The Journal of Economics and Politics

Volume 26 | Issue 1 | Article 3

2021

Information cascades in the classroom: the relationship between in-class feedback and course performance

Amanda C. Cook
Bowling Green State University, accook@bgsu.edu

James Bland
University of Toledo, indycashew@yahoo.com

Andrew Meisner
Purdue University

Information cascades in the classroom: the relationship between in-class feedback and course performance

Cover Page Footnote

We would like to acknowledge helpful comments and suggestions from Xinyan Shi, Kelly Blanchard, Stephen Martin, Sandra Maximiano, an anonymous referee, Georg Schaur, the participants in the 2017 Midwest Economic Association meeting, and support from TopHat in facilitating this research

INTRODUCTION

An important part of teaching is assessing student understanding. While this happens a few times each term with formal assessments, in-class assessment provides a tool to determine how students are grasping new concepts, making connections, and processing information. By utilizing in-class assessments, instructors can decide to spend more or less time on a topic to optimize student learning. Technology allows instructors to pose a question and then tabulate student responses instantaneously in large lectures. Asking questions in class has two benefits to instructors: it allows an instructor to assess whether the majority of the students are understanding the material, and also to assess specific misconceptions associated with alluring incorrect answers.

While posing questions mid-class provides information to the professor, it also provides students an opportunity for meta-assessment. The student thinks `I get this (or not)' and has the opportunity to update their belief based on the immediate feedback provided by getting the questions correct or incorrect. If the student's beliefs of understanding are affirmed (e.g. she believes she understands the material and gets the question correct), she may adapt her strategy when studying for exams based upon this updated signal

This signal of understanding may be noisier if a student sees peer responses when she selects her answer. For example, one software package that allows instructors to ask such questions in class is TopHat. In this software instructors have the option to display or withhold student responses, which motivates the question: is students' behavior influenced by the observed choices of their peers? Such herding behaviour is called an *information cascade* (Banerjee, 1992). For example, if a few students pick an incorrect answer early on, might this cause another student, who would otherwise have answered correctly, to choose the incorrect, but popular, response? Alternatively, if all but a few early answers are correct, a student who would have answered this question incorrectly had they not observed their peers' answers may answer it correctly. Finally, is this harmful to their performance on exams? In this paper we will characterize whether information cascades occur when peer responses are updated in real time, and how the instructor's choice of peer feedback affects students' long-run understanding of material as demonstrated on exams.

The presence or absence of information cascades alter students' responses to TopHat questions. When peer responses are revealed, students are significantly more likely to answer a question correctly (incorrectly) when the correct (incorrect) answer is the modal choice. This means that the performance of an individual student on TopHat questions may not accurately reflect the student's private level of understanding of the material, since their answers are influenced by the responses of their peers. When peer answers are revealed on a question, the performance of the class will likely provide a distorted measure of overall classroom understanding.

In previous literature, cascades are defined as occurring when individuals disregard their private signal and follow the public signal. Previous experiments were conducted in a laboratory environment where both public and private signals were observed. In this experiment, we do not observe private signals. Instead, we vary between treatments whether students can observe their peers' answers when they submit their own. In the Revealed treatment, students observe others' answers as they are submitted. In the Not Revealed treatment, students only observe others' answers after all students have submitted their answer. Hence, in the Not Revealed treatment we observe the counterfactual where students could not have participated in an information cascade. We compare the two treatments to reveal how students respond differentially when they are only responding to a private signal versus when they also have a public signal. The difference between the treatments is a revelation about relative signal strength between public and private signals which we use to identify cascades.

We define a cascade as a 20 percentage point difference between the modal response in the Revealed and Not Revealed treatment. To classify a question as a cascade, we must observe both a 20 percentage point difference between the modal answer in the Revealed treatment and that same answer in the Not Revealed treatment and also a 20 percentage point difference between the modal answer in the Not Revealed treatment and that same answer in the Revealed treatment. Together, these criteria identify questions where the distributions of final answers were sufficiently different between the two treatments. When students cascade on the incorrect (correct) answer in the Revealed treatment, we classify the cascade as negative (positive).

Using this definition, approximately half of the questions in the Revealed treatment resulted in information cascades. These were roughly evenly split between negative and positive cascades. We estimate that a student who participates in a negative cascade has a five-percentage point reduction in their final exam score for each negative cascade in which a student participates.

LITERATURE REVIEW

Information Cascades

The effect of herd behavior in the form of an information cascade was first modelled by Banerjee (1992). In this model, agents moving sequentially update their beliefs based on both their own private signal and actions of agents who precede them. Agents are influenced by the actions of previous decision makers, and may act against their private signal. For some realizations of private signals, *more agents would make the right choice if they had not seen the decisions made before them.* In a similar model Bikhchandani et al. (1992) shows that information cascades are more likely to occur sooner when private signals are more informative, and that in larger groups information cascades are more likely to occur.

The incentives required to observe herd behavior apply to our setting as well. Students may significantly change their beliefs about which answer is correct upon seeing the distribution of answers from their peers, and hence their submitted answers may not reflect the answer they believed to be correct before seeing the choices of others. In an experiment testing the Banerjee (1992) model, Anderson and Holt (1997) observe rational information cascades in 41 of the 56 cases for which the theory predicted they would occur, observing both positive and negative cascades (Anderson and Holt refer to these as "cascades" and "reverse cascades", respectively). Furthermore, the results from Bikhchandani et al. (1992) suggest that if students know (or believe) that high-ability types submit their answers early, then information cascades on the correct answer (positive cascades) are more likely to occur. If high ability students answer first, instructors could infer that more students understand the question than actually do. On the other hand, if high ability students are not always the first movers, we would expect cascades on both the correct and incorrect answers, which we observe in our field experiment.

