# Game-theoretic models identify useful principles for peer collaboration in online learning platforms

**Vineet Pandey**
UC San Diego Design Lab
La Jolla, CA 92093-0404 USA
vipandey@ucsd.edu

**Krishnendu Chatterjee**
Institute of Science and Technology
Am Campus 1 3400
Klosterneuburg Austria
Krishnendu.Chatterjee@ist.ac.at


CSCW '16 Companion, February 27 - March 02, 2016, San Francisco, CA, USA
ACM 978-1-4503-3950-6/16/02.
http://dx.doi.org/10.1145/2818052.2869122

## Abstract

To facilitate collaboration in massive online classrooms, instructors must make many decisions. For instance, the following parameters need to be decided when designing a peer-feedback system where students review each others' essays: the number of students each student must provide feedback to, an algorithm to map feedback providers to receivers, constraints that ensure students do not become *free-riders* (receiving feedback but not providing it), the best times to receive feedback to improve learning etc. While instructors can answer these questions by running experiments or invoking past experience, game-theoretic models with data from online learning platforms can identify better initial designs for further improvements. As an example, we explore the design space of a peer feedback system by modeling it using game theory. Our simulations show that incentivizing students to provide feedback requires the value obtained from receiving a feedback to exceed the cost of providing it by a large factor (greater than 7). Furthermore, hiding feedback from low-effort students incentivizes them to provide more feedback.

## Author Keywords

peer feedback; game theory; MOOCs; online collaboration.

*Simple game formulation for N=3 players*
Consider Alice, Bob and Charlie who are playing the feedback game. *C* and *V* refer to the Cost and Value functions for the feedback system, where $C_{Alice}$ is the cost incurred by Alice in providing a feedback (e.g.: time taken) and $V_{Alice}$ is the value of the feedback provided by Alice (e.g.: number of changes made by a student in response to the feedback).

Alice provides feedback to Bob and receives feedback from Charlie. The payoff for Alice ($P_{Alice}$) is the value of feedback received minus the cost incurred in providing feedback. Hence, $P_{Alice} = V_{Charlie} - C_{Alice}$. Alice's payoff will be high if Charlie submits a high-quality feedback or Alice incurs a small cost when providing feedback (i.e. she submits a poor feedback). Since Alice does not control Charlie's action, she has the incentive to provide a poor feedback to Bob.


## ACM Classification Keywords
K.3.1 [Computer Uses in Education]: Distance learning, Collaborative Learning


## Online peer feedback scales well but with uneven participation
Incentivizing collaboration in online learning platforms is difficult; structuring peer interactions over chat [1] and building a useful peer-feedback system [2] demonstrate the vast number of choices that need to be made by instructors to ensure rapid, useful peer collaboration. *Peer feedback* in massive online classrooms is a technique where students provide subjective feedback on submissions of other students. For instance, the following parameters need to be decided when designing a peer-feedback system where students review each others' essays: the number of students each student must provide feedback to, an algorithm to map feedback providers to receivers, constraints that ensure students do not become *free-riders* (receiving feedback but not providing it), the best times to receive feedback to improve learning, etc. Such a large number of decisions can make it difficult for instructors to find the right tweaks to improve collaboration and provide benefits to students.

While instructors can answer some questions by running experiments with a smaller set of users, this still requires building relevant systems or prototypes to test and gather data. Furthermore, incentivizing users to take part in detailed experiments is difficult. All these factors make the overall process extremely long. For example, the complete cycle of development, debugging, user testing and deployment for a peer feedback research system took a year [2]. Moreover, it is uncertain whether insights generated from small experiments are true at a large scale. For example, solving the free-rider problem can be notoriously difficult in large anonymous networks [3] while users are more likely to co-operate in a small group [4]. *Can we make it easier for instructors to identify the right tweaks by reducing the design space of collaborative learning platforms using automated techniques?*

## Game-theoretic model for peer feedback
To understand how student behavior might be tweaked by design choices, we use game theory to model peer feedback at the scale of a thousand students. We demonstrate ways to identify better design choices, such as techniques to incentivize submitting feedback, **before** building the system. Our simulations do not assume rational behavior from the players (students in the peer feedback system). Instead, different players have strategies that evolve depending on payoffs provided by the game. A strategy that does not provide high payoff gets eliminated. Through simulations, we can predict parameters that help the game move towards favorable outcomes, which corresponds to more students submitting feedback. Previous work has used game theory to model peer reviews to minimize the error in grading assuming the availability of instructor grades [5]. We formally model peer feedback as a simple game without instructor involvement with the goal of increasing the number of students who provide feedback.

