A Game-Theoretic Analysis of Rank-Order Mechanisms for User-Generated Content

Arpita Ghosh
Yahoo! Research
Santa Clara, CA, USA
arpita@yahoo-inc.com

Patrick Hummel
Yahoo! Research
Berkeley, CA, USA
phummel@yahoo-inc.com

ABSTRACT

Many websites rank user-generated content (UGC) using viewer votes, displaying higher quality contributions more prominently and suppressing lower quality ones. Such an allocation of attention constitutes a *mechanism*, which can influence the quality of content elicited from attention-motivated contributors. In this paper, we analyze equilibrium behavior in the widely used *rank-order* mechanism, where contributions are allocated positions on the page in decreasing order of their ratings, and the *proportional* mechanism which distributes attention in proportion to the number of positive ratings, in a game-theoretic model where agents are motivated by attention and the cost of making a contribution is increasing in its quality.

The rank-order mechanism always possesses a symmetric mixed strategy equilibrium in which all agents decide whether to contribute with the same probability, and randomly draw a quality from a common distribution conditional on participating. We investigate equilibrium behavior in the limit of diverging attention, and show that the *lowest* quality that can arise in a mixed strategy equilibrium of the rank-order mechanism becomes optimal as the amount of available attention diverges. We then compare equilibrium qualities in the proportional and the rank-order mechanism and show that the probability an agent chooses a higher quality in the rank-order mechanism than in the proportional mechanism goes to one as the amount of available attention diverges. Thus the rank-order mechanism almost always incentivizes higher quality contributions in equilibrium than the proportional mechanism.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social & Behavioral Sciences

General Terms

Economics

EC'11, June 5–9, 2011, San Jose, California, USA.

Keywords

User generated content (UGC), Attention economics, Quality of online content, Game theory

1. INTRODUCTION

There is a proliferation of user contributed content on the Web, and a multitude of instances where user contributed content adds significant value to websites. The product reviews written by users on Amazon, for instance, are a very valuable component of the service that Amazon provides, while question-and-answer sites such as Yahoo! Answers and StackOverflow or sites aggregating service reviews such as Yelp owe almost all their utility to contributions from users. But while there is a large amount of user contributed content online, not all of it is of the same *quality*— some content is excellent, while some is mediocre and some is outright bad.

Many websites attempt to *rank* content according to its quality, using thumbs-up/thumbs-down style ratings by viewers— this is the case, for example, with comments on Yahoo! News, reviews on Amazon, and posts on Reddit. These websites display higher quality contributions more prominently by placing them near the top of the page and pushing lower quality ones to the bottom. Since content displayed near the top of the page is more likely to be viewed by a user, ranking good content higher leads to a better user experience. But there is also another aspect to displaying better content more prominently: it potentially provides an *incentive* to produce high quality content that might appeal to a contributor's desire for attention. In other words, how contributions are displayed as a function of their estimated quality constitutes a *mechanism* for allocating attention, which might affect the incentives of contributors and influence the quality of their contributions.

What can we understand, using a game-theoretic approach, about how the mechanism used to display content influences the quality of the contributions? In particular, how does the choice of mechanism influence quality when the number of potential viewers, and therefore the potential available attention, grows very large? The diverging attention regime is arguably the most important from the perspective of user-generated content. First, these are the situations where delivering high quality content matters the most from the perspective of viewer welfare. Second, the popular sites are the ones that draw the most attention-motivated contributors, as well as the ones that tend to

attract contributions of varying quality[1]. Indeed, tremendously large amounts of attention are not uncommon for popular content on the web; for instance, the most popular YouTube videos have been viewed over a hundred million times and even days-old 'trending' videos have hundred of thousands of views, numbers that clearly belongs to the diverging attention regime.

In this paper, we analyze two mechanisms that use viewer ratings to allocate attention to content— the widely used rank-order mechanism, where contributions are allocated positions on the page in decreasing order of their ratings, and a *proportional* mechanism [2, 3], which distributes attention in proportion to the number of positive ratings. The rank-order mechanism is ubiquitous on UGC sites all over the Web, while the proportional mechanism is a natural and more 'fair' alternative: if two contributions receive very similar numbers of votes, it only seems fair that they receive similar amounts of attention as well, but this does not necessarily hold in the rank-order mechanism. What happens to equilibrium quality and participation in the rank-order mechanism as the amount of available attention diverges, and how does it compare against the more fair proportional mechanism?

Our Contributions. We analyze equilibrium behavior in the rank-order mechanism in a game-theoretic model where contributors are motivated by attention and have a cost of participation that increases with the quality of their content, as in [3]. Contributors in our model are strategic agents who choose *both* whether to contribute, modeled by their probability of participation, *and* the quality of their contribution conditional on participating, to maximize their payoff. This ensures that participation is *endogenously* determined as a strategic choice by potential contributors, as in [3]: capturing endogenous participation is important in the UGC setting since users have the option of not contributing should they prefer not to do so.

Analyzing the rank-order mechanism is nontrivial because an agent's payoff depends on her choice of quality and other agents' choices of quality in a complicated way— the number of votes m_i an agent i receives is a *random variable* in her quality q_i. The final payoff to i depends on the *rank* of the instantiation of this random variable amongst the m_{-i}, the number of votes received by other agents, which are *also* random variables in the other agents' quality choices.

We first show that symmetric mixed strategy equilibria, where all agents participate with probability β and randomly choose a quality from a common distribution $F(q)$ if contributing, always exist for the rank-order mechanism. We then investigate equilibrium behavior as the amount of attention increases, and show that the *lowest* quality that can arise in a mixed strategy equilibrium tends to the highest possible quality (*i.e.*, the optimal quality) as the available attention diverges. We also show that if there are not too many potential contributors, it is possible to choose the number of votes used by the mechanism to rank contributions in such a way as to guarantee full participation (*i.e.*, $\beta = 1$) in the limit of diverging attention.

[1]For example, Slashdot notes in its comments moderation FAQ that when Slashdot was small, it received a small number of good quality posts— the need for rating comments arose only after it grew larger and more popular, which led to posts of varying quality.

Finally, we compare the rank-order mechanism to the more equitable proportional mechanism. We show that the proportional mechanism always has an equilibrium in which all agents participate with probability β and choose a fixed quality q_p upon participating. However, unlike the rank-order mechanism, whether this equilibrium quality converges asymptotically to the optimal quality *depends* on how the number of potential contributors grows with the number of viewers. We then compare the random equilibrium quality choices in the rank-order mechanism with the deterministic quality choices in the proportional mechanism and show that the probability an agent chooses higher quality in the rank-order mechanism than in the proportional mechanism tends to 1 as the amount of available attention diverges. Thus the rank-order mechanism elicits higher quality contributions in equilibrium than the proportional mechanism.

Related Work. There is a growing body of research on human computing systems and user contributed content, but relatively little of this work addresses the analysis and design of these systems from a game-theoretic perspective [2, 7, 8]. [2] studies the question of designing incentives for online question-and-answer forums, and focuses on incentives for participants to contribute their answers *quickly* instead of delaying their contribution, but does not address the issue of incentivizing high quality contributions. The most relevant paper to our work is [3], which introduces a model to address the quality of user-generated content. This paper shows that a simple mechanism which eliminates contributions that are not uniformly rated highly by all voters achieves optimal quality in the limit as the amount of available attention diverges. Our model has a few technical differences from that in [3], and is also used to instead address the problem of incentives in the widely used rank-order mechanism and compare those to incentives in the proportional mechanism.

There is a large literature in economics on using rank-order tournaments as incentive schemes (e.g. [4], [5], [10], [12]). This literature analyzes the consequences of employee compensation schemes which reward employees based on how their output compares to that of other employees in the organization. While agents in our model are also ranked based on how the number of positive votes they receive compares to that of other agents, there are several important differences between our paper and this work. In our paper, an individual's observable output is the random number of users who vote positively on her contribution, but this type of framework is not captured by the assumptions made in existing work on rank-order tournaments. Also, not all agents need participate in equilibria in our model, but all employees must work in the economics literature. In addition, most existing work on rank-order tournaments focuses on equilibria in which all agents exert a deterministic level of effort in equilibrium,[2] whereas we extensively analyze mixed strategy equilibria. The focus on the limiting case we consider as well as the comparison between the rank-order and proportional mechanisms is also missing in this literature.

There is also a literature in political economy that compares how allocating political power in proportion to the number of votes received rather than in order of which can-

[2]One very rare exception is [12], which briefly mentions that agents may have an incentive to exert random levels of effort, but does not conduct an extensive analysis of such equilibria.

didates receive the most votes affects elections (e.g. [6], [11], [13], [14]). However, these papers do not focus on how these differences would affect incentives for individuals to produce high quality actions.