Chamley and Gale (1994) better reflects the endogenous timing issues in TopHat in an investment game. Gul and Lundholm (1995) present a similar model, but in continuous time. Agents who choose to lock in their actions earlier are the ones with the best information about the true state of the world, and hence learn less from the feedback of others. This result suggests that the early answers will come from the higher ability students. Thus peer feedback in TopHat is not a good signal of students' understanding: cascades on the correct answer should be more likely to occur if the high-ability students are the first movers, but then the instructor only has information about how well the high-ability students understand the question, not the whole class. On the other hand, it is not *ex-ante* obvious that the high-ability students will choose to answer the questions faster: they have the best information, but they also have the lowest opportunity cost of obtaining *better* information by spending more time working on the problem. One could view this as a student exerting effort to get a better private signal of the correct answer. Cingl (2012) allows for endogenous effort to be exerted over time to gain a better private signal and finds that information cascades arose in this setting. However, we find evidence of information cascades on incorrect answers, suggesting that not all early-movers are high ability.

Sgroi (2003) extends the Anderson and Holt (1997) experiment by allowing for endogenous timing. In line with the theory in Gul and Lundholm (1995), those with better private information move earlier, and most information cascades that did form were cascades on the correct action. We find cascades on both correct and incorrect answers. Our mixed results are not consistent with these experimental findings. In our classroom setting, we do not observe a similar pattern where those with better private information respond earlier. As a result, it is less surprising that we observe both positive and negative cascades.

While there are numerous experiments which examine peer effects and information cascades in a laboratory setting, to our knowledge, we are the first to do so in a field experiment. This is likely because in lab experiments, one observes all private signals, whereas in our experiment we infer aggregate private signals of students by comparing our two treatments. Related field experiments tend to focus on identifying the presence of peer effects rather than identifying information cascades. For example, Tucker and Zhang (2007) manipulate the order of search results for wedding vendors by popularity (rather than alphabetically). They find shoppers may value a vendor's popularity as a positive signal. Similarly, Bursztyn et al. (2014) offer socially connected individuals a new financial asset. They find that individuals are significantly more likely to purchase the asset when they know that their social counterpart was interested in the asset. Again, this suggests that individuals' decisions are influenced by the decisions of their peers.

Peer Feedback

In addition to providing frequent feedback, it is possible that TopHat impacts exam performance by affecting students' studying behavior. For example, if students selected an incorrect but popular answer in a negative cascade (where they would have answered correctly otherwise), it may convince them that the concept is very difficult or that getting this question wrong is acceptable, as many others did too. They may skip studying the concept altogether, resulting in a lower exam score. The effect of feedback on student performance has been examined in papers such as Bandiera et al. (2015) who find that more frequent feedback has a positive effect on student performance, especially for high ability students. Similarly, Azmat and Iriberri (2010) find that providing students information about their own grade, relative to the grades of others has a positive effect on student performance, and that the effect is especially large for students whose grades are relatively low. Falk and Ichino (2006), Mas and Moretti (2009), and Blanes i Vidal and Nossol (2011) all find positive effects from students observing information about their own performance relative to the performance of their peers. In short, most of the available literature suggests that providing students with information about their performance relative to the performance of their peers will have a positive effect on future outcomes. The exception to this is Eriksson et al. (2009) which finds no effect of feedback on performance under piece-rate or tournament pay in a laboratory experiment. In our experiment in both treatments, students get information about their performance relative to their peers. Therefore, when students see peer answers, they may be influenced by this information when selecting their answer. We find participating in negative cascades predicts lower exam performance than getting the same questions wrong when they did not see peer answers, suggesting a negative impact of feedback on their performance relative to their peers. These results suggest that the literature is incomplete, in that the type and timing of feedback matters to student performance.

EXPERIMENTAL DESIGN

In this field experiment, we analyse data from two sections of Principles of Microeconomics at [Retracted for review] University. The sections are taught by the same instructor and each has about 50 students enrolled. Each meet three times a week, with one section meeting in the morning and one in the early afternoon. About twice a week in class, the instructor posts a multiple-choice question. Over the course of the semester, there are about 40 questions. Students use technology (cell phones, tablets, and laptops) to respond to TopHat questions by picking one of four multiple choice answers presented at the front of the room. Students have between 3 and 8 minutes to complete each question, depending on the difficulty of the question. The questions varied in objective, sometimes reviewing a concept from the previous class meeting and sometimes assessing how well students are absorbing a newly presented idea in the same class session.

In this experiment we have two treatments, Revealed and Not Revealed. In the Revealed treatment, TopHat continuously updates the distribution of submitted peer responses. In the Not Revealed treatment, students only see the distribution of responses at the end of the question period. The difference in the treatments is whether or not students see peer responses as they submit their answers. For each section, about half the questions are in the Revealed treatment, and half were not. If students in the first section participated in the Revealed treatment, then the second section participated in the Not Revealed treatment, and vice versa. For each question, the section that participated in the Revealed treatment was randomly determined. As our subjects would naturally undertake these tasks if they were not being experimented on, and do not know that they are being experimented on, this study best falls into the category of "natural field experiment" in the taxonomy of Harrison and List (2004).

As students respond to a TopHat question, we collect a time stamp, thereby allowing us to determine which peer responses each student saw when they submitted their answer. For example, in the Revealed treatment, if a student submitter her answer 40 seconds after the first response, she would only see the responses of the first responder and any students who answered within 40 seconds of the first responder. **Error! Reference source not found.** shows what a student might have seen for a given question in the Not Revealed treatment, and

Figure 2 shows the equivalent picture for the Revealed treatment.

