Incentivizing Participation in Online Forums for Education

Arpita Ghosh, Cornell University
Jon Kleinberg, Cornell University

We present a game-theoretic model for online forums for education, where students in a class can post questions to the forum and seek responses from the instructor or other students in the class. We first show that our model predicts the anecdotally observed phenomenon that students' participation in a forum is non-monotone in the instructor's response rate to questions: if the instructor responds at too high a rate students do not respond to their peers' questions, whereas an almost-absent instructor induces very little participation from the students as well. We then investigate the optimal use of a forum for two kinds of questions— single-answer questions and discussion-style questions— which lead to different levels of rewards that can be meaningfully offered. We show that for discussion-style questions, the instructor can choose her response rate so that the expected rate at which the first student response is received increases linearly with the size of the class; for single-answer questions, however, the optimal expected rate of arrival of the first student response remains a constant even as class size diverges. However, this slow response rate can be remedied by mixing question types— as long as there is any positive probability of a discussion-type question in a forum, the instructor can choose her response rate so that the equilibrium rate of the first response from the class diverges with the number of students in the class.

1. INTRODUCTION

The increasing use of online applications for education raises a spectrum of research issues — ranging across the delivery of course content, the evaluation of student work, and the development of resources to stimulate discussion and address student questions. We focus on resources of this latter type — *online forums for education* oriented around discussion and question-answering — which are increasingly found in a number of different educational contexts. In many cases, they exist as supplements to standard offline classes through sites such as Piazza— a rapidly growing online platform for education forums with over 2500 participating classes and already used by over 100,000 students— which provide students with a separate online venue for asking questions and receiving help. In other cases, they are being explored as a key component of massive open online courses (MOOCs), creating communities of students who can discuss course content online.

An important property of such forums, when they work well, is that they are not simply channels for instructors to provide information to students — they also involve students providing significant help to each other, thereby tapping into the peer-learning potential of a large class which, research in education shows, can provide significant learning benefits to students [Topping and Ehly 1998]. Students arrive at the forum with questions or topics they want help with, and fellow students, together with the instructor, provide follow-up discussion and answers. These education forums have much in common with online question-answering (Q&A) sites such as StackOverflow, MathOverflow, or Yahoo! Answers, but also differ in key ways which play important roles in their design:

— Online forums for education have an instructor who can — if she chooses — play a crucial and visible role in the operation of the forum. In addition to providing answers of her own, the instructor can act as a global monitor of the site, and reward students for participation on the forum, both via explicit incentives such as class grades, as well as via social-psychological rewards from instructor approval or endorsement. For these reasons, an instructor in an online forum for educa-

Authors' emails: arpitaghosh@cornell.edu, kleinber@cs.cornell.edu.

DOI: http://dx.doi.org/10.1145/0000000.0000000

tion has access to incentives that operate particularly powerfully relative to moderators on general Q&A sites, who tend to wield a qualitatively lower level of influence than the instructor in an education forum.

— The presence of an instructor who can monitor and potentially penalize poor-quality contributions, as well as the presence of social norms similar to those in a classroom, mean that contribution quality, while still an issue, is arguably much less heterogeneous than in general online Q&A forums on the Web. Indeed, low effort and junk contributions appear to be significantly less common in Piazza and other education forums in comparison with general-purpose forums like Y! Answers.

From the growing use of online forums for education, we see that they can operate well — with specific questions getting resolved quickly, and open-ended discussion questions receiving a strong flow of contributions. But they can also operate poorly, with little student participation and questions that go unanswered for long periods of time. An instructor would like to make use of an online forum so as to promote better outcomes. How can an instructor use her own behavior to incentivize higher student participation? For example, if the instructor actively appears on the site, does this tend to increase student participation (because they are hoping to get the instructor's approval) or decrease it (because they assume that the instructor will answer the questions herself)? Is there an optimal way to trade off these two forces, and how does the best outcome the instructor can elicit through her behavior depend on the size of the class — as the class size becomes large, can the instructor step back and let the class do most of the work on the forum? While such questions appear at the heart of discussions about best practices in these forums [Andresen 2009; Burge et al. 2000; Gray 2004; Hanover Research Council 2009; Mazzolini and Maddison 2003], there has so far not been a formal way of reasoning about them.

1.1. Overview of results

We develop and analyze a game-theoretic model of student participation in online forums for education, and use this model to address questions regarding the optimal use of such forums. While a range of factors, including psychology and user interface design, will clearly play a role in participation on these forums, we argue using our model that a game-theoretic analysis of incentives can provide a useful framework for understanding many of the trade-offs that have emerged in discussions of education forums thus far. In particular, our model requires very few ingredients to define formally yet exhibits rich behavior, including some of the subtle trade-offs in class size, instructor activity, and instructor effort level that have been observed qualitatively in early empirical studies of these forums.

There is a class of n students, and an instructor, each of whom chooses *rates* at which to check the class's online forum for new content; a rate of λ corresponds to repeatedly checking the forum according to a Poisson process with rate λ, truncated to a finite time interval (§2). Checking the forum more frequently, *i.e.*, at a higher rate, incurs a higher *cost*, but potentially also leads to higher benefits, as follows. A student who answers a question is 'rewarded' by the instructor for (usefully, and correctly) answering a question on the forum— this reward could either be points awarded by the instructor that count towards the student's grade in the class, or a social-psychological reward deriving from approval, or endorsement, by the instructor. An instructor rewards students only for answering questions 'in time', *i.e.*, for answering new questions that have arrived since the last time she visited the forum, as explained in §2. (Of course, the instructor could choose to reward students for delayed answers as well, but as our analysis will show, this does not create the right incentives for participation).

Given a rate μ chosen by the instructor, each student therefore chooses the rate λ at which he checks the forum to trade off cost and benefit: choosing a high value of λ comes at a large cost in effort (the student needs to be frequently checking the forum), but increases his expected benefit since there is a higher chance that he will see and answer questions in time to receive rewards from the instructor. In addition, if the nature of the question is such that multiple answers are redundant, a higher rate also increases the student's chance of providing a useful, non-redundant answer, since only useful answers are rewarded by the instructor. This gives us a game where the instructor moves first by choosing her rate μ, and the students move second, choosing their rates λ as a best response to μ, and possibly also as a best response to each other's rates, depending on the type of the question.

We explore how an instructor's choice of behavior, modeled by her rate choice μ, affect both her own utility— the difference between the welfare of the student asking the question and the instructor's cost from her rate choice— as well as the overall behavior of the forum, for two representative question types. We obtain an essentially complete characterization of the forum's behavior for both types of questions, which turn out to lead to very different outcomes.

Open-ended questions. Open-ended or discussion-style questions are those where a student can generate value by answering even if other students have already provided answers earlier, so any student who answers the question 'in time' receives a reward. So here the students are not competing with each other, but only 'racing' against the instructor. We find that the student's response rate λ for open-ended questions is *non-monotone* in the instructor's choice of rate μ: as μ increases upward from 0, it initially causes the student to move faster (to answer the question in time), but as μ becomes too large, the reward from answering is no longer worth the cost it would require to make λ sufficiently large, and the student's choice of rate λ begins dropping back toward 0. We note that this non-monotonic dependence of the student's rate λ on the instructor's rate μ has been motivated in earlier writings about educational forums [Andresen 2009; Gray 2004; Hanover Research Council 2009] on qualitative grounds: when students see a highly active instructor on the forum, they conclude that it's either hopeless or not useful to keep up, and they start to drop out.[1] What is striking here is that this qualitative reasoning is captured formally within our model by the optimization performed by strategic agents.

