
Social Status and the Design of Optimal Badges

Nicole Immorlica

Greg Stoddard

Vasilis Syrgkanis

Abstract

Many websites rely on user-generated content to provide value to consumers. Often these websites incentivize user-generated content by awarding users *badges* based on their contributions. These badges confer value upon users as a symbol of social status. In this paper, we study the optimal design of a system of badges for a designer whose goal is to maximize contributions. We assume users have heterogenous abilities drawn from a common prior and choose how much effort to exert towards a given task. A user's ability and choice of effort determines the level of contribution he makes to the site. A user earns a badge if his contribution surpasses a pre-specified threshold. The problem facing the designer then is how to set badge thresholds to incentivize contributions from users. Our main result is that the optimal total contribution can be well-approximated with a small number of badges. Specifically, with a number of badges just logarithmic in the number of users, the designer can generate at least half as much contribution as the optimal unconstrained (non-badge-based) mechanism.

1 Introduction

A number of popular sites on the web today are driven by user-generated content. Review sites such as Yelp and TripAdvisor need users to rate and review restaurants and hotels, social news aggregators like Reddit rely on users to submit and vote upon articles from around the web, and question and answer sites like Stack Overflow and Quora depend on their users to ask good questions and provide good answers. Stack Overflow, a user-driven Q&A site for programming questions, is one of the most successful of these sites; over five million questions have been asked, with over 60% of those questions receiving a satisfactory answer. An often-cited reason for the high quality of the Stack Overflow community is the clever system of *reputation points* and *badges* it employs. While the main actions on SO are posting and answering questions, a variety of other activities such as flagging posts for moderation, rating contributions, etc are all necessary to maintain a healthy and vibrant discourse on the site. Stack Overflow motivates question and answers through the use of reputation points but all of these secondary activities are motivated by awarding *badges*. A badge is a small symbol displayed on a user's profile and posts and is typically awarded for accomplishing a fixed task, such as answering ten questions, reviewing 100 restaurants, or completing an online course. Most badge systems are not designed to promote competition among users but simply to reward individual accomplishments. As such, badges do not have the exogenously imposed scarcity that ranking systems (top 10 contributors) or contest systems (medals for first, second, and third place) have. In this paper, we examine the incentives created by these badge schemes and focus on designing badge schemes that maximize the user contributions to a site.

Badge systems have become particularly popular in web communities over the past few years. In addition to Stack Overflow, the Huffington Post recently implemented a badge system to reward actions from their commenters; these badges are now prominently displayed next to usernames in the comment sections. Foursquare, a mobile social networking app, awards badges for visiting certain locations or accomplishing certain tasks while there. The Mozilla Foundation is leading an initiative called the Open Badge Project which hopes to set an open standard for awarding, collecting, and displaying badges. Their ultimate goal is provide a persistent collection of badges that can be displayed anywhere on the web as a proof of acquired skills and achievements.

With all this excitement and energy surrounding badge systems, one fundamental question emerges: why do people care about badges? On the surface level, a badge is just a small sticker or group of pixels on a profile, so why should it be that badges incentivize people to exert effort to earn them? In this work, we posit the primary value of a badge is derived from its ability to confer social status to its owner. The most valuable badges not only signal that a user has accomplished a certain task but they also signal that relatively few users have earned them. Indeed, many systems include a top level of badges such as the “gold” badges on Stack Overflow or the “Superuser” badge on Huffington Post, which are described to be difficult to earn and only awarded to the most committed users. As more users earn a particular badge, that badge loses its ability to distinguish a member within the community and thus becomes less valuable. This work studies optimal badge design when badges are a means for establishing relative social status.

The goal of the designer is to maximize the total contributions that users make towards a single task on the site. Then the main question is how to design a set of badges to maximize the total output that users contribute to that task? How many badges does the designer need to set and how should they be awarded? For example, if Yelp wants to maximize the number of reviews they receive, should they award just one badge for contributing at least 50 reviews or should they have one badge for 10 reviews and another badge for 100 reviews? In our model, users have heterogeneous abilities for certain tasks on a site; users of low ability need to exert more effort to produce the same output as high ability users. This effort is costly to users, so they will balance the value they get from earning a badge with the effort it takes to do so. Users all have the same base value for a badge (normalized to 1) but the value of a badge diminishes as more users acquire it. Specifically, if k users earn a particular badge, the value of that badge decreases by a factor of $\frac{1}{k}$. We generically refer to this reduction in value due to the number of people earning a badge as the *congestion* on a badge.

In keeping with the general ethos of badge systems, we restrict the designer to only using *threshold badges*. A threshold badge is one which is awarded to any user who contributes more than a fixed threshold. Thus the designer cannot exogenously impose ex-post scarcity and implement schemes such as awarding only the top 10 contributors or awarding a badge for first place, second place, etc. Instead, the designer may specify badges with different thresholds, i.e. one badge for answering 10 questions, another badge for answering 100 questions, etc. This structure combined agents’ social status concerns allows for discrimination amongst users of varying abilities, such that the users in the highest tier of abilities will earn the top badge, users with ability in the next tier will earn the second highest badge, and so forth.