Figure 1: Not Revealed treatment

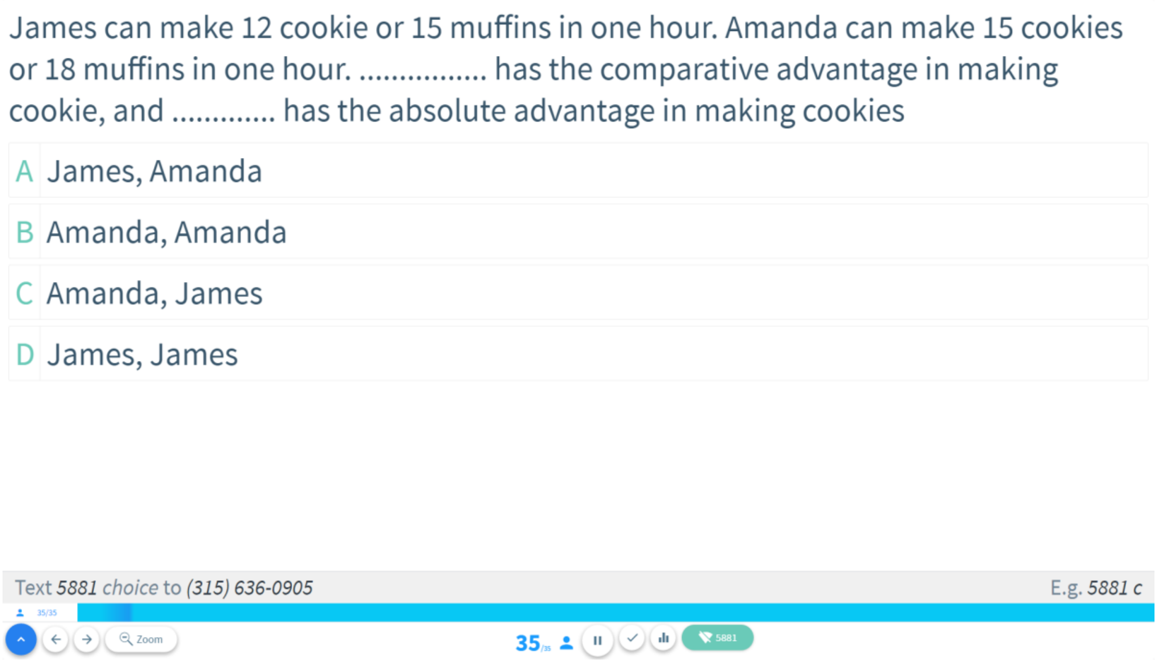

Figure 2: Revealed treatment

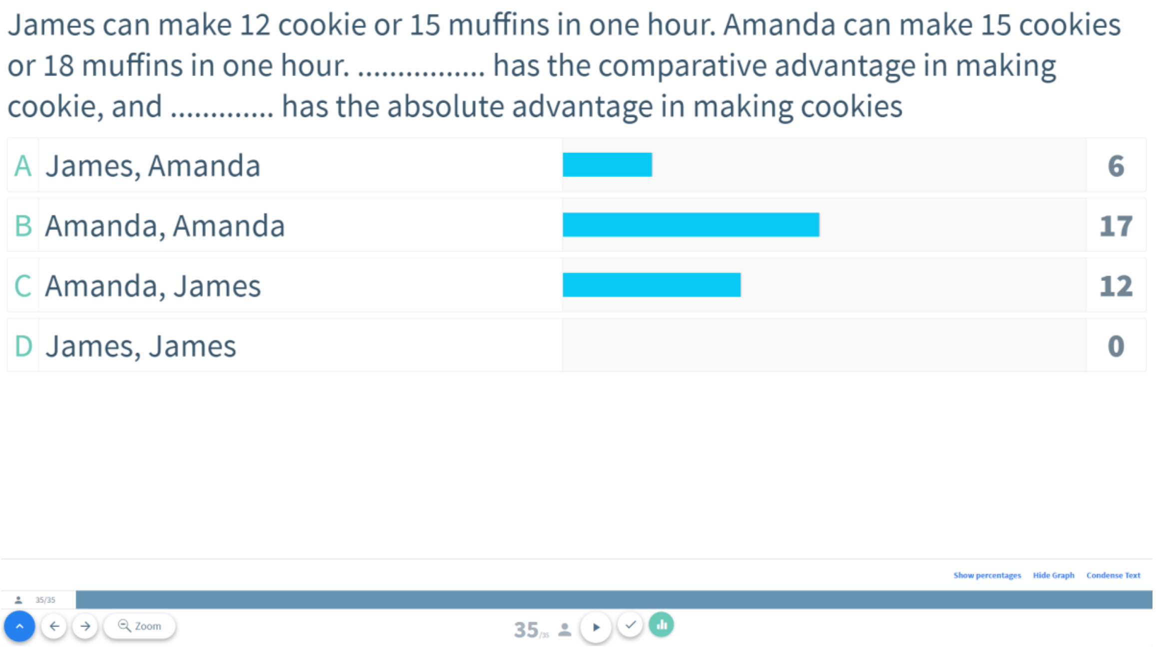

Irrespective of the treatment, students were incentivised with extra credit to participate in TopHat, and were equally incentivised for each correct answer. A student who participated in, and got every TopHat question correct, received approximately a 7-percentage point increase to their final grade. There were 37 questions throughout the term, so a single correct response would produce about a .2 (=7/37) percentage point increase in the student's final grade.