Single-answer questions. Single-answer questions are those with a short well-defined correct answer, where only the first student provides a non-redundant answer and therefore earns a reward. So here, students compete amongst themselves to be the first one to provide an answer, in addition to racing to arrive before the instructor in order to receive their reward. Thus for any announced rate μ by the instructor, the students are playing a simultaneous-move game in which their strategies are their choices of rate. We show that for every μ, this game has a unique symmetric Nash equilibrium in which all students choose the same rate λ^*, and we study this rate as a prediction for how the forum will operate and how the instructor should choose μ. In contrast to the case of open-ended questions where the rate of arrival of the first answer scales linearly with class size n, we find that in equilibrium, the first answer to such single-answer questions arrives at a rate that is almost *independent* of the class size n in equilibrium— each student in a large class slows down his own rate λ^* to an extent great enough that it offsets the gains from having many students trying to answer the question. Moreover, this slow race among the students is very hard for the instructor to influence through her choice of μ: for $n \geq 2$, the equilibrium value λ^* is maximized at $\mu = 0$. This principle that "a large class slows itself down" suggests interesting potential limitations on the power of large classes to provide crowdsourced responses in the single-answer case, and arises from an underlying mechanism that is intuitive: in a large class, each student reasons that his chance of being the first to answer a (single-answer) question is so low that it is not worth checking the class forum very often if it only contains single-answer questions.

We then investigate the instructor's utility $U_I(\mu)$, which is the the utility derived by the question asker minus the instructor's effort cost from checking the forum at rate μ. We discover that as a function of μ, the instructor's utility $U_I(\mu)$ can in fact be *bimodal, with two local maxima*: the optimal utility is achieved either by choosing $\mu \approx 0$ and letting the race among the students do its

[1] Quoting from [Hanover Research Council 2009], "During the discussion process, it is important that instructors continuously manage students' ideas and further facilitate interactions. However, if the online discussion is going well without instructor feedback, it is often best for teachers to wait to jump into the discussion until the students' responses are waning...". Or in the words of an online forum moderator in Gray [Gray 2004], "I was conscious of Burge, Laroque, and Boak's (2000) caution that moderators could not afford to 'get trapped into the 'Atlas syndrome' of holding up the discussion world' and consciously waited for other voices to be heard rather than making postings every day."

work, or by choosing a particular $\mu^* > 0$ in which the students are essentially driven out of the forum and the instructor answers the questions herself. Which of these two local maxima is higher depends on how the instructor values the cost of her own effort. In between, there is an intriguing local minimum in the utility — where, essentially, the instructor is using the forum the "wrong way" — in which the instructor's rate μ suppresses the students' effort in a manner that doesn't increase the overall rate at which the question gets answered.

The fact that the model produces a local minimum between two more efficient modes of operation is intriguing in the way it aligns with empirical evidence from online forums for education. While there are some differences in setting, Andresen discusses the basic trade-off between a mode of operation in which students interact with each other, and a mode in which the instructor principally answers the questions on her own [Andresen 2009]:

> In a study that asks the question of what role an instructor should undertake in an asynchronous discussion forum (sage, guide, or ghost), Mazzolini and Maddison (2003) found that it depends on what the instructor wishes to accomplish. Learner ratings of a course will show that an instructor is more enthusiastic and expert if s/he increases his/her postings. ... However, an instructor that contributes significantly to a discussion tends to decrease the length of discussions (this does not necessarily decrease the quality of the discussion, however) as well as their frequency. What appears to be occurring in this situation is that the instructor can decrease learner-learner interaction because the learners begin to rely on the instructor to answer questions, becoming the expert or sage to 'settle' debates.

Analyzing Mixtures. Finally, we address the question of how an instructor might elicit higher rates from students for single-answer questions. Recall that each student's rate parameter λ, is a global rate that governs how frequently he checks the forum for the class; by definition it is chosen before any particular question is seen. We show that one way to achieve a speedup with class size for single-answer questions is to make sure that the forum contains a mixture of single-answer and open-ended questions— as long as there is any positive probability p that a question in the forum will be an open-ended question (which might be discussion questions posed by the instructor herself), the instructor can choose μ so as to produce a rate λ^* such that the rate of arrival of the first response scales linearly with the size of the class, thus obtaining a speedup with class size which is of the same order as in the case of purely open-ended questions.

1.2. Interpreting the Model

With any stylized model such as the one presented here, it is useful to consider what qualitative conclusions and recommendations it suggests. Our present model is very simple to formulate — it just reflects the balance between the cost of checking the forum frequently and the benefits of getting rewarded for answering questions — and hence the phenomena it captures seem to be arising intrinsically from this tension, rather than from any more complicated modeling artifacts. Moreover, these phenomena, while fairly subtle, align well with evidence from empirical studies of the kind noted earlier.

Our analysis leads to qualitative recommendations (within the stylized setting of our model) for how to optimize the performance of the forum by carefully choosing the instructor's behavior. A first conclusion is that for motivating individual student effort, it is important to find an optimal operating point between motivating students to move quickly and not pushing them to move so quickly that the effort cost leads them to drop out of the system. A second conclusion is that when only a single reward is (meaningfully) available (as in a single-answer question), having many students compete for it may not actually be so productive, because students will scale back the effort they expend in the face of a low probability of receiving the reward. Given this phenomenon, the instructor needs to avoid strategies (such as the local minimum that arises in our single-answer model) where she puts in effort that simply has the effect of suppressing student activity without improving the performance of the forum overall. A final qualitative recommendation from the model is that if the instructor can create reasonable mechanisms by which more than one student is rewarded for the same question, the behavior of the system can transition at least partially toward the faster operation it experiences with open-ended questions; our analysis of mixed-question forums in §5 shows one way to achieve this speedup.

1.3. Related work

As noted above, there has been research into the use of online forums for education focusing on empirical and qualitative evaluation based on experience with such forums [Andresen 2009; Burge et al. 2000; Gray 2004; Hanover Research Council 2009; Mazzolini and Maddison 2003]. There is also a large and growing body of work on general-purpose online Q&A forums such as Y! Answers and StackOverflow starting with Adamic et al. [2008]; most of this literature is empirical, documenting and analyzing usage patterns and user behavior from data or qualitative studies on these forums. There is much less prior work on mathematical models and the design of incentives for online Q&A forums; we discuss these below.

Chen et al. [2009] study Q&A forums such as Y! Answers in a game-theoretic model in which each user has a unique piece of information to contribute to a question, and a user who submits an answer at some time can aggregate all previous answers into her answer. Their work analyzes the equilibria elicited by a best-answer mechanism where the asker awards a point to the best answer received for complements and substitutes information types in terms of the round in which agents supply their answers. They show that a best-answer rule is effective for substitutes information but not for complements information, where an approval-voting scoring rule and a proportional-share scoring rule can enable the most efficient equilibrium. In their model, agents do not have any cost to answering a question, and their strategic choice of when to answer is driven only by the desire to obtain the reward handed out to the best (or best k or value-proportional, depending on the mechanism) answer. Our model differs in two significant aspects, since agents in our model (i) have a cost to the rate at which they check the forum, which is what drives any delays in answering, and (ii) cannot incorporate (or are not rewarded for incorporating) previous contributions into their answer to beat previous answers to obtain a reward. These two differences lead to an entirely different model of strategic behavior and outcomes.

Ghosh and Hummel [2012] focus on the best-answer type of mechanisms used by Q&A sites such as Y! Answers, and investigate *implementability*, asking whether this family of contest-style mechanisms is adequate to elicit optimal outcomes when potential answerers are strategic agents, and how payoffs must be structured to incentivize this desired optimal behavior in terms of the quantity and quality of answers elicited in a simultaneous contribution model. Kumar et al. [2010] study online Q&A sites in the context of network effects in two-sided markets and the rate of arrival of new question-askers (and answerers) as a function of the number of current answerers and question-askers in the system, both empirically and theoretically. They show that the two sides of the market can display radically different patterns, and also investigate competing platforms.

Finally, we note that there is some very recent work on various research challenges arising in the context of online education via massive online open courses (MOOCs). Our work studies models for online forums for education, whether for standard courses taught in 'offline' classrooms or online ones, and intersects with the research on MOOCs to the extent that most MOOCs use Web-based forums as a platform for communication with and between students in the class.