Our main result is that although such threshold badge systems are not optimal, they are approximately optimal. The designer can always implement a threshold badge system that achieves at least half of the total output of the optimal mechanism. To achieve this guarantee, the designer only needs to specify a number of badges that is logarithmic in the number of users, indicating that simple badge systems are give good performance. We prove our results by using a connection between user-generated contributions and all-pay auctions, where the total user contribution is equal to the revenue of the corresponding auction. In particular, we draw an analogy between the allocation function of an all-pay auction and the reduction in a badge’s value in equilibrium due to congestion. Using a trick of Chawla and Hartline [4], we transform our problem to the canonical quasi-linear setting and use tools optimal mechanism design to derive the structure of the optimal contest in this setting. We then show that a simple badge system can approximate the interim

allocation function of the optimal contest and conclude that the total contributions of the badge system must be a good approximation of the optimal mechanism. Last we use the structure of the revenue function to give an efficient algorithm for computing the optimal way to set thresholds if the designer is restricted to setting exactly m badges.

The remainder of the paper is laid out as follows. In section 3 we formally describe our model. In section 4, we characterize the equilibria in this model and describe the relationship to the all-pay auction. In section 5 we derive the form of the optimal mechanism and use that characterization in section 7 to prove that a simple threshold badge mechanism is a 2-approximation of the optimal mechanism. We describe our algorithm to compute the optimal set of m badges in Section 6 and wrap up with discussion and future work in section 8.

2 Related Work

The papers that are most closely related to ours is that of Moldovanu, Sela, and Shi [10] and Dubey and Geanakoplos [5], which consider the same contest for status but the former assumes incomplete information while the latter assumes a complete information game. Although the models are similar, their main results regarding the structure of the optimal contest structure illustrates the differences in the complete versus incomplete information settings. While [10] finds that assigning players unique ranks based on output (for a large class of distributions) is optimal, [5] finds the optimal scheme will often group players with a range of effort into a single rank. Similar to our work, [10] shows that agents into one of two ranks (winners or losers) can elicit at least half as much effort as the optimal ranking. Hopkins and Kornienko [8] consider a model where consumers care about their social status, as determined by their consumption of a “positional” good, and study how shifts in equality in the societal distribution of income affect conspicuous consumption. Although their setting is quite different, there is a similarity in spirit of using tools from auction theory as a means of analyzing behaviors induced by social status concerns.

There’s a growing literature on the role and design of incentives systems in online settings such as question and answer websites (e.g. StackOverflow, Quora, etc). This line of the literature considers how to award virtual points on a website in order to maximize some objective for the designer. In the context of Q&A sites, Ghosh and Hummel [7] show that a simple scheme of giving a large amount of points to the winning answer and a smaller amount of points to each remaining answer can maximize a large class of objective functions. Further, they show always choosing the best contribution to be the winning contribution actually *decreases* the designer’s value. Jain, Chen, and Parkes [9] study the design of Q&A reward schemes in a more complex environment where users contribute answers sequentially and the value of a contribution may vary based on the set of other contributions the designer receives, such as the case where contributions may compliment each other. They show there is no mechanism which can achieve the efficient outcome for every type of valuation for the designer. Hartline et al [4] study crowdsourcing contests where the principal only receives value from the highest quality submission, as is the case with crowdsourcing tasks such as the design of a poster or a logo. Using tools from optimal mechanism design, they show the optimal crowdsourcing contest gives all the prize money to the best contribution and nothing to all other contributions. Cavallo and Jain [3] shows that winner-take-all contests are not welfare-maximizing when there’s uncertainty in the conversion from effort to quality and instead propose a mechanism where every contributor earns a minimum amount of money.

Easley and Ghosh [6] consider a model of earning badges in the presence of noisy observations. They show that when a user’s value for a badge is exogenously determined, the optimal badge structure is to award the badge only to a fixed number of top contributors. They also extend their model to the case when a user’s value for a badge changes based on the number of other users who have earned that badge and interestingly show that a designer can take advantage of user uncertainty over the number of badge earners in order to incentivize more contributions. Depending on the shape of the badge value function (as a function of mass of users who have earned that badge), hiding the number of winners may increase or decrease contributions

in equilibrium. Kleinberg et al [1] find empirical evidence that people are motivated by badges; using the public logs from Stack Overflow, they show that as users get closer to the threshold for winning a particular badge, they spend comparatively more effort performing the action associated with that badge. They also develop a theoretical model describing how a single user may change his behavior to earn badges, and how to design badge thresholds under this model.

3 Badge Utilities with Congestion

We consider the following game theoretic model of contributions to a user-driven site. There are n players each having ability v_i with respect to some task which the designer wishes to incentivize. Ability is a context-specific trait such as programming knowledge in the case of Stack Overflow or inherent ability to write good restaurant reviews. If player i invests effort e_i on this task, then she produces an output of $b_i = v_i \cdot e_i$. A user's output is their observed contribution to the site, such as the number of questions answered or the quality of a restaurant review. The goal of the designer is to maximize the sum of the output of all players $\sum_i b_i$.