To illustrate how we define a cascade — a 20 percentage point difference between the modal response in the Revealed and Not Revealed treatment — we provide examples of student responses in Table 1 with their appropriate classifications of a positive cascade, a negative cascade, or no cascade.

***Table 1: Hypothetical cascade examples*:**

	Question 1		Question 2	
Student Responses	Revealed	Not Revealed	Revealed	Not Revealed
A	65%	25%	75%	50%
B	20%	25%	10%	25%
C	10%	50%	10%	25%
D	5%	0%	5%	0%
Correct Answer	C		A	
Modal response:	A: 65%	C: 50%	A: 75%	A: 50%
Cascade?	Negative		Positive	

Note: To classify a question as a cascade, we must observe both a 20 percentage point difference between the modal answer in the Revealed treatment and that same answer in the Not Revealed treatment and also a 20 percentage point difference between the modal answer in the Not Revealed treatment and that same answer in the Revealed treatment. Together, these criteria identify questions where the distributions of final answers were sufficiently different between the two treatments. When students cascade on the incorrect (correct) answer in the Revealed treatment, we classify the cascade as negative (positive). For perspective, with approximately 40 students participating in each class, the probability of observing an information cascade if all students selected answers randomly is about 2% when we use a 20 percentage point cut off rule.

For example, in Table 1, if for Question 1 in the Revealed section, 65% of the students picked ‘A’ and in the Not Revealed section 25% picked A, there is a 40 percentage point gap between the modal response in the Revealed treatment and the corresponding answer in the Not Revealed treatment. To classify as a cascade, though, we also look at the modal response in the Not Revealed treatment. For Question 1, this is ‘C’, with 50% of students selecting this response. In the Revealed treatment only 10% of the students selected ‘C’, so the difference in responses between the treatments is 40 percentage points. Since both differences between the modal responses differ by more than 20 percentage points, Question 1 is classified as a cascade. In the Revealed treatment, students disproportionately answered this question incorrectly, so it is a negative cascade.

In Question 2, the assessment of cascading requires only a single calculation, as the modal response is the same in both treatments: answer A. In the Revealed treatment ‘A’ was selected 75% of the time, in the Not Revealed treatment, ‘A’ was selected 50% of the time. Since there is a 25-percentage point difference between the two treatments, this is classified as a cascade. Since the cascade occurred on the correct answer, we define this as a positive cascade.

In the remainder of the paper, we answer the following 2 questions:

1. Do we observe positive and negative cascading?
2. Does cascading effect exam scores? We consider the differential effects of positive and negative cascades on performance in the final exam.

RESULTS

Table 2: Summary of the data by treatment.

	Not Revealed	Revealed	Total
THCorrect	0.495 (0.500)	0.526 (0.500)	0.510 (0.500)
Fraction of answers that are modal	0.639 (0.212)	0.846 (0.161)	0.742 (0.215)
Student chose modal response	0.551 (0.498)	0.795 (0.404)	0.672 (0.470)
Fraction of correct peer responses when answer was submitted	0.483 (0.290)	0.504 (0.410)	0.493 (0.355)
Participation in TH	0.672 (0.0664)	0.678 (0.0607)	0.675 (0.0637)

Exam 3 (%)	57.52 (14.29)	57.33 (14.27)	57.42 (14.27)
Summary of cascading behavior			
Early positive cascade count			1.561 (1.035)
Early negative cascade count			1.789 (0.977)
Placebo early negative cascade count			0.719 (0.453)

Standard deviations in parentheses. The Participation in TH variable measures the fraction of enrolled students who participated in a TH question in a given class session. As it was extremely rare for students to attend class and not participate in TH, we can interpret this variable as class attendance.

Table 2 presents summary statistics for some variables in our analysis, split by the Revealed and Not Revealed treatments. $THCorrect_{i,q}$ is an indicator variable equal to 1 if student *i* correctly answered TopHat question *q*. On average, we see no statistically significant difference in correct responses between the two treatments (p=0.191, 2-sided *t*-test for equal means). While this lack of a statistically significant difference may seem surprising, it is important to remember that students may artificially get a correct answer in a positive cascade an artificially get an incorrect answer in a negative cascade. The collective impact of the cascades are opposite in direction, leaving the average correctness approximately the same between the revealed an not revealed treatments. However, the second row indicates a stark difference between the treatments. The variable *Fraction of answers that are modal* reports the fraction of answers that are modal *at the time a student submits their answer*. In the Revealed treatment, students on average saw about 85% of their peers' answers were for the modal answer. In the Not Revealed treatment, where students do not observe peer feedback, this number was 65% on average. We reject that these fractions are equal (p<0.001). When we consider the results from the end of the question period, in the Revealed treatment, students chose the modal response with probability 80%, while in the Not Revealed treatment they were 23 percentage points less likely to do so (p<0.001). In 35 out of 38 questions, the fraction of modal answers was larger in the Revealed treatment (p<0.001, 1-sided 50% Binomial test). We interpret these results as students being influenced by public signals: in the Revealed treatment we observe more students selecting the mode during and at the end of the question period.

The variable "Fraction of correct peer responses when answer was submitted" is equal to the fraction of peer responses that were correct at the time a student submitted an answer. We find no significant difference in this variable between treatments. Table 2 also reports the mean of final exam scores. If these were systematically different by treatment, we may be worried that differently-performing students may be self-selecting out of one treatment (by not participating in some questions). We find no support for such selection, and conclude that on these measures we have adequately balanced the treatment assignment. Finally, we check for randomization over the two sections. The early section participated in 21 Revealed and 17 Not Revealed treatments. We fail to reject that this assignment was generated from a sample of 50% independently and identically distributed Bernoulli random variables (p=0.63, 2-sided Binomial test).