2. MODEL

In this section, we present a game-theoretic model for online forums for education, the central component of which is the strategic choice of *rates* by agents— the presence of an instructor who can monitor answer quality and reward only 'good-enough' contributions allows us to focus on the rate of participation as the strategic choice variable in a model without qualities (or where all answers have equivalent qualities), a simplification which is less justifiable in general online Q&A forums. Students and the instructor choose rates with which to check the online forum for the class, and students are 'rewarded' when the instructor arrives and finds that a student has (usefully) answered a question on the forum. A higher rate incurs a higher *cost*, but can also lead to higher benefit to a student by increasing the probability of arriving before the instructor and thereby earning a reward. This leads to a sequential game where the instructor moves first by choosing her rate, and the students move second, choosing their rate as a best response to the rate chosen by the instructor, and possibly the rate choices of remaining students, as described formally next.

There is a class of n students and an instructor. Questions arrive over a time interval $[0, T]$, and the students and the instructor periodically check the forum and provide answers to these questions, via a process defined below. We refer to the individuals providing the questions as the *question-askers* or simply the *askers*, and the individuals providing the answers as the *students*.[2]

2.1. Rates

Each student and the instructor choose a *rate* at which to check the class's online forum for new content arriving in the interval $[0, T]$. Student j chooses a *rate* λ_j, and repeatedly checks the forum according to a *truncated Poisson process* with rate λ_j — that is, (i) if j has just checked the forum at time t, then he will next check it at time $t + X$, where X is an exponential random variable with rate λ (*i.e.*, with probability density function $f(x) = \lambda e^{-\lambda x}$ for $x \geq 0$); and (ii) j runs this process until the first time that he checks the forum after time T, at which point the process ends. Intuitively, the student's random variable X acts like an "alarm clock" — when the alarm rings at some time t, the student checks the forum and then resets the alarm, so that it next rings at time $t + X$. We will allow a student to choose $\lambda_j = 0$: while one cannot define an exponential random variable of rate 0, we use $\lambda_j = 0$ to simply mean that the student never checks the forum.

We suppose that if a student checks the forum and finds a question that still needs an answer, he can provide one. While obviously not always true, we note that this is not a limiting assumption— our model can be easily extended to a include a probability s_j that a student will have an answer to supply for a random question (the current model corresponds to the special case where $s_j = 1$). Extending the model to include this probability complicates the analysis but not does not change the qualitative conclusions, so we omit it here for clarity given the limited space.

The instructor also chooses her own rate μ at which to check the forum, also according to a truncated Poisson process as described above. When the instructor arrives according to her random process of rate μ, she 'closes' the (new) questions she finds there, supplying an answer (if necessary) to these questions and handing out all appropriate rewards (as described later). Once a question is 'closed', no further rewards can be obtained by answering the question— this could be either because the instructor explicitly disables any further answers to the question, or because no further points will be handed out even if answers can be supplied, or perhaps because no one will read any answers beyond those supplied by the instructor in the case of social-psychological rewards from answering questions.

An instructor is also allowed to choose a rate $\mu = 0$, although with a different interpretation. A zero rate for the instructor means that questions are *never closed* so that a (useful, non-redundant) answer is always assigned a reward no matter how late it arrives

Exponential rates. Exponential random variables have a number of nice properties that makes their analysis particularly tractable. The first is the so-called "memoryless property":

FACT 2.1. *Let X be an exponential random variable with rate λ. Conditioned on the information $X \geq t_0$, the random variable $X - t_0$ is also an exponential random variable of rate λ.*

That is, if we know that the alarm has not rung until time t_0, then starting at time t_0 it behaves as though the alarm had just been freshly reset at t_0. Later in this section, we will use this property

[2]For simplicity, we will consider the askers and the students here as two disjoint populations, but one can also model the case in which the asker on each individual question is drawn from the population of students.

to show that we can reduce the problem of analyzing participants' behavior over the entire time interval to the analysis of their behavior on a per-question basis.

We also note two other standard facts about exponential random variables that will be very useful in parts of the analysis.

FACT 2.2. *If X_1 and X_2 are exponential random variables with rates λ_1 and λ_2, then* $\Pr[X_1 < X_2] = \dfrac{\lambda_1}{\lambda_1 + \lambda_2}$.

FACT 2.3. *If $X_1, \ldots, X_n$ are exponential random variables with rates $\lambda_1, \ldots, \lambda_n$, then $\min(X_1, \ldots, X_n)$ is an exponential random variable with rate $\lambda_1 + \cdots + \lambda_n$.*

2.2. Utilities

The utility is the difference between an agent's (student's or instructor's) benefit and her cost, which we describe next.

Benefits. Rewards are allocated to students by the instructor on a per-question basis. Students are 'rewarded' when the instructor arrives and ascertains that a student has (correctly and usefully) answered a new question on the forum— this reward could either be in the forms of points awarded by the instructor that count towards the student's grade in the class, or because the student derives a social-psychological reward from the instructor's approval or endorsement[3], or from providing a useful answer to her fellow students. Once the instructor arrives and 'closes' a question, no further rewards can be obtained from answering that question— as discussed before, this could be either because the student derives value only when the instructor sees that she has usefully answered a question (either because points are only awarded for such answers, or because the student's reward is the psychological value of the instructor's approval), or because the student derives value from having her fellow students read her answer and no answers beyond the instructor's are read by other students. We note that our model is agnostic on the question of what causes a student to feel rewarded from answering a question, and applies to a spectrum of possible underlying reasons for rewards.

We use v^* to denote the value of each individual reward to a student. Let the random variable R_{jq} be equal to 1 if student j receives a reward for answering question q, and be equal to 0 otherwise. Whether an answer generates value, and therefore merits a reward, is different for the two types of questions discussed in §1. In our definition of *open-ended* or discussion-style questions, an answer adds value even if other answers appear earlier, so that any student who answers before the instructor closes the question can receive a reward (recall our assumption, discussed earlier, that students have correct answers to questions). On the other hand, for *single-answer* questions, only the first student to see the question will answer, because subsequent answers would be redundant (which are not rewarded, so that we can assume that students don't bother to supply redundant answers). We derive the specific forms of the reward functions for these two types of questions in §3 and §4 respectively.

We assume that the student incurs a fixed cost ξ for writing an answer to a question; because a student will only write an answer when the reward received for it is at least ξ, it makes sense to think not about the pure reward v^* but the *net reward* $v = v^* - \xi$.

Costs. We assume that for each student, checking the forum for new content incurs an effort cost of α each time the forum is checked, whether or not there is new content to react to. (If the student does react to new content, this incurs the effort cost of ξ noted above, which is folded into the net reward $v = v^* - \xi$.)

Similarly, the instructor incurs a cost of β each time she checks the forum.

Arrival of Questions and Overall Utility. Questions arrive over the time interval $[0, T]$ according to a Poisson process with rate 1 — that is, we imagine questions being generated by a rate-1 Poisson

[3]Piazza, for instance, explicitly offers an 'instructor endorsement' feature on its forums.

process defined over the positive real numbers, and then we only keep the subset of questions that arrive in $[0, T]$. (The choice of 1 here provides the units of the time scale in which all other rates are measured.) Let Q denote this set of questions, and let C_j denote the set of times at which student j checks the forum. Then student j's overall utility (recall that v is the net reward) is the expectation of

$$\sum_{q \in Q} vR_{jq} - \sum_{h \in C_j} \alpha,$$

reflecting the net rewards from each question and the cost for each checking of the forum.

2.3. Reduction to the Case of a Single Question

Because of the memoryless property of the exponential distribution, we can effectively consider each question that appears on the forum over the duration of the course in isolation, as follows. Recall that each student (and the instructor) runs his process from 0 until the first time that his exponentially distributed "alarm clock" rings after time T. Thus, since each question q appears before time T, all participants' alarm clocks are still running when q arrives, and by the memoryless property we can thus imagine that these alarm clocks are all started at this point of arrival. It follows that for a student j and any two questions q and q', the random variables R_{jq} and $R_{jq'}$ have the same distribution; let R_j denote a random variable with this shared distribution.