Badge Mechanism. Before players invest effort, the designer defines a *badge mechanism*, which is a set of m badges associated with this task. Specifically, each badge $j \in [m]$ is awarded to a user if her output b_i is at least some threshold θ_j . As an example badge mechanism, if the designer sets the badge thresholds $\theta = (10, 100, 200)$, it indicates that he awards badge 1 for answering 10 questions, badge 2 for 100 questions, and badge 3 for 200 questions. Without loss of generality, we assume $\theta_1 \leq \dots \leq \theta_m$. Given a player's output b_i we denote a player's *rank* $r(b_i)$ as the highest badge that player i wins, i.e. the highest threshold θ_j such that $b_i \geq \theta_j$. If a player doesn't win any badge then we assume his rank is zero.

Player Utilities. In this paper, a player's value for a badge is determined endogenously by its ability to signal social status within the community. If a player receives multiple badges, they only value the highest badge they won. Thus their value is determined by their rank. Given a profile of the ranks of all players $r = (r_1, \dots, r_n)$, a player's value for her rank is some continuous decreasing function $f : [0, 1] \rightarrow \mathbb{R}^+$ of the proportion of people that have an equal or higher rank (and zero in the case when she receives no badge). For example, $f(t)$ could be the proportion of the population that a player beats, $f(t) = 1 - t$, or it could be the inverse of the proportion of the population that is at least as good as the player $f(t) = \frac{1}{t}$. We return to these two examples later in the paper.

Let $q_i(r) = \frac{|\{j \in [n] : r_j \geq r_i\}|}{n}$ denotes this proportion. A player's utility is quasi-linear in their value for their rank and the effort they exerted:

$$\hat{u}_i(r, e_i) = f(q_i(r)) \cdot 1_{r_i > 0} - e_i \quad (1)$$

We refer to the first term as the *status value* and we denote it as $x_i(r) = f(q_i(r)) \cdot 1_{r_i > 0}$.

Given a particular badge mechanism, the rank of a player is determined by the profile of outputs $b = (b_1, \dots, b_n)$. Also observe that a player's effort is a function of their output and ability, so we can rewrite a player's utility as:

$$\hat{u}_i(b; v_i) = x_i(r(b)) - \frac{b_i}{v_i} \quad (2)$$

Additional Assumptions. We assume that each player's ability v_i is private information and is drawn independently and identically from a commonly known, atomless, and regular¹ distribution $G(\cdot)$ with support $[0, \bar{v}]$ and density $g(\cdot)$. Hence, any badge mechanism defines a game of incomplete information among the n players. We analyze the Bayes-Nash equilibria of this game. A BNE is a profile of mappings from abilities

¹A distribution G is regular if $v - \frac{1-G(v)}{g(v)}$ is (weakly) monotone in v

v_i to outputs $b_i(v_i)$ such that each player maximizes their utility in expectation over the abilities of the rest of the players. For all v_i, b'_i :

$$\mathbb{E}_{v_{-i}} [u_i(b(v); v_i)] \geq \mathbb{E}_{v_{-i}} [u_i(b'_i, b_{-i}(v_{-i}); v_i)] \quad (3)$$

Since each player is maximizing his utility conditional on his value, we can equivalently assume that a player's utility is:

$$u_i(b; v_i) = v_i \cdot \hat{u}_i(b; v_i) = v_i \cdot x_i(r(b)) - b_i \quad (4)$$

because any strategy which maximizes their original utility \hat{u} also maximizes u . This new formulation allows us to see that this game is equivalent to an all-pay auction where a player's allocation is a function of the ranks assigned to all players.

To summarize the timing of the game, the designer first announces a set of badges θ . Each agent then learns his ability v_i and simultaneously decides how much output b_i to contribute to the site. The designer observes the output of each agent and awards badges based on the thresholds for each badge.

Main Question. The designer of the badge mechanism is interested in maximizing the expected total output produced by the players: $\mathbb{E}[\sum_i b_i(v_i)]$, which is exactly the revenue of the corresponding all-pay auction. We consider two main questions:

1. *Given a distribution of abilities, how should the designer set badge thresholds to maximize expected revenue?*
2. *How many badges are sufficient to closely approximate the revenue achievable by any mechanism in this context?*

4 Uniqueness and Characterization of Bayes-Nash Equilibrium

We begin our analysis by characterizing the Bayes-Nash equilibria induced by any badge mechanism. This model falls within the class of symmetric ranking mechanisms of Hartline and Chawla [4] and therefore has a unique Bayes-Nash equilibrium. This equilibrium is symmetric and monotone, i.e. each player uses the same strategy mapping $b_i(v_i) = b(v_i)$ and this mapping is monotone in the player's ability. Thus we just need to characterize this unique symmetric Bayes-Nash equilibrium in order to analyze this game.

We start by observing that if a player gets rank r_i , her output should be exactly the threshold to win that badge, θ_{r_i} , because producing more output would cost more effort and would not increase her value. Since the equilibrium mapping is monotone in ability, the equilibrium is defined by a set of thresholds in the ability space of the players a_1, \dots, a_m , such that if player i has ability $v_i \in [a_t, a_{t+1})$ then he produces output $b(v_i) = \theta_t$. If $v_i < a_1$ then $b(v_i) = 0$ and if $v_i \geq a_m$ then $b(v_i) = \theta_m$.