Summary of Cascading Behavior

Figure 3: Graphical representation of a positive information cascade.

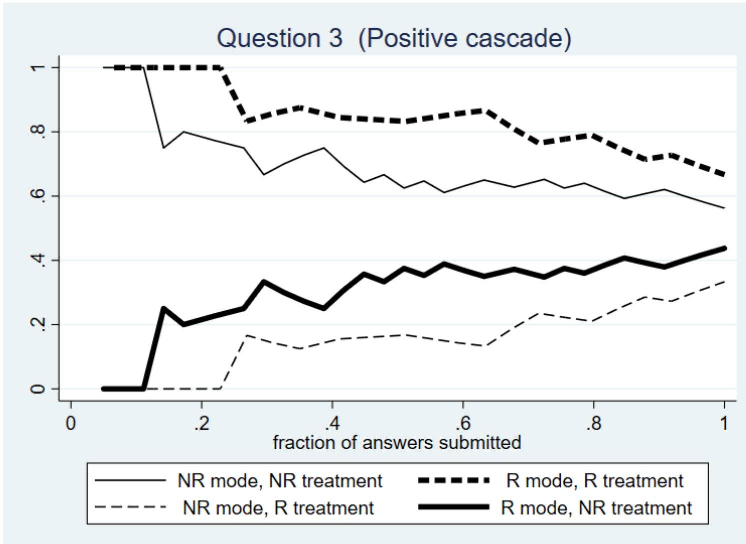

NR mode, NR treatment is the fraction of students who selected the modal answer in the Not Revealed treatment. NR mode, R treatment is the fraction of students in the Revealed Treatment who selected the modal response from the Not revealed treatment. R mode, NR treatment is the fraction of students in the Revealed Treatment who selected the modal answer in the Not Revealed treatment. R mode, R treatment is the fraction of students who selected the modal answer in the Revealed treatment.

Figure 4: Graphical representation of a negative information cascade

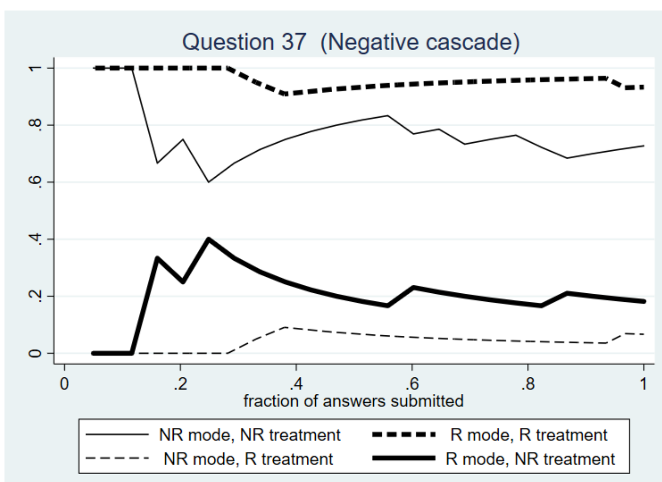

NR mode, NR treatment is the fraction of students who selected the modal answer in the Not Revealed treatment. NR mode, R treatment is the fraction of students in the Revealed Treatment who selected the modal response from the Not revealed treatment. R mode, NR treatment is the fraction of students in the Revealed

Treatment who selected the modal answer in the Not Revealed treatment. R mode, R treatment is the fraction of students who selected the modal answer in the Revealed treatment.

Figure 5: Graphical representation of a TopHat question without an information cascade

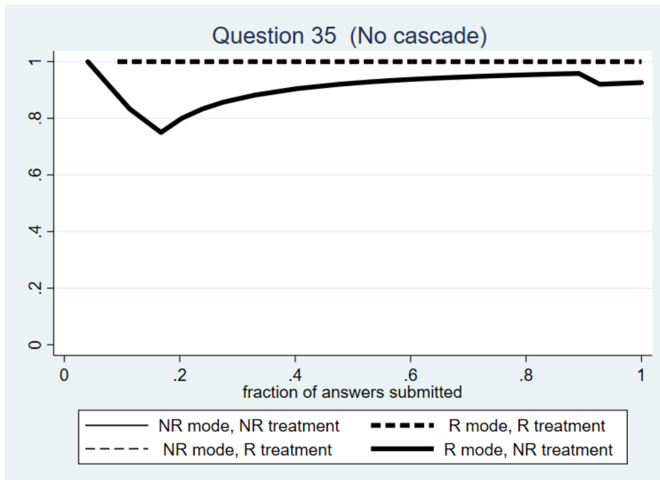

NR mode, NR treatment is the fraction of students who selected the modal answer in the Not Revealed treatment. NR mode, R treatment is the fraction of students in the Revealed Treatment who selected the modal response from the Not revealed treatment. R mode, NR treatment is the fraction of students in the Revealed Treatment who selected the modal answer in the Not Revealed treatment. R mode, R treatment is the fraction of students who selected the modal answer in the Revealed treatment.

Figure 6: Graphical representation of a TopHat question without an information cascade.