Hence, by Wald's Theorem, we can write student j's overall utility as

$$E\left[\sum_{q \in Q} vR_{jq} - \sum_{h \in C_j} \alpha\right] = vE\left[|Q|\right] E\left[R_j\right] - \alpha E\left[|C_j|\right].$$

Since the questions are arriving at rate 1 over the interval $[0, T]$, we have $E\left[|Q|\right] = T$. For $E\left[|C_j|\right]$, we observe that the times at which student j checks the forum consist of draws from a rate-λ_j Poisson process over $[0, T]$, followed deterministically by one additional check of the forum (the first one that the process generates after T). Hence we have $E\left[|C_j|\right] = \lambda_j T + 1$, and so j's utility is

$$vE\left[R_j\right] T - \alpha(\lambda_j T + 1) = (vE\left[R_j\right] - \alpha\lambda_j)T - \alpha.$$

T and α are independent of student j's choice of λ_j, but $E\left[R_j\right]$ depends on λ_j, since a faster rate of checking the forum leads to a greater chance of receiving a reward. We can thus think of j as strategically choosing λ_j to optimize the objective function $vE\left[R_j\right] - \alpha\lambda_j$, and this is the basic form of the problem we consider.

We make one further simplication — we will focus in this exposition on the case in which the net reward $v = 1$, so the objective function representing student j's utility maximization problem is

$$E\left[R_j\right] - \alpha\lambda_j.$$

This is purely for simplicity, and arbitrary values of v are handled completely analogously.

The Instructor's Utility. For the instructor's utility, we can use a direct analogue of the reduction above and focus on her utility for a single question. The instructor's utility on a single question is the difference between her time-discounted value from the answers contributed to the question and her effort cost $\beta\mu$ from checking the forum at a rate μ; as with the students, the form of this utility depends on on whether the question is open-ended or a single-answer question. The instructor derives benefit from the asker of the question getting an answer (for single-answer questions) or answers (in discussion questions) as quickly as possible. We model this by saying that the value of a (useful, non-redundant) answer is discounted exponentially, so that if it arrives at time t after the question was asked, it has value $e^{-\delta t}$ for a parameter $\delta > 0$. Thus, if a question receives k answers — potentially including the instructor's own answer — at times $t_1, t_2, \ldots, t_k$ (note that $k = 1$ for a single-answer question[4]), then the instructor's utility in this realization of the process is

$$\sum_{i=1}^{k} e^{-\delta t_i} - \beta\mu.$$

In both §3 and §4, we will be considering the expected value of this expression, when the times t_i arise from the random process of checking the forum according to the rates λ chosen (in equilibrium) by the students, and the rate μ chosen by the instructor.

2.4. Summary of Parameters

The reduction above allows us to consider the problem of equilibrium rates in the context of a single question that arrives at time 0; all strategies and utilities will be defined in the context of this single question, keeping in mind that this in fact captures the behavior of the system for the full set of questions arriving over $[0, T]$, Beyond the number of students n, the model thus has three basic parameters: α, β, and δ, each with a simple interpretation. The parameters α and β represent how the students and the instructor value the effort they have to put into the forum (for example, a large α means students view checking the forum as a costly activity, which affects how they set λ). The parameter δ models the "impatience" of the asker: a large value of δ means that answers that take longer to come in create significantly less value than answers that come in right away.

3. OPEN-ENDED QUESTIONS

We begin by considering open-ended questions, in which the question-asker benefits from all answers that come in, and hence students will continue to contribute answers until the instructor closes the question. As a result, the students are not competing against each other to provide answers quickly; each student is instead optimizing his choice of rate λ based solely on the instructor's choice of rate μ.

Our first result shows that the student's response rate λ for open-ended questions is *non-monotone* in the instructor's choice of rate μ, and is maximized at a specific intermediate value of μ: as μ increases upward from 0, it initially causes the student to want to move faster (to answer the question before it closes), but as μ becomes too large, the reward from answering is no longer worth the cost it would require to make λ sufficiently large, and the student's choice of rate λ begins dropping back toward 0.

THEOREM 3.1. *A student's optimal choice of rate λ, as a function of the instructor's rate μ, is given by*

$$\lambda(\mu) = \max\{0, \sqrt{\frac{\mu}{\alpha}} - \mu\}.$$

$\lambda(\mu)$ increases on $\mu \in [0, \frac{1}{4\alpha}]$ and decreases beyond, attaining its unique maximum at $\mu^* = \frac{1}{4\alpha}$.

PROOF. Consider a student with rate x and an instructor with rate μ. The student receives a reward if she arrives and answers the question before the instructor arrives; by Fact 2.2, this event has probability $x/(x + \mu)$. So the student's payoff from choosing rate x is

$$U_S(x) = \frac{x}{x + \mu} - \alpha x.$$

[4]As discussed previously, students do not have an incentive to contribute redundant answers, nor do redundant answers generate value to the asker.

Note that U is concave in x, so to maximize U we set the derivative $U'_S(x) = \frac{\mu}{(x+\mu)^2} - \alpha$ to zero and solve for x, which gives

$$x^* = \sqrt{\frac{\mu}{\alpha}} - \mu.$$

This is the best response for a student when $x^* \geq 0$. If $x^* < 0$, then the best response is to choose rate 0, since $U_S(0) = 0$ and $U'_S(0) < 0$ in this case (U_S is concave in x so U'_S is decreasing, and $U'_S(x^*) = 0$ for $x^* < 0$)— this corresponds to the instructor choosing too high a rate μ, specifically $\mu > \alpha^{-1}$ at which point the student simply drops out. Therefore, the student's best response rate when the instructor chooses rate μ is

$$\lambda(\mu) = \max\{0, \sqrt{\frac{\mu}{\alpha}} - \mu\}. \quad (1)$$

Observe that $\lambda(\mu)$ is *non-monotone* in μ: the derivative of $\lambda(\mu)$ with respect to μ, for $\mu \leq \frac{1}{\alpha}$, is $\lambda'(\mu) = \frac{1}{2\sqrt{\alpha\mu}} - 1$, which is positive for $\mu \leq \frac{1}{4\alpha}$ so that λ is increasing for $0 \leq \mu \leq \frac{1}{4\alpha}$, and negative for $\mu > \frac{1}{4\alpha}$, so that $\lambda(\mu)$ decreases with as the instructor increases her rate beyond $\frac{1}{4\alpha}$. In particular, we can ask what value of μ causes the student to answer as quickly as possible; this maximum is precisely where the derivative is 0 (note that λ is concave in μ), which is $\mu^* = \frac{1}{4\alpha}$. □

Note that $\lambda(\mu^*) = \mu^*$, so that when the instructor chooses a rate to maximize the rate chosen as a best response by the students, the instructor and students all have the *same rate* in equilibrium. Note also that at a given $\mu \in [0, \alpha^{-1}]$, the probability the student provides an answer before the instructor closes the question is equal to

$$\frac{\lambda(\mu)}{\lambda(\mu) + \mu} = \frac{\sqrt{\mu\alpha^{-1}} - \mu}{\sqrt{\mu\alpha^{-1}}} = 1 - \sqrt{\alpha\mu},$$

which is decreasing in μ. So for μ close to 0, the student's probability of providing an answer before the question closes — at his optimal rate $\lambda(\mu)$ — is close to 1; this probability decreases to exactly $1/2$ at the optimal $\mu^* = 1/(4\alpha)$, and then continues decreasing to 0 as μ reaches $1/\alpha$.

Overall utility. We now consider how the instructor should choose a rate μ to maximize her utility — recall from Section 2 that this utility consists of a benefit equal to the exponentially discounted arrival times of answers, minus a cost $\beta\mu$ due to the instructor's choice of μ.