So given a badge mechanism with contribution thresholds $\theta = (\theta_1, \dots, \theta_m)$, the resulting BNE is defined by a vector of thresholds on player ability, but it remains to compute those ability thresholds. We compute these thresholds using the Bayes-Nash equilibrium characterization of Myerson for quasi-linear utility, single-dimensional type environments. For a fixed mapping $b(\cdot)$, we denote with $\tilde{x}_i(v_i) = \mathbb{E}_{v_{-i}} [x_i(r(b(v)))]$, the expected status value player i gets, assuming each player uses strategy $b(\cdot)$. $\tilde{x}_i(v_i)$ is often referred to as the interim allocation of player i . Since our setting is completely symmetric, the interim allocation function is the same for all players, and we will denote it with $\tilde{x}(\cdot)$. Now observe that the expected utility of a player with ability v_i , assuming that everyone employs strategy $b(\cdot)$, is :

$$\mathbb{E}_{v_{-i}} [u_i(b(v); v_i)] = v_i \tilde{x}(v_i) - b(v_i) \quad (5)$$

Myerson's equilibrium characterization, when applied to our setting, states that $b(\cdot)$ is an equilibrium only if $\tilde{x}(v_i)$ is monotone in v_i and:

$$b(v_i) = v_i \tilde{x}(v_i) - \int_0^{v_i} \tilde{x}(z) dz \quad (6)$$

Additionally, if the function $b(\cdot)$ spans the whole region of feasible bids, then the latter is an if and only if statement. The latter won't be true in our setting. However, the proposed strategy $b(\cdot)$ will include all the feasible undominated bids, which is all the badge threshold bids. Hence, if we find a strategy that satisfies the above equation then that strategy will be an equilibrium.

Now consider the mapping $b(\cdot)$ that corresponds to some vector of value thresholds a as explained previously. For a player with value $v_i \in [a_t, a_{t+1}]$ we can compute explicitly his expected status value as follows: we know that every player with value $v_j \geq a_t$ will be assigned a rank at least as high as player i . Thus the status value of player i is the same as the social status of a player with value a_t and is simply the expected value of $f\left(\frac{T+1}{n}\right)$ where T is the random variable denoting the number of players other than player i that have value above a_t . This can be computed as follows:

$$\tilde{x}(a_t) = \sum_{\nu=0}^{n-1} f\left(\frac{\nu+1}{n}\right) \cdot \beta_{\nu, n-1}(1 - F(a_t)) \quad (7)$$

where $\beta_{\nu, n}(q) = \binom{n-1}{\nu} \cdot q^\nu \cdot (1-q)^{n-1-\nu}$, denotes the Bernstein polynomial. It is also easy to see that if $f(q)$ is a strictly increasing function then $\tilde{x}(a_t)$ is also strictly increasing. Additionally, since G is an atomless distribution, $\tilde{x}(a_t)$ is continuous and differentiable and $\tilde{x}(\bar{v}) = f(1/n)$ and $\tilde{x}(0) = f(1)$.

Using the step nature of the interim allocation function and the fact that by the initial analysis, a player with value $v_i \in [a_j, a_{j+1}]$ bids θ_j , the equilibrium characterization (6) simply becomes:

$$\theta_j = a_1 \cdot \tilde{x}(a_1) + \sum_{t=2}^j a_t \cdot (\tilde{x}(a_t) - \tilde{x}(a_{t-1})) \quad (8)$$

This relation is depicted in Figure 1. Observe that the above defines a system of m equalities that have




Figure 1: Relation between interim allocation probability and badge thresholds at equilibrium.

a unique solution. An equivalent way of phrasing the above equations is that: $a_1 \cdot \tilde{x}(a_1) = \theta_1$ and for $j \in [2, m]$:

$$a_j \cdot (\tilde{x}(a_j) - \tilde{x}(a_{j-1})) = \theta_j - \theta_{j-1} \quad (9)$$

Thus, given a profile of badge thresholds θ , we can iteratively solve for the profile of ability thresholds a as follows: find a solution to the equation $\theta_1 = a_1 \tilde{x}(a_1)$. Observe that $\bar{v} \cdot \tilde{x}(\bar{v}) = \bar{v} f(1/n)$, $0 \cdot \tilde{x}(0) = 0$,

$v \cdot \tilde{x}(v)$ is continuous increasing. If $\theta_1 < \bar{v}f(1/n)$ then a unique solution exists. Otherwise, no player is willing to bid as high as θ_1 and the recursion stops. Subsequently, find the solution a_2 to the equation: $\theta_2 - \theta_1 = a_2(\tilde{x}(a_2) - \tilde{x}(a_1))$. For similar reason, either a unique such solution exists or no player (not even a player with ability \bar{v}) is willing to bid θ_2 rather than bid θ_1 and we can stop the recursion. Then solve for a_3, \dots, a_m in the same way.

The following corollary of the above discussion summarizes a property that will prove very useful in subsequent sections:

Corollary 1. *For any vector $a = (a_1, \dots, a_m)$ of thresholds in the ability space, there exists a vector of badge thresholds $\theta = (\theta_1, \dots, \theta_m)$, such that at the unique Bayes-Nash equilibrium of the corresponding badge mechanism, a player with ability $v_i \in [a_j, a_{j+1})$ will win badge j , a player with ability $v_i \geq a_m$ will win badge m and a player with ability $v_i < a_1$ will win no badge. When this is the case, we say θ implements a . Similarly, every vector of badge thresholds implements a unique vector of ability thresholds.*

Instead of setting badge thresholds, a designer could optimize over thresholds in the ability space and then compute the badges needed to implement the ability thresholds. Intuitively, this design can be thought of as the direct-revelation mechanism for a badge mechanism. This formulation of the problem is easier because we no longer need to worry about the induced equilibrium behavior. For the remainder of the paper, we focus on this version of the problem.