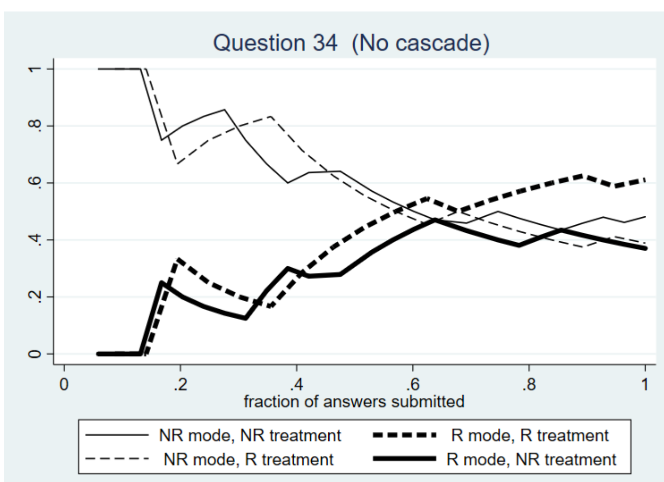

NR mode, NR treatment is the fraction of students who selected the modal answer in the Not Revealed treatment. NR mode, R treatment is the fraction of students in the Revealed Treatment who selected the modal response from the Not revealed treatment. R mode, NR treatment is the fraction of students in the Revealed Treatment who selected the modal answer in the Not Revealed treatment. R mode, R treatment is the fraction of students who selected the modal answer in the Revealed treatment.

We classify approximately half (20 of 37) of TopHat questions as information cascades, roughly evenly split into positive (9 of 37) and negative cascades (11 of 37). Figures 3-6 show examples of these classifications. For the remainder of our analysis, we code the correct answer as "Answer 1". For example, if answer *c* was the correct answer, this question would be coded: *c*=1, *a*=2, *b*=3, *d*=4. On each figure, we differentiate between treatments using solid lines to show the fraction of submitted answers in the Not Revealed treatment, and dashed lines to show this fraction for the Revealed treatment. In order to assess cascades, we need to differentiate between the modal responses in each treatment. The thick dashed line denotes the answer that was the modal response in the Revealed treatment. The thick solid line shows the fraction of students who picked this answer in the Not Revealed treatment. Similarly, the solid thin line shows the modal response for the Not Revealed treatment, and the dashed thin line shows the fraction of students who picked that answer in the Revealed treatment.

Our classification of cascades is based on the difference between the rightmost point of both the thick lines (i.e. comparing the modal answer in the Revealed treatment across both treatments), and the thin lines (i.e. comparing the modal answer in the Not Revealed Treatment across both treatments). Specifically, we require a 20 percentage point difference between these fractions at the end of the question period. Figure 3 provides an example of a positive cascade. Here, in the Revealed treatment approximately 65 percent of students answered correctly (i.e. chose Answer 1). Only 40 percent of students did so in the Not Revealed treatment. Answer 2 was the modal response in the Not Revealed treatment with about 55 percent of the students selecting it. In contrast, only about 35 percent of students selected Answer 2 in the Revealed treatment. Figure 4 is an example of a negative cascade. In the Revealed treatment, Answer 4 was the modal response, but few students selected it in the Not Revealed treatment. Figures 5 and 6 are examples of questions that were not classified as cascades. In Figure 5, almost all students chose the correct answer in both treatments. In Figure 6, the distribution of answers at the end of the question period was quite similar in both treatments.

CASCADING – CONDITIONAL LOGIT ANALYSIS

Next, we analyse the practical effect of revealing peer responses on student outcomes. Specifically, we are interested in how the popularity of the correct (incorrect) answer among a student's peers will affect the odds that the student provides the correct (incorrect) answer. We define $AnswerCount_{i,q,j}$ as the number of students who had submitted answer $j \in \{a, b, c, d\}$ at the time student i submitted their answer to question q. That is, in the Revealed treatment this variable is equal to the number of peer responses to answer *j* displayed on the projector when student i submitted their answer. For practical and theoretical reasons, we model students' decisions using a conditional logit:

Table 3: Conditional logit estimations from the TopHat data.

	(1) Answer	(2) Answer	(3) Answer
Answer count x Revealed treatment	1.070**	1.162***	1.180***
	(0.0223)	(0.0228)	(0.0240)
Answer count	1.131***	0.939***	0.926***
	(0.0129)	(0.0153)	(0.0157)
Observations	7824	7824	7824
Included Fixed Effects	-	question	question and treatment

Exponentiated coefficients; Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Standard errors clustered at the student level. In all models the dependent variable is the answer chosen by the student. The coefficient on *Answer count * Revealed treatment* represents how much more likely a student is to choose an answer, say 'a' in the Revealed treatment when there is one additional peer response on answer 'a'. The coefficient on *Answer count* represents how much more likely a student in a Not Revealed treatment is to pick answer 'a' when an additional peer selects answer 'a'.

$$p\left(y_{i,q,k} \mid X_{i,q} \right) = \frac{\exp\left(\beta X_{i,q,k}\right)}{\sum_{j\in\{a,b,c,d\}} \exp\left(\beta X_{i,q,j}\right)}$$

Where $X_{i,q,k}$is a vector of data describing properties of answer $k \in \{a, b, c, d\}$ in question q for student i. This ratio is the probability of picking answer 'k', any one of our four answer choices. For practical reasons, this model fits well with our data in that the choice data are categorical. Additionally, one can interpret $\beta\left(X_{i,q,a} - X_{i,q,b}\right)$ as the (reduced form) difference between student i's expected utility from choosing answers a and b. We focus on the coefficient on $AnswerCount_{i,q,j}$ and its interactions. β_1, the coefficient on this variable, tells us that if answer a has one more peer response than answer b, (That is: $AnswerCount_{i,q,a} - AnswerCount_{i,q,b} = 1$) then a student will be $100(\exp(\beta_1) - 1)\%$ more likely to choose answer a than answer b, all else held equal. We report the exponentiated coefficients of three specifications, shown in Table 3. In all of these specifications, we include $AnswerCount_{i,q,j}$ (the second row), as well as $AnswerCount_{i,q,j}$ interacted with a dummy variable for the Revealed treatment (first row). Therefore, one can interpret the coefficient on *Answer Count * Revealed Treatment* as representing how much more likely a student is to choose an answer, say 'a' in the Revealed treatment when there is one additional peer response on answer 'a'.