To begin with, let's consider the instructor's utility in the case where she is the only individual providing an answer to the question. If she chooses a rate of x, then the expected benefit of her answer is the expected value of its discounted arrival time, which is

$$\int_0^\infty e^{-\delta t} x e^{-xt} dt = \frac{x}{x + \delta}. \quad (2)$$

Equation (2) has a very clean interpretation — the expected value of the discounted arrival time is equivalent to imagining that the question-asker will give up after an amount of time drawn from an exponential distribution with rate δ (the discount parameter), and the benefit of the answer is the probability that it arrives before the asker gives up. Recall that the discount rate δ models the "impatience" of the asker — a large δ corresponds to a fast rate of decay in value on the left-hand of Equation (2), and it corresponds to the asker giving up quickly on the right-hand side of (2).

Now let's consider the instructor's full utility function. There is a benefit from each answer provided by a student, plus the benefit of the instructor's own answer, minus the instructor's effort cost. The benefit of a student's answer (provided at rate y) is equal to the probability that it arrives both before the asker gives up (at rate δ) and before the instructor closes the question (at rate μ), since the student will not answer a closed question:

$$\int_0^\infty e^{-\delta t} e^{-\mu t} y e^{-yt} dt = \frac{y}{y + \delta + \mu}.$$

When the instructor chooses a rate μ, the students will each choose their best response rate $\lambda(\mu)$ (as in (1)), and the instructor's goal is to maximize her utility function

$$U_I(\mu) = n \cdot \frac{\lambda(\mu)}{\lambda(\mu) + \mu + \delta} + \frac{\mu}{\mu + \delta} - \beta\mu. \tag{3}$$

We now consider the maximum utility achievable in the case when n is large, so that the first term, which is linear in n, is the dominant one in the expression for U_I in (3):

THEOREM 3.2. *When* $\mu^* = \alpha(\sqrt{\delta^2 + \delta\alpha^{-1}} - \delta)^2$, *we have* $U_I(\mu^*) = b^* n + o(n)$ *where* $b^* = 1 - 2\alpha(\sqrt{\delta^2 + \delta\alpha^{-1}} - \delta)$*; also, there is no* μ *and* $b > b^*$ *for which* $U_I(\mu) = bn + o(n)$*. Thus, the rate* μ^* *maximizes the instructor's utility to within a term that is sublinear in* n*.*

Before presenting the proof, it is worth understanding how the quantities μ^* and b^* behave, particularly as we vary the impatience δ of the question-asker. Note that there is a trade-off between speed and the number of answers received in the utility $U_I(\mu)$— when μ is very small, students arrive at a much slower rate but more students are likely to find the question open, so that we obtain a larger number of answers, although at later times. As μ increases, students also increase their response rates $\lambda(\mu)$ and arrive faster, but fewer answers are provided since the question closes earlier.

When the asker's impatience δ is close to 0, the optimal value μ^* is close to 0 and b^* is close to 1: that is, with a very patient asker, the instructor can afford to move slowly to ensure that almost every student contributes close to the full benefit of 1 via their answer.

As δ goes to infinity, it follows from Taylor's Theorem that $\sqrt{\delta^2 + \delta\alpha^{-1}} - \delta$ converges to $1/(2\alpha)$ from below. Hence μ^* converges to $1/(4\alpha)$ and b^* converges to $1 - (2\alpha)(1/(2\alpha)) = 0$ from below. This makes sense: as δ increases, so that the asker becomes more impatient, the instructor needs to trade off quantity against speed to induce the students to move more quickly, so she increases her rate up to the μ^* that, by Theorem 3.1, maximizes the rate of each student. However, each student can only be induced to move at a maximum rate of $1/(4\alpha)$, which is independent of δ, and so as δ goes to infinity the benefit contributed by each student is converging to 0.

We now proceed with the proof of Theorem 3.2.

PROOF. In the expression for U_I in (3), the second and third terms are $o(n)$, so it suffices to find a choice of μ that maximizes the coefficient on n in the first term. Such a choice of μ will lie in the interval $[0, \alpha^{-1}]$, since it needs to cause the students to choose a non-zero rate.

By Theorem 3.1, when the instructor chooses a rate $\mu \in [0, \alpha^{-1}]$, each student chooses the rate $\lambda(\mu) = \sqrt{\mu\alpha^{-1}} - \mu$. Applying this choice of λ in Equation (3), we get $U_I(n) = B(\mu) \cdot n + o(n)$ where

$$B(\mu) = \frac{\sqrt{\mu\alpha^{-1}} - \mu}{\sqrt{\mu\alpha^{-1}} + \delta}.$$

We thus wish to maximize $B(\mu)$ over the interval $[0, \alpha^{-1}]$. The function $B(\mu)$ is concave in μ over this interval, with a unique maximum where $B'(\mu) = 0$. Setting this derivative to 0 yields a quadratic equation in $\sqrt{\mu}$:

$$\alpha^{-1/2}\mu + 2\delta\sqrt{\mu} - \delta\alpha^{-1/2} = 0$$

and hence

$$\sqrt{\mu} = \sqrt{\alpha}(\sqrt{\delta^2 + \delta\alpha^{-1}} - \delta).$$

At this value of μ, we have

$$B(\mu) = \frac{\sqrt{\delta^2 + \delta\alpha^{-1}} - \alpha(\sqrt{\delta^2 + \delta\alpha^{-1}} - \delta)^2}{\sqrt{\delta^2 + \delta\alpha^{-1}}},$$

which simplifies to $B(\mu) = 1 - 2\alpha(\sqrt{\delta^2 + \delta\alpha^{-1}} - \delta)$. □

4. SINGLE-ANSWER QUESTIONS

Single-answer questions, as the name suggests, are questions with one well-defined correct answer, such as "In the homework due tomorrow, were we allowed to assume that x is positive?", or "What is the most abundant isotope of carbon?". In single-answer questions, only the first student to see the question will answer, because subsequent answers would be redundant[5].

For single-answer questions, therefore, students are implicitly competing to be the first to provide an answer— so for any particular rate μ chosen by the instructor, students play a simultaneous-move game amongst themselves in which their strategies are their choices of rate. We show that this game has a unique symmetric Nash equilibrium in which all students choose the same rate $\lambda^*(\mu)$, which we can study as a prediction for how the forum will operate.

Note that even when $\mu = 0$ (so that the instructor doesn't close the question, but only grades the outcome of the race once the class is over), the equilibrium rate λ^* is non-trivial for single-answer questions, due to the competition among the students. We find that at this equilibrium, the first answer to the question arrives at a speed that is almost independent of the class size n. (There is an additive dependence on n that grows proportionally to $1/n$.) Moreover, this slow race among the students is very hard for the instructor to influence through her choice of μ: for $n \geq 2$, we will find that the equilibrium value $\lambda^*(\mu)$ is maximized at $\mu = 0$.

Thus, with a single-answer question, each student in a large class slows down his personal rate, to an extent great enough that it offsets the gains from having many students trying to answer the question. This is a sharp contrast from the case of open-ended questions, where the students are choosing rates without regard to the behavior of the other students, and hence the rate of the fastest answer in that case scales linearly in n. This principle that "a large class slows itself down" suggests interesting potential limitations on the power of large classes to provide crowdsourced responses in the single-answer case, and it arises from an underlying mechanism that is intuitive: in a large class, each student reasons that his chance of being the first to answer a (single-answer) question is so low that it is not worth checking the site very often.

