirely explains the observed effects. For example, if the most attractive vegetarian options were at a food locations not included in this study, vegetarian students

might be migrating out of the data pool and we might be underestimating the effect. Alternatively, if the meat ethics sections that happened to be conducted near mealtimes were also disproportionately held near restaurants with more attractive vegetarian options, we might be overestimating the effect for that subset of students by up to one meal a week.

Third, the voucher portion of the study, which would have provided a more systematic sample of purchase behavior than was available through the receipts, was underpowered due to unexpectedly low redemption rates and for the same reason might reflect non-response bias (systematically differential redemption rates for different types of participants). To improve redemption rates, we altered the design of the study midstream, extending it from Spring to Fall, and we had Dining Services distribute vouchers in the target classrooms. Despite this, power remained low, and we were not able to confirm the effect of instruction on food choice in this portion of the study.

Fourth, the relationship between expressed opinion and observed behavior was smaller than we anticipated (r 's ranging from $-.10$ to $-.22$), raising questions about the accuracy of the opinion measure or the role of change of opinion in mediating change of behavior.

Fifth, the readings, film, and questionnaire focused specifically on the ethics of eating “factory farmed” meat. Although we are not aware of any meat items at the target restaurants that were advertised as “organic,” “grass-fed” or “humanely raised,” the blanket category of “non-vegetarian” obscures differences in the conditions under which farmed animals are typically raised, and the extent to which the label “factory farmed” is appropriate. Although in the U.S. pig and chicken products are overwhelmingly factory farmed throughout their lives, cattle are raised initially on grass, and then in feedlots that, though intensive, are outdoors, while lamb is usually not factory farmed. We do not know what effects better information about such matters would have had or whether students would have chosen humanely raised meat were humane options explicitly made available on campus.

Sixth, it remains unclear what aspect of instruction influenced student opinion and behavior. Teaching assistants were encouraged to teach in their usual style and employed a variety of approaches, likely some more effective than others. Although university-level philosophy discussion sections generally aim to focus on relatively unemotional rational evaluation of the pros and cons of arguments, students might also have been emotionally influenced by the content of the reading, the discussion, or (in some cases) the optional film. Also, if some TAs or fellow students discussed their own vegetarian opinions or experiences, that might have made vegetarianism seem more achievable or socially normal to students in the meat ethics condition than in the control condition. The TAs presumably differed in teaching style in ways that we did not attempt to measure, including in their degree of personal openness about their own behavioral choices, in their evenhandedness in presenting both the pro and con arguments, in their solicitation of debate from students, and in the level of abstraction of their examples – any or all of which might have had an influence on the results. The optional film might have played an important role for those students who chose to view it. It is also possible that merely mentioning facts about meat production that are normally not salient or explicitly discussed might increase cognitive dissonance, motivating behavior change without need of explicit philosophical argument (Bastian & Loughnan, 2017).

We hypothesize that instruction may lead to changes in beliefs, changes in implementation intentions for previously existing beliefs, or changes in spontaneous responses independent of changes in belief or implementation intentions. The treatment group's higher agreement with "eating the meat of factory farmed animals is unethical" is evidence of change in belief, but as we have noted, it is only weakly related to actual food choice. Furthermore, experimenter demand and social desirability might influence responses. Implementation intentions might change even if belief remains the same: A student who before instruction regularly ate meat despite believing that eating the meat of factory farmed

animals is unethical might after exposure to instruction commit to specific steps to reduce meat consumption (e.g., “tomorrow at lunch I will choose the vegetarian alternative”). Likewise, a student who before instruction does not believe that eating meat is unethical might not have their belief changed by instruction but nonetheless form an intention to reduce meat consumption for personal or aesthetic reasons. Finally, neither a student’s belief nor implementation intentions might be changed by instruction and yet the student might reduce meat consumption due to changed spontaneous responses to food choices. After instruction they might, for example, have a less positive implicit attitude toward ham and thus be less likely to choose it when given attractive alternatives. The weak relationship between expressed opinion and food choice ($r = -.13$) suggests that change in belief might not be the sole cause of change in behavior. The only slightly stronger relationship between agreement that “I plan to choose” vegetarian or non-factory farmed meat ($r = -.22$) suggests that implementation intentions might also not be fully explanatory (though there are, for example, “reducitarian” implementation intentions that would not have been captured by that question).