Similarly, one can interpret the coefficient on the "Answer Count" variable as the effect of peer responses in the Not Revealed treatment (i.e. the case where there should be no effect). We include $AnswerCount_{i,q,j}$ alone as we are specifically looking for *differences* between the treatments: the exponentiated coefficient for the interaction term therefore tells us how much more likely a student is to choose an answer in the Revealed treatment compared to the Not Revealed treatment with the same distribution of peer responses. The first column of Table 3 shows this specification, then columns 2 and 3 add question and treatment fixed effects.[i] Depending on the controls used, we estimate that one additional peer response on a particular answer makes a student between about 7% and 18% more likely to choose that answer in the Revealed treatment (all three of these estimates are significant at the 0.1% level). To put this in perspective, suppose that the first four peer responses were for answer a, these models predict that a student will be between 31% and 93% more likely to choose that answer than had they not seen the peer feedback (i.e. $0.31 \approx 1.07^4 - 1$, and $0.93 \approx 1.18^4 - 1$). Our preferred specification is column 3, as it controls for both question and treatment fixed effects, which estimates students are 18% more likely to pick an answer for every peer response on that answer.

TopHat's Predicative Power for Exam Scores

While the presence or absence of peer feedback influences students' choice of answer in TopHat, we now consider the impact of peer responses on students' final exam scores. Table 4 shows regressions of students' final exam scores (measured as a percentage) against various measures of their performance in TopHat.[ii] Column 1 estimates that, controlling for participation, submitting an additional correct TopHat answer predicts an approximately 1.3 percentage point increase in a student's final exam score. In Column 2, we allow this effect to vary by treatment, but fail to reject that correct answers in the Revealed and Not Revealed treatments affect exam performance differently (p=0.809). Note that since we control for participation, the interpretation of these coefficients is the effect on exam scores of correctly answering a question, conditional on participating in that question.

Column 3 investigates the differential impact of positive and negative cascades, by including a count of the number of times a student participated in either a negative or positive cascade. Here we restrict our analysis to the first three negative cascades in each classroom, which occurred in the first half of the semester, however including all negative cascades in the analysis does not change the results substantially or significantly. Participating in one additional negative cascade results in a loss of approximately five percentage points on the final exam. Participating in a positive cascade, however, has no effect on exam scores. One might understand the combination of these two results as a function of underlying student ability. High ability students are more likely to get the correct answer and thus participate in a positive cascade. Low ability students are more likely to be swayed by pier responses, if they do not believe strongly in the quality of their private signal. Thus, low ability students are more likely to participate in a negative cascade. Taken together, it makes sense that negative cascades, which are more likely to draw in low ability students, predict lower exam performance. Similarly, this would explain why positive cascades have no impact on exam scores, as they are more attractive to high ability students who would have done well on the exam in the absence of the revealed treatment. Finally, we estimate Column 4 as a placebo test. It could be that the (incorrect) modal answers in these six questions (three for each class) are good predictors of poor performance, irrespective of whether the student chose that answer with or without seeing their peers' answers. We therefore count the number of times each student in the Not Revealed treatment chose what was the modal response in the other section. This coefficient is not significantly different from zero, indicating that it is the cascading itself, not just the incorrect answers, that predict lower exam scores.

Table 4: Analysis of exam scores.

	(1) Exam 3 (%)	(2) Exam 3 (%)	(3) Exam 3 (%)	(4) Exam 3 (%)
Total TH correct answers	1.340*			
	(0.631)			
Participation	-0.437	-0.419	0.0643	-0.120
	(0.324)	(0.334)	(0.378)	(0.389)
class_id=1	-1.184	-1.374	1.545	1.923
	(3.152)	(3.284)	(4.958)	(4.878)
TH correct answers in Revealed treatment		1.088	0.444	0.531
		(1.199)	(1.257)	(1.210)
TH correct answers in Not Revealed treatment		1.507	1.370	1.352
		(0.961)	(0.925)	(0.878)
Early negative cascade count			-5.172*	-4.705*
			(2.251)	(2.261)
Early positive cascade count			0.722	1.033
			(2.649)	(2.660)
Placebo early negative cascade count				5.316
				(4.093)
Constant	50.16***	50.42***	47.94***	47.86***
	(3.292)	(3.471)	(4.700)	(4.723)
Observations	71	71	71	71
r2	0.130	0.131	0.190	0.211
R & NR coefficients zero		0.121	0.257	0.218
R & NR coefficients equal		0.809	0.597	0.624

Note: Heteroskedasticity-robust standard errors in parentheses. Table 4 shows the impact of TopHat on exam performance. In each specification the dependent variable is the student's final exam score. In Column (1) we estimate the impact of correctly answering a TopHat question on Exam performance, Column (2) includes the impact of the Revealed and Not Revealed treatments, Column (3) is our preferred specification, illustrating the impact of negative cascades on exam scores, while Column (4) demonstrates, through a placebo test, that it is participating in the negative cascade, rather than getting the answers incorrect that causes the reduction in the final exam score.