Our game formulation is not designed to perfectly model real world interactions but to demonstrate two points. First, it is feasible to construct a game to model interactions in a massive online education system among students. Second, we can derive useful insights by evaluating the game using different game-theoretic

**Playing games with Darwin**

*Game parameter:* Benefit ratio = ratio between the value of receiving a feedback and the cost of providing a feedback.

*Player strategies:* Student's frequency of providing feedback. Every game begins with 1000 players whose frequency of providing feedback varies between 0.0 and 0.9 (at 0.1 increments, with 100 learners in each group). The players with frequency 0.0 do not provide any feedback while players with frequency value 0.9 are highly likely to provide feedback.

*Game constraint:* threshold = the number of feedbacks that must be provided by a player to see her own feedback. We tweak the game such that a player cannot receive the feedback on her submissions until she provides a certain number of feedback, effectively creating a reward/punishment setup.

techniques. Improving the model using known results from learning science as well as insights from instructors will further refine the quality and strength of recommendations offered. These recommendations can then be used to build systems that amplify the findings with behavioral cues.

*Darwin: Simulating peer feedback at scale*
Given the parameters of the feedback game and the different strategies employed by players, can we add constraints to the game to incentivize players to provide feedback? To answer this, we built Darwin: a tool that simulates a peer feedback system for 1000 students. Using simulated learners is a known technique; previous work [e.g., 6] has used simulated learners to identify the effectiveness of tutoring systems. Our **independent variables** are the game parameter (benefit ratio: Value of receiving a feedback/ Cost of providing a feedback) and the constraint (threshold: the number of feedbacks that must be provided by a player in order to see her own feedback). Since we need to observe how our choices affect the player strategies, our **dependent variables** are the number of successful players (those who cross the threshold to view their own feedback) and the frequency distribution of population (number of players in each of the 10 frequency bins (0.0, 0.1, …, 0.9)).

## Results

### Effect of varying game parameter: Benefit ratio

For a low benefit ratio (when the value of receiving a feedback is same as the cost of providing it), students modify their strategy to provide less feedback. For a high benefit ratio, such as 10, the number of students providing feedback increases even for high threshold, such as 5. (Figure 1)

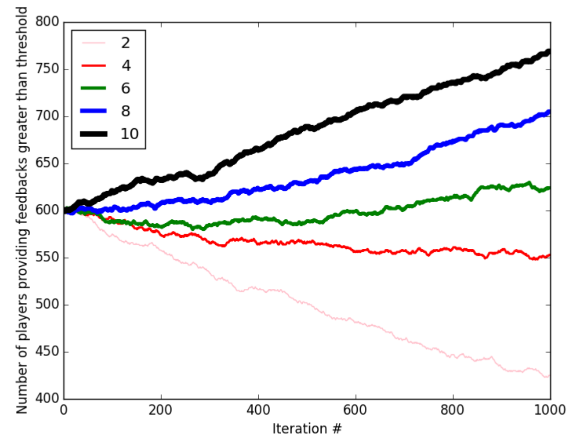


Figure 1: More players (out of 1000 total players) provide feedback when the benefit ratio is high (the lines show different benefit ratio values) i.e. when the value of receiving a feedback is much more than the cost of providing feedback. The threshold was kept constant at 3 across all the runs, but this trend is consistent for any fixed value of threshold.

### Effect of varying the game constraint: threshold

If the threshold is set to 0, then all players resort to providing no feedback since they can trivially cross the threshold with no effort. This demonstrates that selfish choices might lead to overall degradation of the peer feedback system. Also, setting unrealistically high thresholds incentivizes players to provide less effort. Figure 2 shows that a threshold of 11 with a benefit ratio of 8 incentivizes more students to provide no feedback at all (occupy 0.0 bin), since they cannot cross such a high threshold. For a given threshold value, there is an optimum benefit ratio (B*) such that the population converges to a favorable outcome for any chosen benefit ratio greater than B* and to an unfavorable outcome for any value less than B*. Figure 1 shows that for threshold=3, B*=6.