Another convenient way of characterizing a badge mechanism is through the notion of the upper quantile of a distribution. Given a distribution G of abilities, the quantile of ability v_i is the probability that a random sample from G is at least as high as v_i :

$$q(v_i) = 1 - G(v_i) \quad (10)$$

For an atomless continuous distribution G , there is a one-to-one correspondence between abilities and quantiles. Define $v(q_i) = G^{-1}(1 - q_i)$ to be the ability corresponding to quantile q_i . We can also define the interim allocation of a player as a function of his quantile rather than his ability.

$$\hat{x}(q_i) = \tilde{x}(v(q_i)) \quad (11)$$

It is easy to see that there is a one-to-one correspondence between a vector of thresholds $a = (a_1, \dots, a_m)$ in ability space and a vector of thresholds $\kappa = (\kappa_1, \dots, \kappa_m)$ in quantile space. A higher ability corresponds to a lower quantile, with $q(\bar{v}) = 0$ and $q(0) = 1$, and therefore if $a_1 \leq \dots \leq a_m$, we have $\kappa_1 \geq \dots \geq \kappa_m$. Hence, the designer can equivalently think of designing his thresholds in quantile space rather than ability space. By the above observation and Corollary 1, for any vector of thresholds κ in the quantile space, there exists a vector of badge thresholds θ that will implement κ in the unique Bayes-Nash equilibrium.

Given a badge mechanism defined by a vector of quantile thresholds κ , the interim allocation of a player with quantile $q_i \in [\kappa_{t+1}, \kappa_t]$ is simply:

$$\hat{x}(q_i) = \hat{x}(\kappa_t) = \sum_{\nu=0}^{n-1} f\left(\frac{\nu+1}{n}\right) \beta_{\nu, n-1}(\kappa_t) \quad (12)$$

5 Optimal Ranking Mechanism

In the previous section, we observed that a badge mechanism induces a ranking among players; the players who win the highest badge are assigned the highest rank, players who win the second highest badge are assigned the second highest rank, and so on. These badge-induced rankings are just one example of a general mechanism which collects contributions from agents and assigns ranks based on that output. For example, we could collect contributions and then assign the highest contributor to the highest rank, the second highest contributor to the next rank, etc. Another mechanism could assign the 10 highest contributors to the highest

rank, then the next 100 highest contributors to the second highest rank, and the lowest rank to every other player. Each of these mechanisms would incentivize players to invest effort and contribute to the site because it gives them a way to achieve social status. In this section, we characterize the optimal ranking mechanism. We find the optimal mechanism violates our design constraint that we can only use threshold badges and thus badge systems are suboptimal. However, we use this characterization in the next section to prove that a simple badge mechanism provides a good approximation to the optimal mechanism.

We first introduce the optimal mechanism design problem. Suppose the designer is allowed to run an arbitrary mechanism that asks each player to report their ability, and then outputs a rank profile $r = (r_1, \dots, r_n)$ and an output profile $b = (b_1, \dots, b_n)$, such that each player i contributes b_i and is assigned rank r_i . Each player's utility for any such output is of the form $u_i(r, b_i) = v_i \cdot \tilde{x}(r) - b_i$, where $\tilde{x}(r)$ is as defined in the Section 3. The only restriction we make on the mechanism is that each player's expected utility is non-negative, i.e. it is rational for participate in the mechanism. What is the mechanism that produces the highest expected revenue in Bayes-Nash equilibrium?

This scenario is the classic single-parameter optimal mechanism design setting with quasi-linear utilities. By the Revelation Principle the designer can restrict himself to direct Bayesian incentive compatible mechanisms of the form: players report their abilities, the designer decides an allocation of ranks based on the reported abilities, and then asks the players to output some b_i such that truthfully reporting ability is an equilibrium for all players in the resulting game of incomplete information.

Before computing the revenue-optimal allocation, we compute the welfare-optimal allocation. The social welfare in this setting is defined as the sum of all agents utilities plus the utility for the designer. The utility for the designer is equal to the sum of the output for the agents. For a given vector of abilities v , an allocation of rankings r and a vector of agent outputs b , the social welfare is as follows.

$$SW(v, r, b) = SW_{players} + SW_{designer} = \sum_i v_i \cdot \tilde{x}(r) \quad (13)$$

In words, each agent's contribution to the social welfare is equal to his ability weighted by status value. The following lemma shows welfare is maximized² by assigning assigning each player a different rank in decreasing order of ability.