THEOREM 4.1. *For any choice of μ by the instructor, there is a unique symmetric equilibrium $\lambda(\mu)$ for the students. In this equilibrium:*

— *The students don't participate when $\mu \geq \alpha^{-1}$; in this case the instructor answers the question on her own.*
— *When $\mu < \alpha^{-1}$, we have $\lambda(\mu) = \frac{n-1}{2\alpha n^2} - \frac{\mu}{n} + \sqrt{\frac{(n-1)^2}{4\alpha^2 n^4} + \frac{\mu}{\alpha n^3}}$, so that the students do participate along with the instructor, but each at a rate of $\Theta(1/n)$ as a function of the class size n.*
— *For $n \geq 2$, the maximum student rate $\lambda(\mu)$ in equilibrium is achieved when the instructor does not participate ($\mu = 0$).*

PROOF. Consider a student i with rate x: i receives a reward when she arrives before all the remaining $n-1$ students and the instructor. The payoff of i, when the other $n-1$ students all choose rate λ and the instructor chooses rate μ, is (using Fact 2):

$$U(x, \lambda) = \frac{x}{x + (n-1)\lambda + \mu} - \alpha x.$$

[5]Recall that we assumed in our model that a student who sees a question can supply the correct answer to it, and discussed generalizing this assumption in §2. Recall also that students are rewarded only for valuable answers, so we assumed that students do not bother to provide redundant or useless answers.

For x to be a best response, $U(x,\lambda,\mu)$ must be maximized at x. Note that U is concave in x, so setting the derivative to zero is a necessary and sufficient condition for an interior maximum. Taking the derivative and equating to zero, we have $\frac{(n-1)\lambda+\mu}{(x+(n-1)\lambda+\mu)^2} = \alpha$, which gives us $x = \sqrt{\frac{(n-1)\lambda+\mu}{\alpha}} - ((n-1)\lambda+\mu)$. Since we seek a symmetric equilibrium, we want $x=\lambda$. This gives us

$$\sqrt{\frac{(n-1)\lambda+\mu}{\alpha}} = n\lambda+\mu,$$

yielding the following quadratic equation in the variable λ:

$$\alpha n^2\lambda^2 + \lambda(2n\mu\alpha - (n-1)) + (\alpha\mu^2 - \mu) = 0.$$

The solutions to this quadratic are, respectively,

$$(\lambda_1(\mu),\ \lambda_2(\mu)) = \frac{n-1}{2\alpha n^2} - \frac{\mu}{n} \pm \sqrt{\frac{(n-1)^2}{4\alpha^2 n^4} + \frac{\mu}{\alpha n^3}}.$$

Since $\sqrt{\frac{(n-1)^2}{4\alpha^2n^4} + \frac{\mu}{\alpha n^3}} \geq \frac{n-1}{2\alpha n^2}$, it follows that $\lambda_2(\mu) \leq 0$ for any μ. This cannot be the best response since λ is the rate of an exponential and hence must be positive. So the best response of the students when the instructor chooses a rate μ is either $\lambda_1(\mu)$ if $\lambda_1(\mu) \geq 0$, or else it is 0:

$$\lambda(\mu) = \max\{0, \lambda_1(\mu)\} = \max\{0, \frac{n-1}{2\alpha n^2} - \frac{\mu}{n} + \sqrt{\frac{(n-1)^2}{4\alpha^2 n^4} + \frac{\mu}{\alpha n^3}}\}.$$

Now we investigate the function $\lambda_1(\mu)$. At $\mu = 0$, $\lambda_1(0) = \frac{n-1}{\alpha n^2}$. The derivative of λ_1 with respect to μ tells us how the best response rate chosen by the students changes as the instructor increases her rate up from 0, if the best response is to participate:

$$\frac{d}{d\mu}\lambda_1(\mu) = -\frac{1}{n} + \frac{1}{\alpha n^3}\frac{1}{2\sqrt{\frac{(n-1)^2}{4\alpha^2n^4} + \frac{\mu}{\alpha n^3}}} = \frac{1}{n}\left[\frac{1}{2\alpha n^2\sqrt{\frac{(n-1)^2}{4\alpha^2n^4} + \frac{\mu}{\alpha n^3}}} - 1\right].$$

Note that $\frac{d}{d\mu}\lambda_1(\mu)$ is decreasing in μ. At $\mu = 0$, the derivative evaluates to

$$\left.\frac{d}{d\mu}\lambda_1(\mu)\right|_0 = \frac{1}{n}\left[\frac{1}{n-1} - 1\right] \leq 0$$

for all $n \geq 2$. Since the derivative is decreasing in μ, and non-positive at $\mu = 0$, this means that $\lambda_1(\mu)$ is maximized at $\mu = 0$, as long as $n \geq 2$. Now, the value of λ_1 at $\mu = 0$ is $\lambda(0) = \frac{n-1}{\alpha n^2} > 0$, so the value of μ that maximizes $\lambda(\mu) = \max\{0, \lambda_1(\mu)\}$ is $\mu^* = 0$, as long as $n \geq 2$. That is, if there are at least two students in the class (and we need at least two students for there to be a race between students to first supply the correct answer), the instructor can maximize the rate of response of the students by choosing the lowest possible rate for herself, *i.e.*, by setting $\mu = 0$.

Note that for any n, $\lambda_1(\mu)$ becomes 0 at $\mu_0 = \frac{1}{\alpha}$. (Since the derivative of $\lambda_1(\mu)$ is decreasing in μ and non-positive at $\mu = 0$, and $\lambda_1(0) > 0$, $\lambda_1(\mu)$ crosses zero exactly once for $\mu \geq 0$).

Therefore, the students' best response rate $\lambda(\mu)$ behaves as follows:

$$\lambda(\mu) = \begin{cases} \frac{n-1}{2\alpha n^2} - \frac{\mu}{n} + \sqrt{\frac{(n-1)^2}{4\alpha^2 n^4} + \frac{\mu}{\alpha n^3}}, & \text{for } \mu \in [0, \frac{1}{\alpha}] \\ 0, & \text{for } \mu \geq \frac{1}{\alpha}. \end{cases}$$

□

Finally, we consider what the instructor should do to maximize her utility, and here too we find a complex but fairly intuitive phenomenon at work. Recall that the instructor's utility U_I is based on the speed at which the question is answered (either by her or by one of the students), minus the effort cost created by checking the forum at rate μ. When we study this utility $U_I(\mu)$ as a function of μ, we discover that for a small enough effort coefficient β on the part of the instructor, $U_I(\mu)$ is in fact *bimodal*, with two local maxima: the optimal utility is achieved either by setting $\mu = 0$ and letting the race among the students do its work, or by choosing a particular $\mu^* > 0$ in which the students are essentially driven out of the forum and the instructor answers the questions herself. Which of these two local maxima is higher depends on how the instructor values the cost of her own effort. In between, there is an intriguing local minimum in the utility — essentially, when the instructor is using the forum the "wrong way" — in which the instructor's choice of rate μ suppresses the students' effort in a manner that doesn't increase the overall rate (*i.e.*, including both the instructor and the students) at which the question gets answered.

To begin the analysis of this phenomenon, we write down the utilities involved. For single-answer questions, the asker derives utility from the first (and only) answer to her question, which is discounted as before by the time at which this first answer arrives. Again, if the asker discounts time by an exponential factor δ, the expected utility from an answer arriving at time t is $e^{-\delta t}$. The time of arrival of the first answer, when the instructor chooses a rate of μ (and answers the question if she arrives first), and the students choose a rate $\lambda(\mu)$, is exponentially distributed with rate $n\lambda(\mu) + \mu$. So the expected utility to the asker is

$$\int_0^\infty e^{-\delta t}(n\lambda(\mu) + \mu)e^{-(n\lambda(\mu)+\mu)t}dt = \frac{n\lambda(\mu) + \mu}{n\lambda(\mu) + \mu + \delta}.$$

The instructor would like to maximize her own utility which is the difference between the asker's utility, and her own cost to choosing a rate μ, *i.e.*, to choose a rate μ that maximizes

$$U_I(\mu) = \frac{n\lambda(\mu) + \mu}{n\lambda(\mu) + \mu + \delta} - \beta\mu.$$

We now formalize the bimodal properties of U_I in this setting.