**Classical Game Theory provides limited insights**

| **Alice** / Bob | Provides feedback | Abstains |
|---|---|---|
| Provides feedback | **1**, 1 | **-1**, 2 |
| Abstains | **2**, -1 | **0**, 0 |

Table 1: Sample Payoff matrix for a simple model of the feedback game for two players where receiving feedback provides a value of 2 units while providing feedback incurs a cost of 1 unit.

The stable equilibrium for the game shown in Table 1 occurs when both the players abstain from providing feedback. This holds true for any general values for cost and value of a feedback. Hence, Classical game theory predicts that we cannot tweak the parameters to incentivize learners to provide feedback. However, Classical game theory analysis suffers from multiple concerns that make the results difficult to trust: it assumes people make strictly rational choices, and shows poor scaling with the number of players.

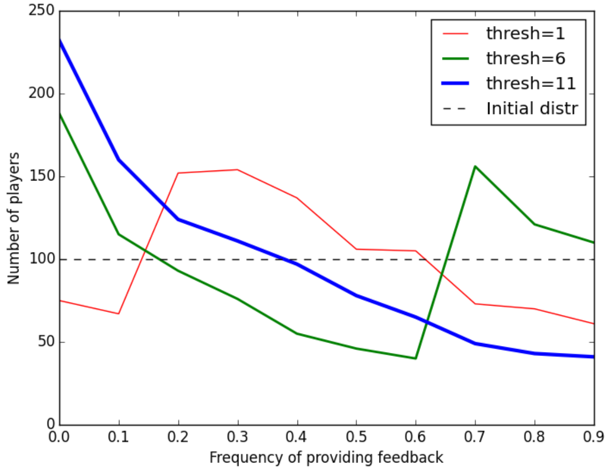


Figure 2: The final distribution of player strategies (frequency of providing feedback) highly depends on the threshold value. When a player needs to provide feedback to 11 players (which is difficult), more players choose to provide very few or no feedbacks at all to save on effort. For the three runs with threshold=1, 6, and 11, the benefit ratio was constant at 8.

To conclude, we have built an evolutionary game-theoretic model of the peer feedback problem to find successful strategies under varying constraints. Our simulations show that a simple model based on hypothetical data can provide intuition about possible student response. Our results demonstrate the difficulty in eliciting feedback from players at a large scale. Going forward, we intend to build a richer model by mining data from MOOCs to identify real player strategies and using insights from instructors to design game parameters and constraints. Finally, we intend to test our insights around incentivizing peer feedback by deploying our system for a real MOOC.


## Acknowledgements

ERC Start Grant Graph Games 279307 supported this research.


## References


1. D. Coetzee, Seongtaek Lim, Armando Fox, Bjorn Hartmann, and Marti A. Hearst. 2015. Structuring Interactions for Large-Scale Synchronous Peer Learning. In Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15). 1139-1152. http://dx.doi.org/10.1145/2675133.2675251
2. Chinmay E. Kulkarni, Michael S. Bernstein, and Scott R. Klemmer. 2015. PeerStudio: Rapid Peer Feedback Emphasizes Revision and Improves Performance. In *Proceedings of the Second (2015) ACM Conference on Learning @ Scale* (L@S '15). 75-84. http://doi.acm.org/10.1145/2724660.2724670
3. Bram Cohen. 2003. Incentives Build Robustness in BitTorrent. Retrieved Oct 29, 2015 from http://www.bittorrent.org/bittorrentecon.pdf
4. Daniele Nosenzo and Simone Quercia, and Martin Sefton. 2015. Cooperation in small groups: the effect of group size. Experimental Economics. http://dx.doi.org/10.1007/s10683-013-9382-8
5. William Wu, Constantinos Daskalakis, Nicolaas Kaashoek, Christos Tzamos, and Matthew Weinberg. 2015. Game Theory based Peer Grading Mechanisms for MOOCs. In Proceedings of the Second (2015) ACM Conference on Learning @ Scale (L@S '15). 281-286. http://doi.acm.org/10.1145/2724660.2728676
6. González-Brenes, José P., Huang, Yun. 2015. Using Data from Real and Simulated Learners to Evaluate Adaptive Tutoring Systems. Proceedings of the Workshops at the 18th International Conference on Artificial Intelligence in Education AIED 2015.