Lemma 2. *If $v_1 \geq v_2 \geq \dots \geq v_n$ then the optimal social welfare is achieved by assigning $r_1 = n, r_2 = n - 1, \dots, r_n = 1$, producing welfare of:*

$$\sum_i v_i \cdot x(r) = \sum_i v_i \cdot f\left(\frac{i}{n}\right)$$

Proof. The statement follows by the following arguments: first it is easy to see that the rank should be monotone non-decreasing in the value, since if $v_i > v_j$ and $r_i < r_j$ then we can increase welfare by swapping the ranks of player i and player j . Additionally, if for some set of values $v_i \geq v_{i+1} \geq \dots \geq v_{i+k}$ we have $r_i = \dots = r_{i+k}$ then by discriminating v_i to be strictly higher than the rest of the values increases the welfare. More concretely, for any $j > i$ we can set $r'_j = r_j + 1$. The satisfaction of all players $j > i$ doesn't change since the number of people that have rank at least as high as them remains the same. Additionally, the satisfaction of player i strictly increases, since the number of people ranked at least as high as him, strictly decreased. ■

Myerson's characterization states that the expected revenue of any mechanism is equal to its expected *virtual welfare*, the sum of each agent's *virtual value*. To continue the analogy between the all-pay auction setting

²This lemma is not characterizing the welfare-optimal because we transformed utilities in equation 4. However, this lemma is useful for characterizing the revenue-optimal mechanism.

and our model, we analogously define the *virtual ability* of an agent with ability v_i to be

$$\phi(v_i) = v_i - \frac{1 - G(v_i)}{g(v_i)} \quad (14)$$

Then the virtual welfare of any ranking mechanism is $\sum_i \phi(v_i) \tilde{x}(v_i)$. The expected total output of any mechanism in any equilibrium profile $(b_1(\cdot), \dots, b_n(\cdot))$ is equal to the virtual welfare.

$$\mathbb{E}_v \left[\sum_i b_i(v_i) \right] = \mathbb{E}_v \left[\sum_i \phi(v_i) \tilde{x}(v_i) \right] \quad (15)$$

This characterization and lemma 2 imply that the optimal mechanism orders players by decreasing virtual ability and assigns them a unique rank, so long as their virtual ability is positive. Players, with negative virtual value are assigned rank 0. In section 3 we assumed the ability distribution G is regular, so virtual ability is monotone in ability and ordering by virtual ability is equivalent to ordering by ability. The optimal mechanism then asks each player to contribute an output of $b_i(v_i)$ according to the payment identity (6).

Corollary 3 (Optimal Ranking Mechanism). *If players are distributed i.i.d. according to a regular distribution G , then the optimal mechanism asks from players to report abilities, assigns rank 0 to any player with ability $v_i \leq \eta$, where η is the solution to the equation $\phi(\eta) = 0$ (monopoly reserve), and assigns distinct ranks to the rest of the players in decreasing order of their abilities. Last it asks from each player to submit an output:*

$$b_i(v_i) = v_i \tilde{x}(v_i) - \int_0^{v_i} \tilde{x}(z) dz \quad (16)$$

where

$$\tilde{x}(v_i) = \sum_{t=0}^{n-1} f\left(\frac{t+1}{n}\right) \beta_{t,n-1}(1 - F(v_i)) \quad (17)$$

Equivalently, the optimal mechanism can be described in the quantile space, as ordering players in increasing quantile after discarding player with quantile higher than the quantile κ^* corresponding to the monopoly reserve η and asking a player with quantile q_i to submit a bid of:

$$b_i(q_i) = v_i(q_i) \hat{x}(q_i) - \int_1^{q_i} \hat{x}(q) dq \quad (18)$$

where

$$\hat{x}(q_i) = \sum_{\nu=0}^{n-1} f\left(\frac{\nu+1}{n}\right) \beta_{\nu,n-1}(q_i) \quad (19)$$

for any $q_i \leq \kappa^*$ and 0 otherwise. The revenue of the mechanism can be computed through the use of the revenue function $R(q) = q \cdot v(q)$ as follows (see Hartline [?] or Bulow and Roberts [2] for details):

$$\mathbb{E}_v \left[\sum_i b_i(v_i) \right] = n \int_0^{\kappa^*} R'(q) \cdot \hat{x}(q) dq = n \left(R(\kappa^*) \hat{x}(\kappa^*) - \int_0^{\kappa^*} R(q) \hat{x}'(q) dq \right) \quad (20)$$

Observe that the interim allocation of a player under a badge mechanism is equal to the interim allocation of the lowest ability player in his badge under the optimal mechanism. Thus the interim allocation of a badge mechanism can be thought of as a downwards rounding of the interim allocation of the optimal mechanism.

6 Optimal Badge Mechanism with Fixed Number of Badges

Before analyzing how badges compare to the optimal hypothetical mechanism, we first turn to the question of computing the optimal badge mechanism, subject to the restriction of using only m badges. We give an algorithmic result, showing that there exists a fully polynomial time approximation scheme that computes the optimal badge thresholds.