THEOREM 4.2. *For n sufficiently large in terms of α, β, and δ, the following holds.*

— *If $\beta < \dfrac{\delta}{(\delta + \alpha^{-1})^2}$, the instructor's utility function $U_I(\mu)$ has two local maxima: one at $\mu = 0$, and the other at a point $\mu > \alpha^{-1}$. In between, U_I has a local minimum at $\mu = \alpha^{-1}$. There is a threshold cost β^* such that the second local maximum has greater utility than the first if and only $\beta < \beta^*$.*

— *If $\beta > \dfrac{\delta}{(\delta + \alpha^{-1})^2}$, the instructor's utility function $U_I(\mu)$ decreases monotonically in μ, with a single maximum at $\mu = 0$.*

PROOF. First we consider what happens if the instructor chooses a rate $\mu < \alpha^{-1}$. Recall from Theorem 4.1 that the students will participate for such μ, and so the combined rate will be $n\lambda(\mu)+\mu$. We use $\phi(\mu)$ to denote this combined rate; so by the definition of the instructor's utility we have $U_I(\mu) = \frac{\phi(\mu)}{\phi(\mu)+\delta} - \beta\mu$, and from Theorem 4.1, we have

$$\phi(\mu) = \frac{n-1}{2\alpha n} + \sqrt{\frac{(n-1)^2}{4\alpha^2 n^2} + \frac{\mu}{\alpha n}}.$$

Differentiating the instructor's utility, we get $U_I'(\mu) = \frac{\delta\phi'(\mu)}{(\phi(\mu)+\delta)^2} - \beta$. Now,

$$\phi'(\mu) = \frac{1}{\alpha n} \cdot \frac{1}{2\sqrt{\frac{(n-1)^2}{4\alpha^2 n^2} + \frac{\mu}{\alpha n}}},$$

which goes to 0 as n diverges, so that the first term in $U'_I(\mu)$ also goes to 0 in n. Thus we see that for sufficiently large n, $U'_I(\mu) < 0$ for $\mu < \alpha^{-1}$. It follows that $\mu = 0$ is a local maximum of the instructor's utility when n is large enough.

Next we consider what happens if the instructor chooses a rate $\mu \geq \alpha^{-1}$. In this case the students will not participate so that the combined rate $\phi(\mu) = \mu$, and so the instructor's utility is given by $U_I(\mu) = \frac{\mu}{\mu+\delta} - \beta\mu$. If we consider this as a function over all $\mu \geq 0$, it is concave, with a unique local maximum at

$$\mu^* = \max\left\{0, \sqrt{\frac{\delta}{\beta}} - \delta\right\}.$$

But this is the correct expression for the instructor's utility only if $\mu \geq \alpha^{-1}$, *i.e.*, once the students have dropped out, so we need to ask when this local maximum μ^* exceeds α^{-1}. The condition for $\mu^* > \alpha^{-1}$ is $\sqrt{\frac{\delta}{\beta}} - \delta > \frac{1}{\alpha}$ and hence $\beta < \frac{\delta}{(\delta + \alpha^{-1})^2}$. We thus have two cases (again keeping in mind that n is sufficiently large):

— When $\beta < \frac{\delta}{(\delta + \alpha^{-1})^2}$, the instructor's utility has a second local maximum μ^* (in addition to $\mu = 0$) at a value of μ beyond α^{-1}— this local maximum corresponds to the instructor answering the question alone with no help from the class, at a rate which optimally trades off her cost to checking the forum and her benefit from the asker receiving an answer quickly.

— On the other hand, when $\beta > \frac{\delta}{(\delta + \alpha^{-1})^2}$, then $U_I(\mu)$ is monotonically decreasing as soon as $\mu \geq \alpha^{-1}$— this is because the concave function $\frac{\mu}{\mu+\delta} - \beta\mu$ is decreasing beyond its maximum, *i.e.*, for $\mu \geq \mu^*$, and $\mu^* < \alpha^{-1}$. Hence $\mu = 0$ is the only local maximum for this range of β.

Finally, in the case when $\beta < \frac{\delta}{(\delta + \alpha^{-1})^2}$, we can ask which of the two local maxima 0 or μ^* achieves a higher utility. The utility at the local maximum $\mu = 0$ is

$$U_I(0) = \frac{\phi(\mu)}{\phi(\mu) + \delta} = \frac{\frac{n-1}{\alpha n}}{\frac{n-1}{\alpha n} + \delta} \leq \frac{1}{1 + \alpha\delta}.$$

The utility at the other maximum μ^* is given by

$$U_I(\mu^*) = (1 - \sqrt{\beta\delta})^2.$$

Now we observe that as β goes to 0, so that the instructor's effort becomes less and less costly, the utility $U_I(\mu^*)$ converges to 1 while $U_I(0)$ remains bounded below 1.

Thus there is a β^* such that when $\beta < \beta^*$, the second local maximum, when the instructor answers the question on her own, has the higher utility. In the limit as $n \to \infty$, the utility $U_I(0)$ is converging to $\frac{1}{1 + \alpha\delta}$, and so the cross-over point where $U_I(0) = U_I(\mu^*)$ is converging to the solution of

$$(1 - \sqrt{\beta\delta})^2 = \frac{1}{1 + \alpha\delta},$$

which is

$$\beta^* = \frac{1}{\delta}(1 - \frac{1}{\sqrt{1 + \delta\alpha}})^2.$$

□

5. MIXED FORUMS

In this section, we address the question of how an instructor might elicit higher rates from students for single-answer questions, and show that one way to offset some of the slow behavior arising from single-answer questions is by ensuring that the forum contains a mixture of single-answer and open-ended questions.

We imagine that each question is open-ended with probability $p > 0$ and single-answer with probability $1 - p$. Each student's λ parameter, however, is a global rate that governs how frequently he checks the site; by definition it must be chosen before any particular question is seen. When the student checks the site (according to rate λ), he may sometimes find an open-ended question, and sometimes a single-answer question.

As before, we find that there is a unique symmetric equilibrium λ^* for any choice of μ. Our main result here is that the instructor can choose μ so as to elicit a rate λ^* that is roughly p times what it is in the case of purely open-ended questions, and this is tight for large n. In this way, the rate at which the $1 - p$ fraction of single-answer questions get answered benefits significantly from the presence of even a small positive fraction $p > 0$ of open-ended questions — they are essentially sped up by a linear amount (in n) when p is any constant.

THEOREM 5.1. *Let p be the probability that a question is open-ended. The instructor can choose a rate μ such that the equilibrium student rate λ^* satisfies $\lambda^* \geq \frac{p}{4\alpha}$. Moreover this is essentially tight in the limit, in the following sense: For any $\varepsilon > 0$, there exists a large enough n so that with at least n students, there is no μ the instructor can choose so as to achieve $\lambda^* \geq \frac{(1+\varepsilon)p}{4\alpha}$.*

PROOF. Suppose the instructor chooses rate μ; we want to determine the best response rate λ chosen by the students. Using the payoffs derived in §3 and §4, the payoff to a student who chooses rate x when the instructor has rate μ and the remaining $n - 1$ students choose rate λ is

$$U(x) = p\frac{x}{x+\mu} + (1-p)\frac{x}{x+(n-1)\lambda+\mu} - \alpha x.$$

As usual, for x to be a best response, a necessary condition is for $U'(x)$ to be zero, which gives us

$$p\frac{\mu}{(x+\mu)^2} + (1-p)\frac{(n-1)\lambda+\mu}{(x+(n-1)\lambda+\mu)^2} - \alpha = 0.$$

Since we are looking for a symmetric equilibrium, we set $x = \lambda$, which gives us the following equation in λ:

$$p\frac{\mu}{(\lambda+\mu)^2} + (1-p)\frac{(n-1)\lambda+\mu}{(n\lambda+\mu)^2} = \alpha. \tag{4}$$

Note that (4) can be interpreted as a bivariate equation in the two variables λ and μ; for a given value of μ, if there is a solution $\lambda(\mu) \geq 0$ to this equation, then $\lambda(\mu)$ is a best response to μ.