2. MODEL

Content. Each unit of content, or contribution, has a *quality* q, where $q \in [0, 1]$ is the probability that a viewer will like the contribution, *i.e.*, rate it as 'good' or 'useful'. The quality q of a contribution is not directly observable in our model, but it influences the number of positive votes the content receives.

Each contribution is rated by T viewers, where we later allow T to be a parameter that can be chosen by the mechanism. We assume that viewers are not strategic, and simply provide this binary feedback non-strategically by truthfully answering the question of whether they found the contribution to be good or useful.

Given a contribution i with quality q_i, the number of positive votes it receives is a random variable. We let m_i denote the number of positive votes received this contribution, and note that the distribution of m_i is binomial with parameters (T, q_i). We also let m_{-i} denote the vector of the numbers of positive votes received by other contributions.

Contributors. There is a pool of potential contributors, or agents, of size K. Each agent can choose the *probability* with which she will contribute, as well as the *quality* of her contribution should she decide to participate. We denote the probability that agent i decides to contribute by β_i and the quality she chooses when she contributes by q_i. Since each potential contributor decides whether to participate or not probabilistically, the number of actual contributors is a random variable. We denote the instantiation of this random variable by k, and note that the distribution of k depends on K as well as $\beta_1, \ldots, \beta_K$.

Contributors are *strategic*: they choose both β and q strategically to maximize their expected payoffs given the potential costs and benefits from contributing. We describe these payoffs next.

If an agent chooses not to contribute, then she incurs no cost but also receives no benefit, so her net payoff is 0. Otherwise, the agent must pay a cost reflecting the effort she expended to produce content, and obtains a benefit reflecting the attention she was able to get as a result of producing that content.

The *cost* incurred by a contributor depends on the quality of the content she chooses to produce: the cost of producing content of quality q is $c(q)$, which is an increasing function of q (*i.e.*, producing higher quality content is more costly). We make the following assumptions:

1. The cost function c is continuously differentiable in q for all $q < 1$.
2. $c(0) > 0$ and $\lim_{q \to 1} c(q) = \infty$.
3. c is convex and $c'(0) = 0$.

The condition $c(0) > 0$ indicates that making a contribution, even a very poor one, takes more effort than not participating at all. The assumption that $\lim_{q \to 1} c(q) = \infty$, as in [3], indicates that producing perfect content, which corresponds to making every viewer happy, is nearly impossible. The final assumption is a standard assumption for cost functions in the economics literature and indicates that it is more difficult to improve higher-quality content, while improving awful content takes almost no effort at all.

The *benefit* derived by a contributor, as in [3], depends on the amount of attention she receives, which can depend in general both on the rating of her contribution and the number and ratings of other contributions. Let A denote the total amount of attention that is available to distribute amongst contributions[3]. We assume that an agent's benefit is exactly equal to the attention that she receives. If contributor i is allocated a fraction $\alpha_i(m_i, m_{-i}) \geq 0$ of the attention A, then her benefit from receiving this attention is $V(m_i, m_{-i}) = \alpha_i(m_i, m_{-i})A$. A *mechanism* in this setting determines how to partition the available attention amongst the various contributions; thus the value of $V(m_i, m_{-i})$ is decided by the mechanism.

A contributor i's *payoff* from generating content of quality q_i is the difference between her expected benefit and cost,

$$\pi(q_i, q_{-i}, \beta_{-i}) = E[V(m_i, m_{-i})|(q_i, q_{-i}, \beta_{-i})] - c(q_i),$$

where $q_{-i} = (q_1, \ldots, q_{i-1}, q_{i+1}, \ldots, q_k)$ denotes the quality choices of the other contributors, and β_{-i} denotes the participation probabilities of the remaining contributors. (Here we assume that i has already decided to contribute, so her participation probability does not affect this payoff.) Note that the expectation in this payoff is taken over the random number of contributors k as well as the random variables m_i that are determined in part by the parameters q_i.

We assume that every viewer who visits a page or site brings at least one unit of attention, but possibly more, corresponding to the fact that each viewer will read at least one contribution but some read more. Thus A is at least as large as the number of viewers. We will also assume throughout that the number of potential contributors K is no larger than the number of potential viewers and therefore the total amount of attention, *i.e.*, $K \leq A$.

Solution Concept. Since agents' payoff functions are symmetric in the parameters of the game, we focus throughout on symmetric equilibria. In a symmetric equilibrium, each contributor participates with the same probability and follows the same strategy of quality choices conditional on participating. Formally, a *symmetric mixed strategy equilibrium* is a probability β and a distribution F over qualities q such that when every agent contributes with probability β, and chooses a quality drawn from the CDF $F(q)$ conditional on contributing, no agent can increase her expected payoff by deviating from this strategy, *i.e.*, by changing either the probability of participation or the distribution over qualities with which to contribute.

Asymptotics. We will be particularly interested in the qualities and participation levels in equilibrium in the limiting case as the amount of attention A goes to infinity. We will assume that this diverging amount of attention comes from a diverging number of viewers; therefore, the number of potential contributors, K, as well as the number of viewers available to vote on contributions, T, can increase with A as well.

We will sometimes write $K(A)$ and $T(A)$ to make explicit the possible dependence of K and T on A, and assume that as A diverges, $K(A)$ diverges and $T(A)$ can be chosen to di-

[3]If the agents did not know the exact value of A, but instead only knew its distribution, then the equilibrium analysis would stay the same except that A would be replaced with $E[A]$.

verge as well. This corresponds to the observation in practice that as sites grow more popular in terms of viewership (A increases), they attract more users who are interested in contributing ($K(A)$), and more viewers are available to rate the generated content ($T(A)$). For instance, the growing popularity of YouTube has led to an increasing number of videos posted to YouTube, as well as an increasing number of ratings on these videos. We emphasize, though, that $T = T(A)$ is a parameter chosen by the mechanism, and specifically can be chosen to be small compared to A, since the mechanism can choose to use only a subset of the available ratings (for instance, the first T votes) to rank contributions. Thus, the actual number of ratings is only an upper bound on T, and not necessarily the value of T chosen by the mechanism.

3. PRELIMINARIES

First note that if $A \leq c(0)$, *i.e.*, the total available attention is smaller than the cost of producing even zero quality content, then no agents would want to participate since an agent's payoff from participating is nonpositive even if there is no other agent with whom attention must be shared. The case where no agents participate is not interesting, so we will assume that

$$A > c(0)$$

for the remainder of this paper. This assumption guarantees that $\beta > 0$ in any symmetric equilibrium[4].

A pair $(\beta, F(q))$ with $\beta \in (0, 1]$ constitutes a symmetric mixed strategy equilibrium if the following two conditions are satisfied:

1. No agent wants to choose a different β: For this to hold when $\beta = 1$, the expected payoff to an agent from participating must be nonnegative when she chooses a quality from $F(q)$ and the remaining agents participate with probability $\beta = 1$ and choose quality from $F(q)$. For this to hold for $\beta \in (0, 1)$, an agent's expected payoff from participating must be exactly zero since otherwise she would want to either always participate or never participate.

2. For any q in the support of F, no agent can obtain a strictly larger payoff by choosing some other quality $q' \in [0, 1)$ instead of q upon participating, given that remaining agents use the strategy $(\beta, F(q))$.

We next elaborate on the conditions that must hold for $\beta \in (0, 1]$ to be an equilibrium participation probability in a symmetric equilibrium.

If all agents participate with probability $\beta = 1$ and draw their qualities from the same distribution $F(q)$ upon participating, then by symmetry, each agent receives the same amount of attention as every other agent in expectation. Therefore, each agent's expected benefit is $\frac{A}{K}$ ($\beta = 1$ so $k = K$, *i.e.*, all agents contribute) and each agent's expected cost is $E[c(q)|q \sim F(q)]$. In order for an agent to not be able to profitably deviate by changing her participation probability, we must thus have $\frac{A}{K} \geq E[c(q)|q \sim F(q)]$.

Now suppose all other agents participate with probability $\beta \in (0, 1)$ and agent i decides that she will participate and use the same quality distribution $F(q)$ as the other agents. Then agent i obtains an expected benefit of $E[\frac{A}{k} |\ i$ enters; other agents enter with probability $\beta]$, where the expectation is over the random decisions of other agents to contribute. This holds because each agent draws qualities from the same distribution, so by symmetry each agent's expected benefit from participating is $\frac{A}{k}$ when $k - 1$ of the $K - 1$ remaining contributors participate. The expected cost of using strategy $F(q)$, conditional on participating, is $E[c(q)|q \sim F(q)]$. Thus in order for agent i to be indifferent between participating and not participating, we must have that $E[\frac{A}{k}|\ i$ enters; other agents enter with probability $\beta] = E[c(q)|q \sim F(q)]$.