We hypothesize that instruction might have had its effect on belief and behavior through three partly separable channels: (1) rational argument or the presentation of facts, (2) emotional appeal, and (3) social influence. Rachels’ (2004) “basic argument” for vegetarianism is that the modern meat-production business involves substantial animal suffering and that our enjoyment of the way meat tastes is not good enough reason to justify the amount of suffering. His article provides a detailed but nontechnical exposition of that argument and a defense of that argument against some common objections. Students might have become convinced on rational grounds that this argument is sound, that objections to it fail, that no excusing conditions apply in their own case, and consequently that their eating factory farmed meat is unethical. Change in belief through rational argumentation is one way that philosophers have traditionally sought to change behavior, and it might in some cases be

psychologically effective (May, 2018). Emotional appeal, while not strictly separable from rational argument (arguments might sometimes legitimately appeal to emotion), might also help drive the effect: Students might, for example, feel disgust, guilt, anger, or sadness when reading about or viewing video footage of animal suffering, and those emotions might influence their implicit or explicit attitudes toward meat consumption independently of their assessment of the force of the rational arguments. Studies of charitable giving, for example, suggest that invoking positive or negative emotions might substantially influence donation amounts (Bagozzi and Moore 1994; Erlandsson, Nilsson, Västfjäll 2018). Moral behavior is also likely to be influenced by social expectations. People might conform to norms in part based on the empirical expectation that others are doing the same (Schultz, Nolan, Cialdini, Goldstein, & Griskevicius, 2007; Bicchieri, 2017; Schwitzgebel, 2019a), and moral exemplars might be an important source of moral change (Engelen, Thomas, Archer, & van de Ven, 2018; Han, Kim, Jeong, & Cohen, 2017; Schnall, Roper, & Fessler, 2010; Zagzebski, 2017). Even if vegetarianism is a minority view, endorsement of it by one or more peers or exemplars might inspire others who have been considering the possibility to give it a try. If a teaching assistant and other students appear to embrace vegetarian attitudes or practices, this might make the avoidance of meat seem more desirable, achievable, or socially normal.

All three of these mechanisms of change might be culturally or subculturally specific: We chose meat ethics for study in part because of anecdotal reports that philosophical instruction on this topic influences students' behavior (in contrast to what we seem to see on other topics). University students in southern California might be especially prepared by their cultural context for behavior and opinion change on this topic. Another context might produce different results – possibly even “backfire” results (as we saw with opinion about charitable giving) if students aren't prepared for change and the instruction seems heavy-handed. All three of these mechanisms might also be influenced in potentially

unpredictable ways by teaching style. In some contexts, emphasizing one side of an argument might be more convincing than an even-handed presentation. In other contexts, even-handedness might help convince students that all objections have been fairly considered. If the student has a positive attitude toward a vegetarian TA, personal interaction in section meetings might be persuasive, perhaps especially if the TA reveals personal details about their opinion or behavior. On the other hand, if the student has a negative attitude toward a vegetarian TA, personal interaction in section meetings might aggravate “do-gooder derogation” (Minson & Monin, 2012).