CONCLUSION

In this field experiment, we document the presence and consequences of information cascades in an undergraduate classroom. The results of this experiment highlight that there are trade-offs for instructors and students in seeing the distribution of student answers during in-class assessment. For teachers, allowing students to view peer responses may cause an obscured view of what their students truly understand. This can lead to inefficient allocation of class time as wellunderstood concepts are covered again, and poorly-understood concepts are skipped over. For students, picking the wrong answer because it is popular may unduly discourage the student, making them think the concept is ‘unknowable’ as everyone got it wrong. In fact, students who cascade on an incorrect answer are likely to have a final exam score which is 5 percentage points lower. We see no such effect for getting the same question wrong in the Not Revealed treatment. Therefore, we conclude it is participation in the cascade that causes the impact on understanding.

As educators, we may care about predicting student performance (beyond causal inference) and thus may like a diagnostic tool to help identify poorly performing students. To this end, getting a TopHat question incorrect[iii] predicts a lower final exam score. What is unknown is if participating in a negative cascade has additional predictive power for poor exam performance. While this is not a question we can answer with our data, it is possible that using a ‘Revealed’ treatment in class may aid educators in identifying ‘at-risk’ students early in the course. However, we would recommend using this tool cautiously, as it is clear that when students participate in a negative cascade, they do worse on the final exam! We hypothesize this could be a peer effect where they deem the question so difficult that it isn’t worth investing study time on this topic, since many of their peers got it wrong as well. Alternatively, being one of many students who answered incorrectly in a negative cascade may provide misleading information about relative ranking in the class. Colloquially, if I am part of the herd, I won’t do too badly relative to everyone else.

Alternatively, we might imagine that students do some updating regarding the quality of peer information over the course of the class. It may be the case where higher ability students are better able to take advantage of peer information then their lower ability counterparts. This would increase the likelihood that a lower ability student participates in a negative cascade later in the course.

Within the context of this study, we did not investigate how the probability of having a negative information cascade changed depending on question characteristics. Future research could investigate if presenting a question on concepts learned in the same class versus a review question at the start of the next class period change the probability of a negative information cascade. In this study, we randomly assigned the Revealed or Not Revealed treatment, but question type and timing were not a focus of our experiment design.

This experiment documents information cascades outside of the laboratory. While there is a host of literature which investigates information cascades in the lab, observing them in the field widens our view of how information cascades impact decision making and the welfare of groups making decisions.

REFERENCES

Anderson, L. R. and C. A. Holt (1997). Information cascades in the laboratory. *American Economic Review*, 87(5), 847-862.

Azmat, G. & N. Iriberri (2010). The importance of relative performance feedback information: Evidence from a natural experiment using high school students. *Journal of Public Economics*, 94(7), 435-452.

Bandiera, O., V. Larcinese & I. Rasul (2015). Blissful ignorance? A natural experiment on the effect of feedback on students' performance. *Labour Economics*, 34, 13-25.

Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3), 797-817.

Bikhchandani, S., D. Hirshleifer & I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5), 992-1026.

Blanes i Vidal, J. & M. Nossol (2011). Tournaments without prizes: Evidence from personnel records. *Management Science*, 57(10), 1721-1736.

Bursztyn, L., F. Ederer, B. Ferman & N. Yuchtman (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica*, 82(4), 1273-1301.

Chamley, C. & D. Gale (1994). Information revelation and strategic delay in a model of investment. *Econometrica*, 62(5), 1065-1085.

Cingl, L. (2012). Does herd behavior arise more easily under time pressure? Experimental approach, IES Working Paper.

Eriksson, T., A. Poulsen & M. C. Villeval (2009). Feedback and incentives: Experimental evidence. *Labour Economics*, 16(6), 679-688.

Falk, A. & A. Ichino (2006). Clean evidence on peer effects. *Journal of Labor Economics*, 24(1), 39-57.

Gul, F. & R. Lundholm (1995). Endogenous timing and the clustering of agents' decisions. *Journal of Political Economy* 103(5), 1039-1066.

Harrison, G. W. & J. A. List (2004). Field experiments. *Journal of Economic Literature*, 42(4), 1009-1055.

Mas, A. & E. Moretti (2009). Peers at work. *American Economic Review*, 99(1), 112-145.

Sgroi, D. (2003). The right choice at the right time: A herding experiment in endogenous time. *Experimental Economics* 6(2), 159-180.

Tucker, C. & J. Zhang (2007). Long tail or steep tail? A field investigation into how online popularity information affects the distribution of customer choices. Sloan Working Papers, Sloane School of Management. (mit.edu).

---

[i] In the conditional logit model, as there are four outcomes, there are three “question fixed effects” per question, corresponding to answers 2, 3, and 4. Since answer 1 is the base case, the exponentiated fixed effect for answer $l \in \{2,3,4\}$ is equal to the probability of choosing answer l divided by the probability of choosing answer 1 (setting all covariates equal to zero). Hence, these fixed effects allow each question to have different distributions of answers when peer feedback is withheld.

[ii] We also investigated the power of students' answers to individual TopHat questions in predicting their answers to similar exam questions. We use a logistic regression of correct exam questions against various data collected in TopHat. We find that correctly answering a TopHat question has no predictive power for success on a similar exam question. Similar to our estimates in Table 4, we find that participating in a negative information cascade predicts that a student will be less likely to correctly answer a similar exam question. However, getting a TopHat question incorrect does not predict similar exam questions. Again, it is the negative cascading, not simply incorrectly answering TopHat questions, that predicts more incorrect answers to exam questions.

[iii] See Column (1) of Table 4. We have the coefficient for THCorrect (1.340), so the impact of getting a question incorrect would be -1.340 percentage points on the final exam.