To summarize, the necessary conditions to have an equilibrium in which each agent participates with probability $\beta \in (0, 1]$ and chooses quality drawn from the distribution $F(q)$ upon participating are the following:

1. If $\beta \in (0, 1)$, the expected benefit from participating must equal the expected cost:

$$E[\frac{A}{k}\ |i \text{ enters; other agents enter with probability } \beta] = E[c(q)|q \sim F(q)].$$

2. If $\beta = 1$, the expected benefit from participating is not smaller than the expected cost:

$$\frac{A}{K} \geq E[c(q)|q \sim F(q)].$$

3. For any q in the support of F, no agent can obtain a strictly larger payoff by choosing some other quality $q' \in [0, 1)$ instead of q upon participating, given that remaining agents use the strategy $(\beta, F(q))$.

4. RANKING MECHANISM

We first consider *ranking* mechanisms. A ranking mechanism arranges contributions in decreasing order of the number of positive votes they receive and allocates more attention to contributions which are ranked higher. The common practice of displaying comments or answers on a webpage in decreasing order of the number of positive votes received, used by most sites hosting UGC, corresponds to such a ranking mechanism since more users view content near the top of the page than the bottom. In this section, we analyze the mixed strategy equilibria of ranking mechanisms, and then prove results about how equilibrium quality and participation probability behave as the amount of available attention becomes large. We first formally define ranking mechanisms.

Definition 4.1 (Ranking Mechanism $M_r(T, \alpha)$). *Let $\alpha_j(k) \geq 0$ be a sequence of numbers that is nonincreasing in j for all k and satisfies $\sum_{j=1}^{k} \alpha_j(k) = 1$, where the values of α_j* do not depend *on the qualities q for any j, although they can depend on the* number *of contributors k.*

Suppose there are k contributors and each contribution is voted on by T viewers. The ranking mechanism ranks contributions in decreasing order of the number of positive votes received (with ties broken randomly) and awards the jth ranked contribution attention $\alpha_j A$.

Note that both α, which specifies the distribution of attention amongst the ranks, as well as T, the number of votes used to determine the rankings, are *parameters* of the mechanism that can be chosen to achieve desirable properties.

We first state the following simple proposition.

[4] $\beta = 0$ cannot be an equilibrium participation probability since then an agent can improve her current payoff of 0 by entering and producing zero quality content.

PROPOSITION 4.1. *Suppose $q_i > q_j$, and let the number of votes received by i and j be m_i and m_j respectively. Then* $\lim_{T\to\infty} Pr(m_i > m_j) = 1$.

Thus a higher quality contribution receives a larger number of positive votes, and is therefore ranked higher, than a lower quality contribution in the limit as the number of votes T goes to infinity.

We first show that the ranking mechanism always has a symmetric mixed strategy equilibrium in which all agents choose the same probability of entry, and choose a quality from the same distribution if they decide to contribute.

THEOREM 4.1. *For any values of A, K, T, and α, there exists a symmetric mixed strategy equilibrium in which all contributors enter with probability β and choose a quality drawn from the same cumulative distribution function $F(q)$ conditional on contributing.*

PROOF. First note that no player in this game would ever choose a quality $q > c^{-1}(A)$, as a player could always obtain a strictly greater expected payoff by not participating than by participating and choosing a quality $q > c^{-1}(A)$. Thus any mixed strategy equilibrium to the game in which players are restricted to choosing $q \in [0, c^{-1}(A)]$ is also a mixed strategy equilibrium of the original game.

Now note that this modified game in which players are restricted to choosing $q \in [0, c^{-1}(A)]$ is a symmetric game in which each player has a pure strategy space that is compact and Hausdorff. Also note that each player's expected payoff in this modified game is continuous in the actions of the players. It thus follows from Theorem 1 of [1] that there exists a symmetric mixed strategy equilibrium of this modified game. This in turn implies that there is a symmetric mixed strategy equilibrium of the original game. □

Thus it always makes sense to analyze symmetric mixed strategy equilibria of the form discussed in §2. While mixed strategies are more difficult to analyze than pure strategies, the use of mixed strategies as a solution concept is particularly reasonable in this setting as it allows for the possibility that different contributors produce different quality content. This correlates well with observation in practice, where not all contributions, for example answers to a question on StackOverflow or comments on a news article, are of the same quality.

In general, a mechanism can have multiple mixed strategy equilibria. A natural question to ask is whether any of these equilibria is '*bad*', in the sense of inducing low-quality equilibrium contributions. Our first result illustrates that as long as there is at least a small difference between the amount of attention allocated to contributions that are ranked differently, such bad Nash equilibria cannot exist. Specifically, we show that if T, the number of votes used to rank the contributions, becomes large as the amount of attention grows, contributors will choose quality arbitrarily close to 1 in the limit as the amount of attention diverges.

THEOREM 4.2. *Suppose that $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ for all $j \leq k-1$, $\alpha_k(k) = 0$ for sufficiently large k, and $\lim_{A\to\infty} T(A) = \infty$. Then, for any $q^* < 1$, the probability an agent who contributes chooses quality $q > q^*$ goes to 1 as A goes to infinity.*

In a mixed strategy equilibrium, a contributor can draw any quality in the support of the equilibrium distribution: the theorem says that the *lowest* quality in this equilibrium distribution tends to the optimal quality as the amount of attention diverges.

Since the precise equilibrium strategies that agents use may vary with A, we will henceforth use $\beta(A)$ to denote the dependence of the equilibrium participation probability, and $F_A(q)$ to denote the dependence of the equilibrium distribution from which contributors choose their quality, on the available attention A. We now prove the result.

PROOF OF THEOREM 4.2. Suppose by means of contradiction that there exists some $q^* < 1$ and some $\gamma > 0$ such that the probability a contributor chooses quality $q \leq q^*$ is at least γ for an infinite number of A. If q_A denotes the minimum value of the set of all q in the support of $F_A(q)$, then $q_A \leq q^*$ holds for all such A. For small $\epsilon > 0$, let $p_A(\epsilon) = F_A(q_A + \epsilon)$ denote the probability that a contributor chooses some quality $q \leq q_A + \epsilon$ for a given A.

Our proof breaks down into three steps. We first show that if we restrict attention to a subsequence of A for which the probability a contributor chooses quality $q \leq q^*$ is at least γ, it must be the case that $\lim_{A\to\infty} p_A(\epsilon) \leq C\epsilon$ for some constant C and small $\epsilon > 0$, where the limit is taken along this subsequence. We then show that if $\lim_{A\to\infty} p_A(\epsilon) \leq C\epsilon$, then it must be the case that $\lim_{A\to\infty} \frac{\beta(A)K(A)}{A} = 0$ along this subsequence. Finally, we show that if $\lim_{A\to\infty} p_A(\epsilon) \leq C\epsilon$ and $\lim_{A\to\infty} \frac{\beta(A)K(A)}{A} = 0$ along this subsequence, then it cannot be the case that q_A is in the support of $F_A(q)$.

Step 1: We first show that if we restrict attention to a subsequence of A for which the probability a contributor chooses quality $q \leq q^*$ is at least γ, it must be the case that $\lim_{A\to\infty} p_A(\epsilon) \leq C\epsilon$ for some constant C and small $\epsilon > 0$, where the limit is taken along this subsequence.

To see this, suppose by means of contradiction that there is no constant C such that $\lim_{A\to\infty} p_A(\epsilon) \leq C\epsilon$ for small $\epsilon > 0$ along this subsequence. Note that if a contributor chooses quality $q = q_A$, then the probability she receives a higher ranking than a particular other contributor who chooses quality $q \leq q_A + \epsilon$ is no greater than $\frac{1}{2}$. So if $\eta_A(\epsilon)$ denotes the expected fraction of contributors who choose quality $q > q_A + \epsilon$ and receive a lower ranking than this contributor for a given A, then the expected total fraction of contributors who receive a lower ranking than her is no greater than $\frac{p_A(\epsilon)}{2} + \eta_A(\epsilon)$.

Now note that if this contributor instead uses quality $q = q_A + 2\epsilon$ for some $\epsilon > 0$, then the probability she receives a higher ranking than a particular other contributor who chooses quality $q \leq q_A + \epsilon$ goes to 1 in the limit as T goes to infinity. The expected fraction of contributors who choose quality $q > q_A + \epsilon$ and receive a lower ranking than a contributor who uses quality $q = q_A + 2\epsilon$ is at least as large as the expected fraction of contributors who choose quality $q > q_A + \epsilon$ and receive a lower ranking than a contributor who uses quality $q = q_A$. Thus if she instead uses quality $q = q_A + 2\epsilon$ for some $\epsilon > 0$, then the expected fraction of contributors who choose quality $q > q_A + \epsilon$ and receive a lower ranking than her is at least $\eta_A(\epsilon)$. From this it follows that the expected total fraction of contributors who receive a lower ranking than this contributor when she chooses quality $q = q_A + 2\epsilon$ is at least $p_A(\epsilon) + \eta_A(\epsilon)$ in the limit as T goes to infinity.