These data can be reconciled with Schwitzgebel and Rust’s (2014) noneffects in at least two ways. As Schwitzgebel (2019a) notes, to the extent ethicists’ moral behavior is guided by social conformity with non-ethicist peers, ethicists would not be expected to behave differently than their non-ethicists peers, even as their philosophical expertise grows and their opinions change. In contrast, students’ opinions about peer behavior might change considerably as a result of ethics instruction, with behavior following suit. Alternatively but not incompatibly, Nahmias (2012) has suggested that Schwitzgebel’s null results for ethicists may be compatible with moral behavioral change among philosophy students if professors tend to be settled in their ways, having already undergone, as undergraduates, all the moral change that exposure to philosophy is likely to inspire.

In follow-up research, we hope to better explore the basis of opinion and behavior change by exploring the effects of different teaching styles, collecting a larger pool of data allowing for more powerful time-course analysis, expanding the opinion questions to obtain a fuller picture of the basis and extent of opinion change, and more carefully examining the possible effects of watching or not watching the film. We urge other researchers interested in the practical effects of teaching philosophical ethics in the university curriculum to develop observational measures of ordinary ethical behavior rather than relying exclusively on laboratory behavior or self-report.

Declarations of Interest: none.

Funding: The project was financed by a grant from Princeton's University Center for Human Values and an Academic Senate Grant from University of California at Riverside.

Acknowledgements

This project required many hands. We are thankful to instructors Luca Ferrero, Heinrik Hellwig, and Andreja Novakovic; teaching assistants Taylor Doran, Marie Evanston, Mara Garza, Dan Green, John Green, Tom Hanauer, Rotem Hermann, Greg Hertz, Daisy Laforce, Luke Luttmann, Max McCoy, Deborah Nelson, Carlo Palombo, Jeremy Pober, and Avery Snelson; research assistant Seth Margolis; the staff at Oracle, especially Ryan Kelley and Annie Van Pelt; Dining Services, Information Technology, and Philosophy Department staff at U.C. Riverside, especially Karen Fiorenza, David Henry, Bart Kats, Gerardo Sanchez, Pat Sorokowski, and Cassie Taan; and Karen Xu at U.C.R.'s Statistical Consulting Collaboratory.

References

- Abend, G. (2014). *The moral background*. Princeton, NJ: Princeton University Press.
- Adams, C. J. (1990/2015). *The sexual politics of meat*. New York: Bloomsbury.
- Bagozzi, R. P., & Moore, D. J. (1994). Public service advertisements: Emotions and empathy guide prosocial behavior. *Journal of Marketing*, 58, 56–70.
- Bastian, B., & Loughnan, S. (2017). Resolving the meat-paradox: A motivational account of morally troublesome behavior and its maintenance. *Personality and Social Psychology Review*, 21, 278-299.
- Berkeley Wellness (2017). What type of vegetarian are you? *Berkeley Wellness* (Mar. 28). URL: <https://www.berkeleywellness.com/healthy-eating/food/article/what-type-vegetarian-are-you>.
- Bicchieri, C. (2017). *Norms in the wild*. Oxford: Oxford University Press.
- Camosy, C. (2013). *For love of animals*. Cincinnati, Ohio: Franciscan Media.
- Cooney, N. (2013). Report: Which factory farming video is more effective? Blog post at *Humane League Labs Blog* (Jul. 9). URL: <http://www.humaneleaguelabs.org/blog/2013-07-19-which-factory-farming-video-is-more-effective> [accessed Oct. 22, 2019].
- Cushman, F., Young, L., & Greene, J. D. (2010). Multi-system moral psychology. In J. M. Doris (Ed.), *The moral psychology handbook*. Oxford: Oxford University Press (pp. 47-71).
- DeGrazia, D. (1996). *Taking animals seriously*. Cambridge: Cambridge University Press.
- Dhammika, B. (2019). How I became a vegetarian. URL: <https://www.bhantedhammika.net/to-eat-or-not-to-eat-meat/how-i-became-a-vegetarian> [accessed December 3, 2019].
- Ellemers, N., van der Toorn, J., Paunov, Y., & van Leeuwen, T. (2019). The psychology of morality: A review and analysis of empirical studies published from 1940 through 2017. *Personality and Social Psychology Review*, 23, 332–366.