We first identify a choice of μ for which the equilibrium value $\lambda(\mu)$ is at least $\frac{p}{4\alpha}$. It turns out that this can be achieved using a μ for which $\lambda(\mu) = \mu$, as follows. In order to have an equilibrium in which $\lambda(\mu) = \mu$, we simply need a pair (μ, λ) satisfying Equation (4) for which $\mu = \lambda$. Setting $\mu = \lambda$ in (4), we get

$$p\frac{\lambda}{4\lambda^2} + (1-p)\frac{n\lambda}{(n+1)^2\lambda^2} = \alpha,$$

leading to

$$\lambda = \frac{p + \frac{4}{n+1}(1-p)}{4\alpha} \geq \frac{p}{4\alpha}, \tag{5}$$

as desired.

This gives us a lower bound on the fastest possible equilibrium λ, by demonstrating the value of λ achieved at a particular carefully-chosen equilibrium. Next we will show that as n becomes large, this lower bound is essentially the best possible.

To do this, fix any constant $\varepsilon > 0$, choose σ small enough that $(1-\sigma)(1+\varepsilon) > 1$, and then choose n large enough in terms of ε, p, and σ so that $\frac{4}{pn} < \sigma$. Now suppose by way of contradiction that for some choice of μ, there is a pair (μ, λ) that solves the bivariate equation (4) with $\lambda \geq \frac{(1+\varepsilon)p}{4\alpha}$.

Then we have

$$\begin{aligned}
\alpha &= p\frac{\mu}{(\lambda+\mu)^2} + (1-p)\frac{(n-1)\lambda+\mu}{(n\lambda+\mu)^2} \\
&\leq p\frac{\mu}{(\lambda+\mu)^2} + (1-p)\frac{n\lambda+\mu}{(n\lambda+\mu)^2} \\
&= p\frac{\mu}{(\lambda+\mu)^2} + (1-p)\frac{1}{(n\lambda+\mu)} \\
&\leq p\frac{\mu}{(\lambda+\mu)^2} + \frac{1}{n\lambda} \\
&\leq p\frac{\mu}{(\lambda+\mu)^2} + \frac{4\alpha}{pn} \\
&< p\frac{\mu}{(\lambda+\mu)^2} + \sigma\alpha,
\end{aligned}$$

and hence

$$p\frac{\mu}{(\lambda+\mu)^2} > (1-\sigma)\alpha.$$

Writing $r = \frac{p}{(1-\sigma)\alpha}$, this implies an upper bound on λ via the inequality $\lambda < \sqrt{r\mu} - \mu$. Now, over all $\mu \geq 0$, the right-hand side of this inequality is maximized at $\mu = r/4$, where it takes the value $r/4$, and hence we have

$$\lambda < r/4 = \frac{p}{4(1-\sigma)\alpha} < \frac{(1+\varepsilon)p}{4\alpha}.$$

This last inequality contradicts our assumption that $\lambda \geq \frac{(1+\varepsilon)p}{4\alpha}$. $\square$

Thus, as long as there is any positive probability p that a question in the forum will be an open-ended question (which might be discussion questions posed by the instructor herself), the instructor can choose μ to produce a rate λ^* such that the rate of arrival of the first response to any question on the forum scales linearly with the size of the class, obtaining a speedup with class size which is of the same order as in the case of purely open-ended questions.

6. DISCUSSION

In this paper, we introduced a game-theoretic model for online forums for education, and investigated how an instructor can influence student participation and utility in these forums.

For single-answer questions, we find that the instructor has very little influence over the rate at which students participate in the forum, especially as the class size grows larger and larger— here, the rate of response is driven almost entirely by the students' competition amongst each other to arrive first to supply the answer. However, when the instructor is optimizing for the speed at which questions are answered, taking into account her own ability to spend effort on producing answers, a more complex effect emerges in which the most efficient behaviors of the forum are clustered into two modes — one in which the instructor steps back and lets the students compete to answer, and one in which she effectively "goes it alone."

The instructor's participation plays a very different role in open-ended questions. First, the instructor must choose a rate $\mu > 0$ if she would like the students to at all contribute the forum— if $\mu = 0$, the students choose a rate $\lambda(\mu) = 0$ as well. Second, the students' response rate λ varies non-monotonically in the instructor's rate; as noted earlier, this provides an interesting strategic basis for the observation that students will increase their effort in the presence of an active instructor up to a point, but then start to give up if the instructor is too active [Gray 2004; Hanover Research Council 2009].

It is also interesting to note the way in which the student's rates range very widely as the instructor's rate ranges over values in the interval $\mu \in [0, \frac{1}{\alpha}]$, where α is the students' cost parameter: each student's rate increases from 0 at $\mu = 0$ until the optimal value $\mu = \frac{1}{4\alpha}$ that elicits the fastest responses from the students, and then falls back to 0 at $\mu = \frac{1}{\alpha}$. Thus, the overall response rate of the class, $n\lambda$, ranges all the way from 0 to $\frac{n}{4\alpha}$ and back to 0 over this interval $\mu \in [0, \frac{1}{\alpha}]$. When α is small, so that students consider effort to be inexpensive, this can be a reasonably gradual effect in μ. But for large α, the effect happens over a narrow interval; this means that for a class where students have a high cost to frequently checking the forum, the behavior of the forum is very delicate and changes dramatically over a very small interval of the instructor's rate in $[0, \frac{1}{\alpha}]$. Such sensitivity to the underlying parameters connects in intriguing ways to types of delicate behavior observed in real online forums, where small shifts in activity level can lead to wide changes in participation rate.

Among the natural questions for further work, a fundamental one that connects to this point about sensitivity is to ask how an instructor might learn the optimal rate to use for her forum. While we do not attempt to formally address this question here, note that the response rate for *single-answer* questions is essentially insensitive to the instructor's choice of rate, but does depend on the students' cost parameter α in our stylized model. This suggests that an instructor could initially use single-answer questions to learn the students' cost parameters from their response rates to these questions, and use this to infer the optimal rate μ^* that she should use for open-ended questions.

More broadly, developing a model for learning parameter values to optimize student participation, as well as modeling learned behavior and evolution of response rates in such online education forums is an interesting open direction for further work.

REFERENCES

ADAMIC, L., ZHANG, J., BAKSHY, E., AND ACKERMAN, M. 2008. *Knowledge sharing and Yahoo! Answers: Everyone knows something.* Proc. 17th International World Wide Web Conference.

ANDRESEN, M. A. 2009. *Asynchronous discussion forums: success factors, outcomes, assessments, and limitations.* Educational Technology & Society, 12 (1).

BURGE, E., LAROQUE, D., AND BOAK, C. 2000. *Baring professional souls: Reflections on Web life.* Journal of Distance Education, 15(1).

CHEN, Y., JAIN, S., AND PARKES, D. 2009. *Designing Incentives for Online Question and Answer Forums*, Proc. ACM EC.

GHOSH, A. AND HUMMEL, P. 2012. *Implementing Optimal Outcomes in Social Computing: A Game-theoretic Approach*, Proc. 21st International World Wide Web Conference.

GRAY, B. 2004. *Informal Learning in an Online Community of Practice*. J. Distance Educ.. 19(1).

HANOVER RESEARCH COUNCIL. 2009. *Best Practices in Online Teaching Strategies*.

KUMAR, R., LIFSHITS, Y., AND TOMKINS, A. 2010. *Evolution of two-sided markets*, Proc. 3rd ACM International conference on Web Search and Data Mining (WSDM '10).

MAZZOLINI, M. AND MADDISON, S. 2003. *Sage, guide or ghost? The effect of instructor intervention on student participation in online discussion forums.* Computers & Education, 40(3).

TOPPING, K. AND EHLY, S. 1998. *Peer-Assisted Learning*, Lawrence Erlbaum.