Thus, choosing quality $q = q_A + 2\epsilon$ instead of $q = q_A$ re-

sults in this contributor receiving a higher ranking than an expected fraction of at least $\frac{p_A(\epsilon)}{2}$ more contributors. Therefore, if there are k participating contributors, the expected number of contributors that she beats increases by at least $\frac{p_A(\epsilon)k}{2}$ as a result of this change. Moving up in the rankings by one spot increases a contributor's payoff by $\Theta(\frac{A}{k^2})$ since $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ for all $j \leq k-1$. Thus choosing quality $q = q_A + 2\epsilon$ instead of $q = q_A$ increases her expected benefits by $\Theta(\frac{p_A(\epsilon)k}{2}(\frac{A}{k^2})) = \Theta(\frac{p_A(\epsilon)A}{k})$. So if $b_A(\epsilon)$ denotes the change in expected benefits from choosing quality $q = q_A + 2\epsilon$ instead of choosing quality $q = q_A$, there is no constant C such that $\lim_{A\to\infty} b_A(\epsilon) \leq C\epsilon$ for small $\epsilon > 0$.

Now let $c_A(\epsilon)$ denote the added cost a contributor incurs by choosing quality $q = q_A + 2\epsilon$ instead of choosing quality $q = q_A$. Note that there exists some constant C such that $c_A(\epsilon) \leq C\epsilon$ for small $\epsilon > 0$ since $q_A \leq q^* < 1$ implies $c'(q_A) \leq c'(q^*)$, which is finite. Combining this with the result in the previous paragraph shows that there exists some large A and some small $\epsilon > 0$ such that a contributor obtains a strictly larger expected payoff from choosing the quality $q = q_A + 2\epsilon$ instead of choosing quality $q = q_A$. This contradicts the fact that q_A is in the support of $F_A(q)$ and proves that there exists some constant C such that $\lim_{A\to\infty} p_A(\epsilon) \leq C\epsilon$ for small $\epsilon > 0$ along this subsequence.

Step 2: Now we show that if there exists some constant $C > 0$ such that $\lim_{A\to\infty} p_A(\epsilon) \leq C\epsilon$ for small $\epsilon > 0$ along this subsequence, then it must be the case that $\lim_{A\to\infty} \frac{\beta(A)K(A)}{A} = 0$ along this subsequence.

To see this, note that if there exists some constant $C > 0$ such that $\lim_{A\to\infty} p_A(\epsilon) \leq C\epsilon$ for small $\epsilon > 0$ along this subsequence, then as T becomes large, the expected fraction of contributors who receive a lower ranking than a contributor who participates with quality $q = q_A$ goes to zero, for the following reason. Let $\epsilon(T)$ denote the largest value of $\epsilon > 0$ such that the probability of an agent with quality q_A receiving a higher ranking than one with quality $q_A + \epsilon(T)$ is at least $\frac{1}{T}$; then, $\lim_{T\to\infty} \epsilon(T) = 0$. Thus the expected fraction of contributors who receive a lower ranking than a contributor who participates with quality q_A is no greater than $p_A(\epsilon(T)) + \frac{1}{T}$ for any T. This tends to zero as T goes to infinity.

Thus if $g(A)$ denotes the expected fraction of contributors who receive a lower ranking than a contributor who participates with quality q_A, then $\lim_{A\to\infty} g(A) = 0$. Now if there are k contributors, then the expected number of contributors who receive a lower ranking than one who participates with quality q_A is $g(A)k$. Combining this with the facts that $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ for all $j \leq k-1$ and $\alpha_k(k) = 0$ for sufficiently large k shows that a contributor who participates with quality q_A obtains an expected amount of attention $\Theta(g(A)k\frac{A}{k^2}) = \Theta(g(A)\frac{A}{k})$.

Now $\frac{k}{\beta(A)K(A)}$ converges in probability to 1 as $A \to \infty$, so $\Theta(g(A)\frac{A}{k}) = \Theta(g(A)\frac{A}{\beta(A)K(A)})$. Thus if $\lim_{A\to\infty} \frac{\beta(A)K(A)}{A} = 0$ does not hold along this subsequence, then $\lim_{A\to\infty} g(A) = 0$ implies it must be the case that a contributor who contributes with quality q_A receives an expected amount of attention that approaches zero as A goes to infinity. But this means that an agent could obtain a strictly higher payoff by not contributing than by contributing with quality q_A for some large A, contradicting the fact that q_A is in the support of $F_A(q)$. This contradiction shows that $\lim_{A\to\infty} \frac{\beta(A)K(A)}{A} = 0$ must hold along this subsequence.

Step 3: Finally we show that if $\lim_{A\to\infty} \frac{\beta(A)K(A)}{A} = 0$ holds along this subsequence, then it cannot be the case that q_A is in the support of $F_A(q)$.

To see this, let $\gamma_A = F_A(q^*)$ denote the probability that a contributor chooses a quality $q \leq q^*$ for a given A, and note that $\gamma_A \geq \gamma$ for all A in the subsequence. Also note that if a contributor uses quality $q = q^* + \epsilon$ instead of using quality $q = q_A$, then this costs her no more than $c(q^* + \epsilon) - c(q_A) \leq c(q^* + \epsilon)$.

Now if this contributor uses quality $q = q_A$, then the probability she receives a higher ranking than a particular other contributor who chooses quality $q \leq q^*$ is no greater than $\frac{1}{2}$. Thus if δ_A denotes the expected fraction of contributors who choose quality $q > q^*$ and receive a lower ranking than this contributor using quality q_A, then the expected total fraction of contributors that receive a lower ranking than her is no greater than $\frac{\gamma_A}{2} + \delta_A$.

Now, if she instead uses quality $q = q^* + \epsilon$ for some $\epsilon > 0$, then the probability she receives a higher ranking than a particular other contributor who chooses quality $q \leq q^*$ goes to 1 in the limit as T goes to infinity. The expected fraction of contributors who choose quality $q > q^*$ and receive a lower ranking than a contributor who uses quality $q = q^* + \epsilon$ is at least as large as the expected fraction of contributors who choose quality $q > q^*$ and receive a lower ranking than a contributor who uses quality $q = q_A$. Thus if the contributor instead uses quality $q = q^* + \epsilon$ for some $\epsilon > 0$, then the expected fraction of contributors who choose quality $q > q^*$ and receive a lower ranking than the contributor is at least δ_A. From this it follows that the expected fraction of contributors who receive a lower ranking than this contributor when she chooses quality $q = q^* + \epsilon$ is at least $\gamma_A + \delta_A$ in the limit as T goes to infinity.

Thus choosing quality $q = q^* + \epsilon$ instead of $q = q_A$ results in receiving a higher ranking than an expected fraction of at least $\frac{\gamma_A}{2}$ additional contributors. Thus if there are k participating contributors, this contributor increases the expected number of agents she beats by at least $\frac{\gamma_A k}{2}$ as a result of this change. As before, moving up in the rankings by one spot increases the payoff by $\Theta(\frac{A}{k^2})$ since $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ for all $j \leq k-1$. Thus if a contributor chooses quality $q = q^* + \epsilon$ instead of choosing quality $q = q_A$, then she increases her expected benefits by $\Theta(\frac{\gamma_A k}{2}(\frac{A}{k^2})) = \Theta(\frac{A}{k}) = \Theta(\frac{A}{\beta(A)K(A)})$.

But for sufficiently large A in the subsequence, it follows that this increase in expected benefits is greater than $c(q^* + \epsilon)$ since $\lim_{A\to\infty} \frac{A}{\beta(A)K(A)} = \infty$ in the subsequence but $c(q^* + \epsilon)$ is a constant independent of A. Thus from this it follows that a contributor obtains a strictly greater expected payoff from choosing quality $q = q^* + \epsilon$ instead of choosing quality $q = q_A$ for sufficiently large A in the subsequence. This contradicts the existence of some such q^* and proves the desired result. □

This theorem uses the condition that $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$. It is worth noting that the result that contributors produce arbitrarily high quality content does not depend crucially on this assumption. If we instead replace the assumption that $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ for all $j \leq k-1$ in Theorem 4.2 with the weaker assumption that $\alpha_j(k) - \alpha_{j+1}(k) = \Omega(k^{-2})$ for all $j \leq k-1$, then the result

would go through with a very similar proof. That is, there is no problem inducing high quality content as long there is at least some minimum difference between the attention allocated to content ranked differently.