- Engelen, B., Thomas, A., Archer, A., & van de Ven, N. (2018). Exemplars and nudges: Combining two strategies for moral education. *Journal of Moral Education, 47*, 346-365.
- Feltz, S., & Feltz, A. (2019). The knowledge of animals as food scale. *Human-Animal Interaction Bulletin, 7*, 19-45.
- Goodall, J., McAvoy, G., & Hudson, G. (2005). *Harvest for hope*. New York: Warner Books.
- Haddad, E. H., & Tanzman, J. S. (2003). What do vegetarians in the United States eat? *American Journal of Clinical Nutrition, 78*, 626S-632S.
- Hadot, P. (1995). *Philosophy as a way of life*. Malden, MA: Blackwell.
- Haidt, J. (2012). *The righteous mind*. New York: Random House.
- Han, H., Kim, J., Jeong, C., & Gohen, G. L. (2017). Attainable and relevant moral exemplars are more effective than extraordinary exemplars in promoting voluntary service engagement. *Frontiers in Psychology, 8*, DOI: 10.3389/fpsyg.2017.00283
- Huemer, M. (2019). *Dialogues on ethical vegetarianism*. New York: Routledge.
- Ivanhoe, P.J. (2000). *Confucian moral self-cultivation*. Indianapolis: Hackett.
- Kant, I. (1785/1996). Groundwork of the metaphysics of morals, M. J. Gregor (Trans.). In M. J. Gregor (Ed.), *Practical philosophy*. Cambridge: Cambridge University Press (pp. 37-108).
- Kohlberg, L. (1981). *Essays on moral development, vol. 1*. San Francisco: Harper & Row.
- Korsgaard, C. (2018). *Fellow creatures*. Oxford: Oxford University Press.
- May, J. (2018). *Regard for reason in the moral mind*. Oxford: Oxford University Press.
- Mill, J. S. (1859/2003). *On liberty*. In M. Warnock (Ed.), *Utilitarianism and On liberty, 2nd ed.* Malden, MA: Blackwell.

- Miller, K. (2019). The 7 types of vegetarian diets from lacto-ovo to vegan, explained by a nutritionist. *Women's Health* (Nov. 13). URL: <https://www.womenshealthmag.com/food/a29700788/types-of-vegetarians>
- Minson, J. A., & Monin, B. (2012). Do-gooder derogation: Disparaging morally motivated minorities to defuse anticipated reproach. *Social Psychological and Personality Science*, 3, 200-207.
- Nahmias, E. (2012). Comments on Schwitzgebel & Rust's "The self-reported moral behavior of ethics professors". Oral presentation at the Society of Philosophy and Psychology, Boulder, Colorado.
- Narvaez, D. (2014). *Neurobiology and the development of human morality*. New York: W.W. Norton.
- Nussbaum, M. (1997). *Cultivating humanity*. Cambridge, MA: Harvard.
- Peacocke, J. (2018). Measuring change in diet for animal advocacy. *Humane League Labs Report* E009R01. URL: <https://thehumaneleague.org/article/E009R01-measuring-diet-animal-advocacy> [accessed Oct. 22, 2019].
- PETA [People for the Ethical Treatment of Animals] (2015). *What is animal liberation? Philosopher Peter Singer's groundbreaking work turns 40*. URL: <https://www.peta.org/about-peta/learn-about-peta/ingrid-newkirk/animal-liberation> [accessed December 3, 2019].
- Pollan, M. (2006). *The omnivore's dilemma*. New York: Bloomsbury.
- Rachels, J. (2004). The basic argument for vegetarianism. In S. F. Sapontzis (Ed.), *Food for thought*. Amherst, NY: Prometheus Books (pp. 70-80).
- Regan, T. (1983). *The case for animal rights*. Berkeley: University of California Press.
- Rest, J. R., Narvaez, D., Thoma, S. J. & Bebeau, M. J. (1999). DIT2: Devising and testing a revised instrument of moral judgment. *Journal of Educational Psychology*, 91, 644-659.
- Schnall, S., Roper, J., & Fessler, D. M. T. (2010). Elevation leads to altruistic behavior. *Psychological Science*, 21, 315-320.