Using our equilibrium characterization we can compute the expected revenue produced by a set of badges with thresholds $\kappa = (\kappa_1, \dots, \kappa_m)$ in the quantile space as follows: For notational convenience assume that $\kappa_0 = 1$ with $\hat{x}(\kappa_0) = 0$ and $\kappa_{m+1} = 0$ with $\hat{x}(\kappa_{m+1}) = 1$. If a player has quantile $q_i \in [\kappa_j, \kappa_{j+1}]$ then by transforming Equation (8) we get that he submits a bid of $\theta_j = \sum_{t=1}^j v(\kappa_t) (\hat{x}(\kappa_t) - \hat{x}(\kappa_{t-1}))$. The latter happens with probability $\kappa_j - \kappa_{j+1}$. Thus a player's expected payment is:

$$\begin{aligned}
\mathbb{E}_{v_i}[b(v_i)] &= \sum_{j=1}^m \sum_{t=1}^j (\kappa_j - \kappa_{j+1}) v(\kappa_t) (\hat{x}(\kappa_t) - \hat{x}(\kappa_{t-1})) \\
&= \sum_{t=1}^m \sum_{j=t}^m (\kappa_j - \kappa_{j+1}) v(\kappa_t) (\hat{x}(\kappa_t) - \hat{x}(\kappa_{t-1})) \\
&= \sum_{t=1}^m \left(v(\kappa_t) (\hat{x}(\kappa_t) - \hat{x}(\kappa_{t-1})) \sum_{j=t}^m (\kappa_j - \kappa_{j+1}) \right) \\
&= \sum_{t=1}^m v(\kappa_t) (\hat{x}(\kappa_t) - \hat{x}(\kappa_{t-1})) (\kappa_t - \kappa_{m+1}) \\
&= \sum_{t=1}^m \kappa_t \cdot v(\kappa_t) \cdot (\hat{x}(\kappa_t) - \hat{x}(\kappa_{t-1})) \\
&= \sum_{t=1}^m R(\kappa_t) \cdot (\hat{x}(\kappa_t) - \hat{x}(\kappa_{t-1}))
\end{aligned}$$

Since players are drawn from the same distribution the expected total revenue of the badge mechanism is simply n times the expected revenue contribution of each individual player. Observe that the contribution of a badge t to the revenue depends only on the values of κ_t and κ_{t-1} . Specifically, if we denote with

$$T(q, q') = R(q) \cdot (\hat{x}(q) - \hat{x}(q')) \quad (21)$$

Then the revenue of a badge mechanism with m thresholds can be written as:

$$\mathbb{E}_v \left[\sum_i b(v_i) \right] = n \sum_{t=1}^m T(\kappa_t, \kappa_{t-1})$$

The latter implies a dynamic programming approach for computing the optimal quantile thresholds (which would then yield the optimal badge thresholds). Let $C(x, m)$ denote the optimal revenue achievable with m badges subject to the constraint that the last badge has quantile threshold x . Then we obtain the recursive equation:

$$C(x, m) = \max_{y \in [x, 1]} \{T(x, y) + C(y, m - 1)\} \quad (22)$$

with $C(x, 1)$ simply being the revenue of the badge mechanism with a single quantile threshold of x . If the quantiles were discrete, taking values among a set of T finite values, then the latter would imply a dynamic programming algorithm that runs in time $O(T^2 \cdot m)$. Hence, we can discretize the quantile space in multiples of δ , which would yield some error ϵ that is polynomial in δ and would require $O(m/\delta^2)$ time to run.

6.1 Addition of Badges and Revenue Monotonicity

Given the above characterization of revenue in a badge mechanism, a natural question is whether an extra badge can actually hurt revenue. We show that this cannot happen as long as the quantile threshold of this extra badge is above the monopoly quantile of the distribution.

Lemma 4. *Adding an extra badge with quantile threshold $\tilde{\kappa} \geq \kappa^*$ can only increase the revenue of an existing badge mechanism.*

Proof. The revenue of badge mechanism with quantile thresholds $\kappa = (\kappa_1, \dots, \kappa_m)$ can be rewritten as:

$$REV_\kappa = \sum_{t=1}^m (R(\kappa_t) - R(\kappa_{t+1})) \cdot \hat{x}(\kappa_t) \quad (23)$$

where we assumed that $\kappa_{m+1} = 0$ and hence $R(\kappa_{m+1}) = 0$.

Consider augmenting the badge mechanism with an extra badge with quantile threshold $\tilde{\kappa}$ and such that $\tilde{\kappa} \in [\kappa_{t+1}, \kappa_t]$ for some $t \in [1, m]$, with the convention that $\kappa_{m+1} = 0$ and $\kappa_0 = 1$. Then the revenue of this new mechanism is:

$$REV_{\kappa+\tilde{\kappa}} = \sum_{\sigma \neq t} (R(\kappa_\sigma) - R(\kappa_{\sigma+1})) \cdot \hat{x}(\kappa_\sigma) + (R(\kappa_t) - R(\tilde{\kappa})) \cdot \hat{x}(\kappa_t) + (R(\tilde{\kappa}) - R(\kappa_{t+1})) \cdot \hat{x}(\tilde{\kappa})$$

By regularity of the distribution, we know that $R(\cdot)$ is monotone in the region $[0, \kappa^*]$. Since we assumed that $\tilde{\kappa} \leq \kappa^*$ we get that $R(\tilde{\kappa}) \geq R(\kappa_{t+1})$. By monotonicity of $\hat{x}(\cdot)$ we also have: $\hat{x}(\tilde{\kappa}) \geq \hat{x}(\kappa_t)$. Combining the above we get the Lemma:

$$\begin{aligned} REV_{\kappa+\tilde{\kappa}} &\geq \sum_{\sigma \neq t} (R(\kappa_\sigma) - R(\kappa_{\sigma+1})) \cdot \hat{x}(\kappa_\sigma) + (R(\kappa_t) - R(\tilde{\kappa})) \cdot \hat{x}(\kappa_t) + (R(\tilde{\kappa}) - R(\kappa_{t+1})) \cdot \hat{x}(\kappa_t) \\ &= \sum_{\sigma \neq t} (R(\kappa_\sigma) - R(\kappa_{\sigma+1})) \cdot \hat{x}(\kappa_\sigma) + (R(\kappa_t) - R(\kappa_{t+1})) \cdot \hat{x}(\kappa_t) = REV_\kappa \end{aligned}$$