To understand the intuition behind this result, first note that in any equilibrium it must be the case that a large fraction of participating contributors produce content with quality close to q_A, where q_A is the minimum quality in the support of $F_A(q)$: if only a small fraction of agents produce content with quality close to q_A, then a contributor who produces quality q_A will achieve a lower ranking than almost all other contributors, and will obtain almost no attention. Thus such an agent could achieve a higher expected payoff by not contributing, and it would not be possible for q_A to be in the support of $F_A(q)$.

But if a large fraction of participating contributors are producing content with quality close to q_A, then for large T, an agent can ensure that she will achieve a higher ranking than a significant number of additional contributors by producing content with quality $q = q_A + \epsilon$ for some small $\epsilon > 0$ (in contrast with choosing quality $q = q_A$). Thus if q_A is bounded away from 1, an agent could profitably deviate by producing content with quality $q = q_A + \epsilon$ instead of content with quality $q = q_A$ for some small $\epsilon > 0$. From this it follows that q_A must be close to 1 for large A.

Next we show that one can choose the number of individuals who vote on the content, T, in such a way to induce all contributors to participate in the limit, as long as the pool of potential contributors does not grow too quickly with the number of viewers.

THEOREM 4.3. *Suppose that* $\lim_{A\to\infty} \frac{K(A)}{A} = 0$ *and* $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ *for all* $j \leq k-1$. *Then there exists a sequence* $\{T(A)\}_{A=1}^{\infty}$ *such that* $\lim_{A\to\infty} T(A) = \infty$ *and* $\beta = 1$ *in equilibrium for sufficiently large* A.

We prove this theorem in the appendix. This result indicates that if $T(A)$ does not grow too quickly with A, then all contributors will participate for sufficiently large A (note that $T(A)$ still diverges so that quality still becomes optimal, albeit at a slower rate). While this result makes use of a technical assumption that $\lim_{A\to\infty} \frac{K(A)}{A} = 0$, we believe this is the most critical case to ensure large participation, as this is the case where there is only a relatively small number of agents who may contribute.

To understand the intuition behind this result, note that if $T(A)$ does not increase too quickly with A, then contributors will have a relatively lower incentive to choose high quality because producing higher quality content does not result in an especially large increase in one's probability of being ranked highly. But when contributors are not producing exceedingly high quality content, the cost to participating with a quality that will be competitive in the mechanism is relatively small, and agents find it worthwhile to participate. Thus if $T(A)$ does not increase too quickly with A, then all agents will participate for sufficiently large A.

Implementing Ranking Mechanisms. Theorem 4.2 requires a mechanism designer to choose the values of α to satisfy $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ and $\alpha_k(k) = 0$ to induce high quality. The needed difference in attention between two contributions can be achieved as follows: To increase the difference in attention between two contributions, one can show one contribution prominently more often than the other, and to decrease the difference one can show the two contributions in prominent positions roughly equally often. One can also easily ensure that $\alpha_k(k) = 0$ by simply never showing a contribution that is ranked last by the initial voting.

The analysis in Theorem 4.3 suggests that a mechanism designer will want to choose T so that $T(A)$ does not grow too quickly with A in order to induce full participation. This can also be achieved in practice by simply ignoring votes beyond the first T votes. Thus the type of restrictions on T and α used in these results should be easily implementable in practice.

5. PROPORTIONAL MECHANISM

Since all agents produce similar qualities in equilibrium in the rank-order mechanism, it might seem unfair to give significantly greater amounts of attention to agents who receive higher numbers of positive votes, even though the difference between the numbers of positive votes the contributors receive is likely to be very small. A natural and more equitable alternative is to reward contributors in proportion to the number of positive votes they receive— in such a proportional mechanism, two agents who both receive very similar numbers of positive votes also receive similar amounts of attention[5]. However, while the proportional system might allocate attention in a more equitable manner, the mechanism also significantly changes incentives for agents to produce high quality content in equilibrium. In this section, we analyze the equilibria of the proportional mechanism.

First we formally define the proportional mechanism.

DEFINITION 5.1 (PROPORTIONAL MECHANISM $M_p(T)$). *Suppose k agents contribute, each contribution is voted on by T viewers, and each participating contributor i receives m_i positive votes. Then the proportional mechanism gives the ith contributor a share $\frac{m_i}{\sum_{j=1}^{k} m_j}$ of the available attention if $m_i > 0$ for some i. If $m_i = 0$ for all i, every contributor receives a share $\frac{1}{k}$ of the available attention.*

Note that in the proportional mechanism, only the number of votes used by the mechanism, T, is available as a parameter of the mechanism that can be varied to achieve desirable incentives. We first prove a result on the nature of the equilibrium in the proportional mechanism.

THEOREM 5.1. *For any values of A, K, and T, there exists a symmetric equilibrium to the proportional mechanism in which all contributors participate with probability β and choose the same quality q conditional on contributing.*

PROOF. Define $\beta(q)$ to be the unique value of $\beta \in [0,1]$ that satisfies either

$$c(q) = E[\frac{A}{k} | \ i \text{ enters; other agents enter with probability } \beta]$$

if this holds for some $\beta \in (0,1)$, $\beta(q) = 0$ if $c(q) \geq A$, and $\beta(q) = 1$ if $c(q) \leq \frac{A}{K}$. Note that $\beta(q)$ is unique and well-defined for all $q \in [0,1)$ and that $\beta(q)$ is continuous in q.

Define $b(q_i, q_{-i}, \beta)$ to be the expected amount of attention that a contributor i obtains if contributor i contributes with

[5]Lazear [9] notes that fairness can lead to desirable equilibrium outcomes in certain settings where malicious sabotage is a concern.

quality q_i, all other agents who participate contribute with quality q_{-i}, and all other agents participate with probability β. Thus $b(q_i, q_{-i}, \beta) = E[\frac{m_i}{m_i+\sum_{j\neq i} m_j}|\beta, q_i, q_{-i}]$. Note that $b(q^*, q^*, \beta(q^*))$ is continuously differentiable in q^* for all $q^* \in [0, 1)$ since the probability that a particular realization of positive votes (m_i, m_{-i}) takes place is a continuously differentiable function of (q_i, q_{-i}, β).

Now note that $c'(0) < \frac{\partial b}{\partial q_i}(0, 0, \beta(0))$ since $c'(0) = 0$ and $\frac{\partial b}{\partial q_i}(0, 0, \beta(0)) > 0$ since the proportional mechanism assigns equal attention to all participants if all participants obtain zero positive votes. Also note that $\lim_{q^*\to 1} c'(q^*) > \lim_{q^*\to 1} \frac{\partial b}{\partial q_i}(q^*, q^*, \beta(q^*))$ because $\lim_{q^*\to 1} c'(q^*) = \infty$ but $\frac{\partial b}{\partial q_i}(q^*, q^*, \beta(q^*))$ is finite for all $q^* \in [0, 1]$. Since $c'(q^*)$ and $\frac{\partial b}{\partial q_i}(q^*, q^*, \beta(q^*))$ are continuous in q^* for all $q^* \in [0, 1)$, it follows from the intermediate value theorem that there exists some $q^* \in (0, 1)$ such that $c'(q^*) = \frac{\partial b}{\partial q_i}(q^*, q^*, \beta(q^*))$. This shows that there is a quality q^* at which the *local* conditions for not deviating from the quality choice are satisfied, if all participants enter with probability $\beta(q^*)$. To show that this pair of values indeed constitutes an equilibrium, we need to show that (i) no agent wants to choose a different participation probability β', and (ii) there are no *non-local* profitable deviations of quality.

First note that it must be the case that $\beta(q^*) > 0$ for this q^*. If $\beta = 0$, then we know that $\frac{\partial b}{\partial q_i}(q^*, q^*, \beta) = 0$, so $c'(q^*) \neq \frac{\partial b}{\partial q_i}(q^*, q^*, \beta(q^*))$ since the fact that $c(q)$ is increasing and convex in q implies $c'(q^*) > 0$ for $q^* > 0$. Thus $\beta(q^*) = 0$ cannot hold for this q^* and we have $\beta(q^*) > 0$.

Now if $\beta(q^*) \in (0, 1)$, then by definition of $\beta(q)$, it must be the case that

$$c(q^*) = E[\frac{A}{k}|\, i \text{ enters; other agents enter with probability } \beta]$$

and agents are indifferent between participating with quality q^* and not participating. And if $\beta(q^*) = 1$, then $c(q^*) \leq \frac{A}{K}$ and all agents prefer to participate with quality q^* than to not participate. Thus if all agents contribute with probability $\beta(q^*)$ and produce quality q^*, no agent can profitably deviate by participating with a different probability from $\beta(q^*)$.

Now we show that no agent can profitably deviate by contributing with a quality $q_i \neq q^*$ conditional on participating. Note that if an agent participates with quality q_i and all other agents participate with quality q^* and contribute with probability $\beta(q^*)$, then the marginal change in an agent's payoff from contributing with a slightly larger quality is $\frac{\partial b}{\partial q_i}(q_i, q^*, \beta(q^*)) - c'(q_i)$. Since $c(q_i)$ is convex, it follows that $-c'(q_i)$ is decreasing in q_i for all $q_i \in [0, 1)$.