- Schönegger, P. & Wagner, J. (2019). The moral behavior of ethics professors: A replication-extension in German-speaking countries. *Philosophical Psychology*, 32, 532-559.
- Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2007). The Constructive, Destructive, and Reconstructive Power of Social Norms. *Psychological Science*, 18, 429-434.
- Schwitzgebel, E. (2007). Human nature and moral education in Mencius, Xunzi, Hobbes, and Rousseau. *History of Philosophy Quarterly*, 24, 147-168.
- Schwitzgebel, E., (2013). *Do ethics classes influence student behavior?* Manuscript. URL: <https://faculty.ucr.edu/~eschwitz/SchwitzAbs/EthicsClasses.htm> (accessed Oct. 22, 2019).
- Schwitzgebel, E. (2019a). Aiming for moral mediocrity. *Res Philosophica*, 96, 347-368.
- Schwitzgebel, E. (2019b). Most U.S. and German ethicists condemn meat-eating (or German philosophers think meat is the wurst). Blog post at *The Splintered Mind* (Mar. 22). URL: <https://schwitzsplinters.blogspot.com/2019/03/most-us-and-german-ethicists-condemn.html> (accessed Oct. 22, 2019).
- Schwitzgebel, E., and Ellis, J. E. (2017). Rationalization in moral and philosophical thought. In J.-F. Bonnefon & B. Trémolière (Eds.), *Moral inferences*. New York: Routledge (pp. 170-190).
- Schwitzgebel, E., & Rust, J. (2014). The moral behavior of ethics professors: Relationships among self-reported behavior, expressed normative attitude, and directly observed behavior. *Philosophical Psychology*, 27, 293-327.
- Schwitzgebel, E., & Rust, J. (2016). The behavior of ethicists. In J. Sytsma & W. Buckwalter (Eds.), *A companion to experimental philosophy*. Malden, MA: Wiley Blackwell (pp. 225-233).
- Scruton, R. (2004). The conscientious carnivore. In S. F. Sapontzis (Ed.), *Food for thought*. Amherst, NY: Prometheus Books (pp. 81-91).

- Singer, P. (1975/2009). *Animal liberation*. New York: Harper Perennial.
- Singer, P. (1999/2007). The Singer solution to world poverty. In J. Rachels & S. Rachels (Eds.), *The right thing to do*. Boston: McGraw Hill (pp. 138-144).
- Singer, P. (2019). The risk of a mass switch from red meat to chicken. *New York Daily News* (January 23, 2019). URL: <https://www.nydailynews.com/opinion/ny-oped-the-risk-of-a-mass-switch-from-red-meat-to-chicken-20190123-story.html> [accessed December 3, 2019].
- Vinnari, M., Montonen, J., Härkänen, T., & Männistö, S. (2009). Identifying vegetarians and their food consumption according to self-identification and operationalized definition in Finland. *Public Health Nutrition*, 12, 481-488.
- Watts, L. L., Medeiros, K.E., Mulhearn, T.J., Steele, L.M., Connelly, S., & Mumford, M.D (2017). Are ethics training programs improving? A meta-analytic review of past and present ethics instruction in the sciences. *Ethics & Behavior*, 27, 351-384.
- Wright, J.C. (in press). Symposium on *The character gap*. *Journal of Philosophical Research*.
- Zagzebski, L. (2017). *Exemplarist moral theory*. Oxford: Oxford University Press.