Now we show that $\frac{\partial b}{\partial q_i}(q_i, q^*, \beta(q^*))$ is decreasing in q_i. Note that the marginal benefit from receiving an additional positive vote is smaller when m_i is larger for any fixed values of m_{-i} since the difference between $\frac{m_i+1}{m_i+1+\sum_{j\neq i} m_j}$ and $\frac{m_i}{m_i+\sum_{j\neq i} m_j}$ is decreasing in m_i. Now if $G(m_i|q_i)$ denotes the distribution of the values of m_i given that agent i is producing quality q_i, then $q_i' > q_i$ implies that $G(m_i|q_i')$ first order stochastically dominates $G(m_i|q_i)$. Thus if $q_i' > q_i$, then the expected benefit from receiving an additional positive vote is smaller when an agent is producing with quality q_i' than it is when an agent is producing with quality q_i. From this it follows that $\frac{\partial b}{\partial q_i}(q_i, q^*, \beta(q^*))$ is decreasing in q_i.

Putting this all together shows that $\frac{\partial b}{\partial q_i}(q_i, q^*, \beta(q^*)) - c'(q_i)$ is decreasing in q_i. Thus since $\frac{\partial b}{\partial q_i}(q^*, q^*, \beta(q^*)) - c'(q^*) = 0$, it follows that $\frac{\partial b}{\partial q_i}(q_i, q^*, \beta(q^*)) - c'(q_i) > 0$ when $q_i < q^*$ and $\frac{\partial b}{\partial q_i}(q_i, q^*, \beta(q^*)) - c'(q_i) < 0$ when $q_i > q^*$. But this means that an agent prefers to contribute with quality q^* than with any quality q_i satisfying either $q_i < q^*$ or $q_i > q^*$. From this it follows that an agent cannot profitably deviate by choosing some quality different from q^* and there is an equilibrium in which all agents participate with probability $\beta(q^*)$ and contribute with quality q^* conditional on participating. □

Next we investigate asymptotic equilibrium quality choices in the proportional mechanism. In contrast to the ranking mechanism, whether equilibrium quality choices converge to the optimal quality in the proportional mechanism *depends* on how quickly the number of potential contributors grows with the number of viewers. If this number grows slowly enough that $\frac{K(A)}{A}$ goes to 0 as A goes to infinity, then equilibrium quality converges 1. But if not, equilibrium quality remains strictly less than one in the limit.

Before proving this, we first state a lemma regarding the derivative of the cost function at the equilibrium quality, which we will use for both our remaining results.

LEMMA 5.1. *Let $\beta(A)$ and $q_p(A)$ denote equilibrium participation probabilities and quality choices in the proportional mechanism for a given A. Then $c'(q_p(A)) = \Theta(\frac{A}{\beta(A)K(A)})$.*

This lemma is proven in the appendix. Now we use this lemma to illustrate how equilibrium quality choices vary in the proportional mechanism as the amount of available attention diverges.

THEOREM 5.2. *If $\lim_{A\to\infty} \frac{K(A)}{A} = 0$, then $\lim_{A\to\infty} q_p(A) = 1$ and $\beta = 1$ for sufficiently large A. If $\liminf_{A\to\infty} \frac{K(A)}{A} > 0$, then $\limsup_{A\to\infty} q_p(A) < 1$.*

PROOF. Note that if $\lim_{A\to\infty} \frac{K(A)}{A} = 0$, then the fact that $c'(q_p(A)) = \Theta(\frac{A}{\beta(A)K(A)}) = \Omega(\frac{A}{K(A)})$ implies $\lim_{A\to\infty} c'(q_p(A)) = \infty$ and $\lim_{A\to\infty} q_p(A) = 1$. Also note that if $\beta(A) < 1$ for some large A, then the equilibrium conditions for indifference between contributing and not contributing indicate that $c(q_p(A)) = \Theta(\frac{A}{\beta(A)K(A)})$. Combining this with the fact that $c'(q_p(A)) = \Theta(\frac{A}{\beta(A)K(A)})$ shows that $\frac{c'(q_p(A))}{c(q_p(A))} = \Theta(1)$. But since $\lim_{q\to 1} c(q) = \infty$, $\lim_{q\to 1} \log c(q) = \infty$ as well, $\lim_{q\to 1} \frac{d}{dq} \log c(q) = \infty$, and $\lim_{q\to 1} \frac{c'(q)}{c(q)} = \infty$. Thus the fact that $\lim_{A\to\infty} q_p(A) = 1$ implies that $\lim_{A\to\infty} \frac{c'(q_p(A))}{c(q_p(A))} = \infty$, meaning $\frac{c'(q_p(A))}{c(q_p(A))} = \Theta(1)$ cannot hold. Thus for sufficiently large A, it is not possible to have $\beta(A) < 1$, and it must be the case that $\beta = 1$ for sufficiently large A.

Now suppose that $\liminf_{A\to\infty} \frac{K(A)}{A} > 0$. In this case, if $\beta = 1$ for sufficiently large A, then the fact that $c'(q_p(A)) = \Theta(\frac{A}{\beta(A)K(A)})$ implies that $c'(q_p(A)) = \Theta(1)$ and $\limsup_{A\to\infty} q_p(A) < 1$. And if $\beta(A) < 1$ for some infinite subsequence of A, then an identical argument to that given in the previous paragraph shows that $\frac{c'(q_p(A))}{c(q_p(A))} = \Theta(1)$ and $\limsup_{A\to\infty} q_p(A) < 1$ along this subsequence.

From this it follows that $\liminf_{A\to\infty} \frac{K(A)}{A} > 0$ implies $\limsup_{A\to\infty} q_p(A) < 1$. $\square$

We note that both regimes in this theorem, $\frac{K(A)}{A} \to 0$ and $\frac{K(A)}{A} \to r > 0$, are of interest in the context of user-generated content. In question-and-answer sites such as Yahoo! Answers or StackOverflow, the number of users $K(A)$ who can answer a question is often significantly smaller than the number of users who consume the answer, possibly via a search engine, and $\frac{K(A)}{A} \to 0$ is likely. On the other hand, in settings like posts on discussion forums or comments on blogs where many consumers are also producers, the number of contributors may not be negligible compared to the number of viewers who consume the content, *i.e.*, $\frac{K(A)}{A}$ is not vanishingly small. The theorem says that the proportional mechanism elicits the optimal quality in the first kind of setting, but not in the second.

6. COMPARING MECHANISMS

In this section, we compare equilibrium qualities in the ranking mechanism and the proportional mechanism. We already know from Theorems 4.2 and 5.2 that when $\lim_{A\to\infty} \frac{K(A)}{A} > 0$, the lowest quality in the support of the equilibrium distribution in the ranking mechanism converges to 1, but not in the proportional mechanism. For this case, therefore, the ranking mechanism leads to higher quality contributions than the proportional mechanism in the limit of diverging attention. We now complete this comparison by investigating the case where $\lim_{A\to\infty} \frac{K(A)}{A} = 0$. We show that in this case also, the ranking mechanism elicits higher quality contributions than the proportional mechanism in equilibrium.

First we prove a simple technical lemma.

LEMMA 6.1. *The minimum value of $\sum_{i=1}^n x_i^2$ subject to the constraints $x_i \geq 0$ and $\sum_{i=1}^n x_i = 1$ is $\frac{1}{n}$.*

PROOF. Note that subject to the constraints $x_i \geq 0$ and $\sum_{i=1}^n x_i = 1$, the expression $\sum_{i=1}^n x_i^2$ is minimized when $x_i = \frac{1}{n}$ for all i. To see this, note that if $x_i = \frac{1}{n} + \delta_i$ for some $\{\delta_i\}_{i=1}^n$ satisfying $\sum_{i=1}^n \delta_i = 0$, then $\sum_{i=1}^n x_i^2 = \sum_{i=1}^n (\frac{1}{n} + \delta_i)^2 = \sum_{i=1}^n (\frac{1}{n^2} + \frac{2}{n}\delta_i + \delta_i^2) = \frac{1}{n} + \sum_{i=1}^n \delta_i^2$. But $\frac{1}{n} + \sum_{i=1}^n \delta_i^2$ is minimized when $\delta_i = 0$ for all i. From this it follows that the expression $\sum_{i=1}^n x_i^2$ is minimized when $x_i = \frac{1}{n}$ for all i, and the minimum value of $\sum_{i=1}^n x_i^2$ subject to the constraints $x_i \geq 0$ and $\sum_{i=1}^n x_i = 1$ is $\frac{1}{n}$. $\square$

Now we use this to show that agents almost always choose higher quality in the ranking mechanism than in the proportional mechanism when the amount of attention available becomes large.

THEOREM 6.1. *Suppose $\lim_{A\to\infty} \frac{K(A)}{A} = 0$, $\lim_{A\to\infty} T(A) = \infty$, and α in the ranking mechanism satisfies $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ for all $j \leq k-1$. Let $q_r(A)$ denote the (possibly random) quality chosen by the contributors in some equilibrium of the ranking mechanism, and let $q_p(A)$ be an equilibrium quality in the proportional mechanism. Then $\lim_{A\to\infty} Pr(q_r(A) > q_p(A)) = 1$.*

PROOF. Suppose by means of contradiction that $Pr(q_r(A) \leq q_p(A)) \geq \gamma > 0$ for an infinite number of A. Restrict attention to values of A satisfying $Pr(q_r(A) \leq q_p(A)) \geq \gamma$, and define $q^*(A)$ to be the largest quality such that $Pr(q_r(A) < q^*(A)) \leq \frac{\gamma}{2}$ for any such A. From this it follows that $Pr(q^*(A) \leq q_r(A) \leq q_p(A)) \geq \frac{\gamma}{2}$ for all A in this subsequence. We also know from Theorem 4.2 that $\lim_{A\to\infty} q^*(A) = 1$.

Now suppose that a contributor deviates from $F_A(q)$ by making the following change: if she draws a quality $q \in [q^*(A), q_p(A)]$, then she instead chooses a quality $q + \epsilon$ for some infinitesimal amount $\epsilon > 0$; if she draws a quality $q \notin [q^*(A), q_p(A)]$, then she makes no change to her quality. We seek to show that this is a profitable deviation for a contributor for sufficiently large A in the subsequence. To do this, we first show that the additional expected cost from this deviation is bounded above by the marginal costs of producing higher quality content in the proportional mechanism, and then show that the expected benefit in the ranking mechanism from this deviation exceeds these marginal costs.

First note that this change in quality costs a contributor an additional amount no greater than $E[c(q+\epsilon) - c(q) | q \in [q^*(A), q_p(A)]]$ which, in the limit as $\epsilon \to 0$, converges to $\epsilon E[c'(q) | q \in [q^*(A), q_p(A)]] \leq \epsilon c'(q_p(A)) = \Theta(\epsilon \frac{A}{K(A)})$ (since $\beta(A) = 1$ for large A in the proportional mechanism when $\lim_{A\to\infty} \frac{K(A)}{A} = 0$). Next we calculate the expected benefits from this increase in quality.

Suppose a contributor i chooses a quality $q \in [q^*(A), q_p(A)]$. (This happens with probability at least $\gamma/2 > 0$ by assumption). Then there exists an infinite sequence $\hat{q}(A)$ such that $\lim_{A\to\infty} \hat{q}(A) = 1$ and $Pr(\frac{m_i}{T} \geq \hat{q}(A)) \geq \frac{1}{2}$ for all A. Thus if two contributors both choose qualities in $[q^*(A), q_p(A)]$, then the probability that they both receive at least $\hat{q}(A)T$ positive votes is at least $\frac{1}{4}$. Moreover, conditional on both receiving at least $\hat{q}(A)T$ positive votes, the probability a contributor receives any particular number of votes is the same for both contributors.

From Lemma 6.1, we know that if two contributors both receive at least $\hat{q}(A)T$ positive votes and they both receive any particular number of votes with the same probability, then the probability that both contributors receive an equal number of votes is at least $\frac{1}{(1-\hat{q}(A))T}$. Thus if one contributor receives an additional vote when both receive at least $\hat{q}(A)T$ positive votes, then this additional vote increases her probability of being ranked ahead of the other contributor by at least $\frac{1}{2(1-\hat{q}(A))T}$. (An agent is ranked ahead of another agent with probability $\frac{1}{2}$ when both receive the same number of votes and with probability 1 when one receives a strictly larger number of votes, so the additional vote increases the agents's probability of being ranked ahead of the other agent by a factor of $\frac{1}{2}$.)

Now the probability both contributor i and some other contributor choose qualities $q \in [q^*(A), q_p(A)]$ is at least $(\frac{\gamma}{2})^2$. And we have seen that the probability that both contributors receive at least $\hat{q}(A)T$ positive votes if they choose qualities $q \in [q^*(A), q_p(A)]$ is at least $\frac{1}{4}$. Combining this with the results in the previous paragraph shows that an additional vote for contributor i increases the probability of her being ranked ahead of a particular other contributor by at least $(\frac{\gamma}{2})^2 \frac{1}{4} \frac{1}{2(1-\hat{q}(A))T(A)} = \frac{\gamma^2}{32(1-\hat{q}(A))T(A)}$.

Now increasing quality by an infinitesimal amount $\epsilon(A) = o(\frac{1}{T(A)})$ leads to an additional positive vote with probability $\Theta(\epsilon(A)T(A))$. Therefore, increasing quality by $\epsilon(A)$ increases the probability of being ranked

ahead of a particular other contributor by an amount $\Omega(\epsilon(A)T(A)\frac{\gamma^2}{32(1-\hat{q}(A))T(A)}) = \Omega(\frac{\epsilon(A)}{1-\hat{q}(A)})$.

This implies that a contributor's change in expected ranking from increasing quality by $\epsilon(A)$ when $q \in [q^*(A), q_p(A)]$ is $\Omega(\frac{\epsilon(A)k}{1-\hat{q}(A)})$. And moving up in the rankings by one spot increases one's payoff by $\Theta(\frac{A}{k^2})$ since $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ for all $j \leq k-1$. Thus if an agent increases her quality by $\epsilon(A)$ when $q \in [q^*(A), q_p(A)]$, then her expected benefit increases by an amount $\Omega(\frac{\epsilon(A)k}{1-\hat{q}(A)}\frac{A}{k^2}) = \Omega(\frac{\epsilon(A)A}{(1-\hat{q}(A))\beta(A)K(A)})$.

Thus, increasing quality by $\epsilon(A)$ when $q \in [q^*(A), q_p(A)]$ leads to an additional benefit of $\Omega(\frac{\epsilon(A)A}{(1-\hat{q}(A))\beta(A)K(A)})$, at an additional cost of $\Theta(\frac{\epsilon(A)A}{K(A)})$. Since $\lim_{A\to\infty} q(A) = 1$, it follows that this is a profitable deviation for sufficiently large A. This contradicts the assumption that there is a sequence of equilibria for which $Pr(q_r(A) \leq q_p(A)) \geq \gamma > 0$ for an infinite number of A and proves the desired result. □

To understand the intuition behind this result, suppose for the time being that agents chose quality in the ranking mechanism according to a symmetric pure strategy, $q_r(A)$. Then if an agent makes a small change in quality from $q = q_r(A)$ to $q = q_r(A) + \epsilon$ for some $\epsilon > 0$, the agent goes from obtaining an average ranking in expectation to almost certainly being ranked near the very top. Thus such a change dramatically increases an agent's expected payoff. By contrast, in the proportional mechanism, increasing one's quality by $\epsilon > 0$ does relatively little to improve the expected attention that an agent receives. Thus incentives to produce higher quality content are greater in the ranking mechanism than in the proportional mechanism. The proof of the theorem illustrates how this intuition can be extended when the agents choose quality according to a mixed strategy.

7. DISCUSSION

In this paper, we have analyzed the widely used rank-order mechanism for displaying user contributed content in a model with strategic attention-driven contributors, and shown that the rank-order mechanism elicits contributions of optimal quality in the limit as the amount of available attention diverges. By contrast, whether equilibrium quality in the proportional mechanism becomes optimal depends on the relative rates of growth of the number of potential contributors and the number of viewers. Even when equilibrium quality in the proportional mechanism tends to the optimal quality, quality is almost always lower in the proportional mechanism than in the ranking mechanism. Thus, despite being more equitable, the proportional system creates inferior incentives for eliciting high quality contributions than the ranking mechanism.

We note that the ranking mechanism is a relatively robust mechanism, in the sense that the presence of even a large number of voters who are error prone or vote 'good' or 'bad' on all content uniformly does not affect the limiting equilibrium quality, as long as the number of voters voting according to quality continues to diverge. This is in contrast to the mechanism analyzed in [3], where even a small fraction of raters who vote thumbs down on content can lead to no content being displayed on the site.

There are a number of interesting directions for further work; we discuss two specific directions. On most sites, there are almost always some contributors who produce consistently low quality contributions, despite receiving little or no attention for it. This indicates that some subset of contributors have zero cost for producing low-quality content. An ideal mechanism in this setting would continue to elicit high quality contributions and high participation from the remaining contributors: an interesting question is how effective the ranking mechanism is when such contributors are present, and whether other mechanisms might be more effective. A second interesting direction regards questions related to malicious voters. The ranking mechanism is robust to voters who do not vote according to the model as long as they do this uniformly for all content. However, it is less clear how robust the results are in a model in which some malicious voters try to bring up specific contributions or put down others. Addressing this is an open question.