TEACHING MEAT ETHICS

Online Supplement Containing Parametric Analyses of Meat Scores

As mentioned in the main body of the article, our preregistration initially envisioned both categorical and parametric analyses. As it happened, the parametric and categorical analyses yielded virtually identical results. Since the categorical analyses are simpler and more intuitively interpretable, to keep the main body of the paper less complicated, we report the parametric analyses here.

Vouchers. Voucher purchases at The Barn were given one of five ranks on a “meatiness” scale: red meat (beef, pork, lamb: coded as 5), poultry (chicken, turkey: 4), seafood (fish, shrimp: 3), vegetarian but not vegan (2), or vegan (1), with the entire transaction coded based on the highest-number menu item. Rank analysis is justified in terms of dietary restrictiveness. A vegan diet forbids all animal products, which is stricter than a vegetarian diet which permits milk and eggs but forbids eating the meat of any animals, which is in turn stricter than a “pescatarian” diet which permits eating seafood and animal products such as milk and eggs but forbids all other meats, which is in turn stricter than a policy of only avoiding “red meat”. Distinctions among these types of vegetarian or semi-vegetarian diets are well-known among U.S. residents who are interested in the ethics of meat, as expressed in popular treatments such as Berkeley Wellness (2017) and Miller (2019). Beef, pork, and lamb violate all four of these dietary policies; poultry violates three dietary policies; seafood violates two; cheese violates one; and lentils violate none, justifying the rank ordering. We note, with Singer (2019), that this rank ordering might not be *ethically* justified: Chicken might not actually be an ethically better choice than beef, etc. The question is only whether there is an empirical tendency, given common cultural understandings of the various forms and strengths of vegetarianism or semi-vegetarianism, for exposure to material on meat ethics to move students to purchase foods lower in this rank ordering.

Parametric analysis of mean ranks found no difference between the charity and meat ethics conditions ($M_{charity} = 4.25$, $M_{meateth} = 4.15$, pooled $SD = 1.00$, $t(371) = 0.96$, $p = .34$, 95% CI = -0.10 to +0.30), matching the null result of the categorical analysis. About half of the students made red meat purchases: 42% in the meat ethics condition and 51% in the ethics of charity condition (two-proportion z test: $z = 1.73$, $p = .08$).

Dining Card Data. Dining card data did not permit clear separation between vegan and vegetarian items, so we used a four point rank scale from 0 (vegetarian) to 3 (red meat), with the floor set at 0 instead of 1 to reduce the misleading appearance of comparability with the voucher scale.

The charity ethics control group's purchases had a mean rank of 0.68 before the discussion section meeting vs. 0.72 after (pooled $SD = 1.15$, $t = -1.01$, $DF = 6331$, $p = .31$; $d = 0.03$). Considering only purchases of \$4.99 or more, the mean rank was 1.24 before vs. 1.28 after (pooled $SD = 1.30$, $t = -0.74$, $DF = 2773$, $p = .46$; $d = 0.03$). In contrast, the meat ethics treatment group mean ranks were 0.68 before vs. 0.61 after (pooled $SD = 1.13$, $t = 2.60$, $DF = 7307$, $p = .009$; $d = 0.06$) and, limiting to purchases of \$4.99 or more, 1.24 before vs. 1.09 after (pooled $SD = 1.29$, $t = 2.97$, $DF = 3204$, $p = .003$; $d = 0.12$). The null result for the control group and the positive results for the treatment group thus match what we found in the categorical analyses reported in the main article. The Cohen's d effect sizes reported in this supplement are not straightforwardly comparable with the odds ratios for the categorical data reported in the main article.

The main article employed a logistic regression to confirm main findings that were reported in terms of z tests on proportions. A parallel linear regression predicts meat score (0 to 3) from condition (control = 0, meat ethics = 1), time (before intervention = 0, after = 1), and the interaction variable condition*time. Matching the findings reported in the main body of the article, only the interaction variable is statistically significantly predictive ($b = -0.19$; $p = .011$).