Acknowledgements

We thank Preston McAfee, David Reiley, and Michael Schwarz for helpful discussions and the anonymous reviewers for comments.

8. REFERENCES

[1] J.G. Becker, D.S. Damianov, *On the Existence of Symmetric Mixed Strategy Equilibria*, Economics Letters, 90(1), 84-87, 2006.

[2] Y. Chen, S. Jain, D. Parkes, *Designing Incentives for Online Question and Answer Forums*, ACM EC, 2009.

[3] A. Ghosh, R.P. McAfee, *Incentivizing High-Quality User Generated Content*, 20th International World Wide Web Conference (WWW), 2011.

[4] J.R. Green, N.L. Stokey, *A Comparison of Tournaments and Contracts*, Journal of Political Economy, 91(3), 349-364, 1983.

[5] B. Holmstrom, *Moral Hazard in Teams*, Bell Journal of Economics, 13(2), 324-340, 1982.

[6] P. Hummel, *Proportional Versus Winner-Take-All Electoral Vote Allocations*, Public Choice, forthcoming.

[7] S. Jain, D. Parkes, *The Role of Game Theory in Human Computation Systems* (Position Paper), HCOMP 2009.

[8] S. Jain, D. Parkes, *A Game-Theoretic Analysis of the ESP Game*, Working Paper.

[9] E.P. Lazear, *Pay Equality and Industrial Politics*, Journal of Political Economy, 97(3), 561-580, 1989.

[10] E.P. Lazear, S. Rosen, *Rank-Order Tournaments as Optimum Labor Contracts*, Journal of Political Economy, 89(5), 841-864, 1981.

[11] A. Lizzeri, N. Persico, *The Provision of Public Goods under Alternative Electoral Systems*, American Economic Review, 91(1), 225-239, 2001.

[12] B.J. Nalebuff, J.E. Stiglitz, *Prizes and Incentives: Towards a General Theory of Compensation and Competition*, Bell Journal of Economics, 14(1), 21-43, 1983.

[13] T. Persson, G. Tabellini, *The Size and Scope of Government: Comparative Politics with Rational Politicians*, European Economic Review, 43(4), 699-735, 1999.

[14] J.M. Snyder, *Election Goals and the Allocation of Campaign Resources*, Econometrica, 57(3), 637-660, 1989.

APPENDIX

Proof Of Theorem 4.3. Suppose that there is an infinite sequence of values of A such that, for each A, there is a symmetric equilibrium to the ranking mechanism with $\beta(A) < 1$. We know from the conditions on indifference between entry and exit that this implies that $E[c(q)|q \sim$

$F_A(q)] = \Theta(\frac{A}{\beta(A)K(A)})$ for all A in this sequence. Thus if $q^*(A)$ denotes the largest quality in the support of $F_A(q)$, then we know that $c(q^*(A)) = \Omega(\frac{A}{\beta(A)K(A)})$ for all A in this sequence. Thus there exists some function $g(A) = \Theta(\frac{A}{K(A)})$ independent of $T(A)$ such that $c(q^*(A)) \geq g(A)$ and $q^*(A) \geq c^{-1}(g(A))$ for all A in this sequence. Thus if $\hat{q}(A) = c^{-1}(g(A))$, then $\lim_{A\to\infty} \hat{q}(A) = 1$ and $q^*(A) \geq \hat{q}(A)$ for all A in this sequence.

Now note that if a contributor chooses some quality $q = q^*(A) + \epsilon(A)$ instead of choosing quality $q = q^*(A)$ for some infinitesimal amount $\epsilon(A) = o(\frac{1}{T(A)})$ (which may either be positive or negative), then the contributor changes the probability that she receives an additional positive vote by an amount $\Theta(\epsilon(A)T(A))$. Thus the contributor changes the probability that she receives a higher ranking than a particular other contributor by an amount $O(\epsilon(A)T(A))$, and the contributor changes the expected number of other contributors that the contributor receives a higher ranking than by an amount $O(\epsilon(A)T(A)\beta(A)K(A))$. And when a contributor moves up in the rankings by one spot, the contributor increases her payoff by an amount $\Theta(\frac{A}{k^2})$ since $\alpha_j(k) - \alpha_{j+1}(k) = \Theta(k^{-2})$ for all $j \leq k-1$. Thus if a contributor chooses some quality $q = q^*(A) + \epsilon(A)$ instead of choosing quality $q = q^*(A)$ for some infinitesimal amount $\epsilon(A)$, then the contributor changes her expected benefits by an amount $O(\epsilon(A)T(A)\beta(A)K(A)\frac{A}{(\beta(A)K(A))^2}) = O(\epsilon(A)T(A)\frac{A}{\beta(A)K(A)})$.

Thus if $b_A(\epsilon(A))$ denotes the difference between the expected benefits from attention that a contributor obtains by choosing $q = q^*(A) + \epsilon(A)$ instead of choosing quality $q = q^*(A)$ for some infinitesimal amount $\epsilon(A)$, then $|b_A(\epsilon(A))| = O(|\epsilon(A)T(A)\frac{A}{\beta(A)K(A)}|)$, meaning $b'_A(0) = O(T(A)\frac{A}{\beta(A)K(A)})$. But in order for a contributor to not be able to profitably deviate from choosing quality $q = q^*(A)$, it is necessary that $b'_A(0) = c'(q^*(A))$. Thus $c'(q^*(A)) = O(T(A)\frac{A}{\beta(A)K(A)})$.

Recall that $c(q^*(A)) = \Omega(\frac{A}{\beta(A)K(A)})$. Combining this with the result in the previous paragraph shows that $\frac{c'(q^*(A))}{c(q^*(A))} = O(T(A))$. Thus if $\overline{q}(A)$ denotes a minimizer of $\frac{c'(q)}{c(q)}$ subject to the constraint $q \in [\hat{q}(A), 1]$, and $T(A) = \Theta\left(\log \frac{c'(\overline{q}(A))}{c(\overline{q}(A))}\right)$, then this implies that $q^*(A) < \hat{q}(A)$ for sufficiently large A, which would contradict the fact that $q^*(A) \geq \hat{q}(A)$ for all A in this sequence.

Now since $\lim_{q\to 1} c(q) = \infty$, it follows that $\lim_{q\to 1} \log c(q) = \infty$, $\lim_{q\to 1} \frac{d}{dq} \log c(q) = \infty$, and $\lim_{q\to 1} \frac{c'(q)}{c(q)} = \infty$. Thus $\lim_{A\to\infty} \overline{q}(A) = 1$ implies $\lim_{A\to\infty} \frac{c'(\overline{q}(A))}{c(\overline{q}(A))} = \infty$. Combining this with the results in the previous paragraphs shows that if $T(A) = \Theta\left(\log \frac{c'(\overline{q}(A))}{c(\overline{q}(A))}\right)$, then $\{T(A)\}_{A=1}^{\infty}$ is a sequence satisfying $\lim_{A\to\infty} T(A) = \infty$ such that $\beta(A) = 1$ must hold for sufficiently large A. $\square$

Proof Of Lemma 5.1.

$$E[V(m_i, m_{-i})|(\beta, q_i, q_{-i})] = E[\frac{m_i}{m_i + \sum_{j\neq i} m_j} A|(\beta, q_i, q_{-i})].$$

Note that for a given value of k, as T goes to infinity, $\frac{m_i}{m_i+\sum_{j\neq i} m_j} = \frac{m_i/T}{m_i/T+\sum_{j\neq i} m_j/T}$ converges in probability to $\frac{q_i}{q_i+\sum_{j\neq i} q_j}$. Now if $q_j = q^*$ for all $j \neq i$, $\frac{q_i}{q_i+\sum_{j\neq i} q_j} = \frac{q_i}{q_i+(k-1)q^*}$. Thus as A and $K(A)$ go to infinity, $\frac{q_i}{q_i+(k-1)q^*}$ converges in probability to $\frac{q_i}{q_i+\beta(A)K(A)q^*}$. Thus $\frac{d}{dq_i} E[V(m_i, m_{-i})|(\beta, q_i, q^*)] = \Theta(\frac{d}{dq_i} \frac{q_i A}{q_i+\beta(A)K(A)q^*})$ for large A, $K(A)$, and $T(A)$. But $\frac{d}{dq_i} \frac{q_i A}{q_i+\beta(A)K(A)q^*} = \Theta(\frac{A}{\beta(A)K(A)})$. Thus in equilibrium it must be the case that $c'(q_p(A)) = \Theta(\frac{A}{\beta(A)K(A)})$. $\square$