■

7 Approximating Optimal Revenue with Small Number of Badges

In this section, we show that simple badge mechanisms can achieve a constant approximation of the output generated by the optimal mechanism, under very generic assumptions on the social status function $f : [0, 1] \rightarrow \mathbb{R}^+$. We break our analysis based on the convexity of the status function. We show that when social status is a concave function of the proportion of players ranked at least as high (i.e. the marginal increase in status decreases as you beat more and more players), then a single badge can achieve a 4-approximation to the optimal mechanism. On the contrary, when status is a convex function then a $O(\log(N))$ badges are necessary and sufficient to achieve a constant approximation.

7.1 Decreasing Marginal Status Functions

Theorem 5. *If the social status function $f : [0, 1] \rightarrow \mathbb{R}^+$ is concave, then a single badge, with quantile threshold equal to half the monopoly quantile, can achieve a 4-approximation to the optimal revenue.*

Proof. Summarizing the analysis in the previous two sections we know that the optimal mechanism achieves revenue of:

$$OPT = n \left(R(\kappa^*) \cdot \hat{x}(\kappa^*) - \int_0^{\kappa^*} R(q) \hat{x}'(q) dq \right) \quad (24)$$

where $\hat{x}(q) = \sum_{\nu=0}^{n-1} f\left(\frac{\nu+1}{n}\right) \beta_{\nu, n-1}(q)$.

On the other hand a single badge mechanism at quantile κ achieves revenue of

$$APX = n \cdot R(\kappa) \cdot \hat{x}(\kappa) \quad (25)$$

We will show that if we set $\kappa = \frac{\kappa^*}{2}$ then $4 \cdot APX \geq OPT$.

First we point out that if the social status function is concave, then $\hat{x}(q)$ is also a concave function. This follows from known facts about binomial distributions and derivatives of Bernstein polynomials (see Appendix). Thus $\hat{x}(q)$ is a decreasing concave function of q . From this fact it follows that:

$$\hat{x}\left(\frac{\kappa^*}{2}\right) \geq \frac{\hat{x}(0) + \hat{x}(\kappa^*)}{2} \geq \frac{\hat{x}(0)}{2} \quad (26)$$

Additionally, by regularity of the distribution we know that $R(q)$ is a concave function and additionally, for any $q \leq \kappa^*$, $R(q)$ is increasing (since, κ^* is defined as the point where $R'(\kappa^*) = 0$). Additionally, $R(0) = 0$. From this it follows that:

$$R\left(\frac{\kappa^*}{2}\right) \geq \frac{R(\kappa^*)}{2} \quad (27)$$

Combining the two we get that the revenue of the single badge mechanism is at least:

$$APX \geq n \cdot \frac{R(\kappa^*) \cdot \hat{x}(0)}{4} \quad (28)$$

By the fact that $R(q)$ is increasing in the region $[0, \kappa^*]$ we have that $R(q) \leq R(\kappa^*)$. Since $\hat{x}(q)$ is decreasing in q (i.e. $\hat{x}'(q) \leq 0$) we have that: $R(q)\hat{x}'(q) \geq R(\kappa^*)\hat{x}'(q)$. Thus, we can also upper bound the revenue of the optimal mechanism:

$$OPT \leq n \left(R(\kappa^*) \cdot \hat{x}(\kappa^*) - R(\kappa^*) \int_0^{\kappa^*} \hat{x}'(q) dq \right) = n \cdot R(\kappa^*) \cdot \hat{x}(0) \quad (29)$$

The theorem then follows. ■

Example. To gain some intuition on the latter theorem, we examine the special case where the social status function is simply the proportion of players that a player beats: i.e.

$$f(t) = 1 - t. \quad (30)$$

In this case the interim allocation takes the simple form: $\hat{x}(q) = \frac{n-1}{n}(1 - q)$. Thus the revenue of the optimal mechanism becomes:

$$OPT = (n - 1) \left(R(\kappa^*)(1 - \kappa^*) + \int_0^{\kappa^*} R(q) dq \right) \quad (31)$$

Whilst, the revenue of a single badge mechanism is simply $APX = R(\kappa) \cdot (1 - \kappa)$ and more generally, the revenue of a badge mechanism with m badges is

$$R(\kappa_1) \cdot (1 - \kappa_1) + \sum_{t=2}^m R(\kappa_m) \cdot (\kappa_{m-1} - \kappa_m)$$

. The proof of the existence of a threshold κ that achieves a 4-approximation, has a nice graphical interpretation in this case as depicted in Figure 2. Theorem 5 in this example of a linear status function, boils down to showing that for a concave revenue curve there exist a κ such that the shaded area in the right figure is at least $1/4$ of the shaded area in the left figure. ■




Figure 2: Left figure depicts the revenue of the optimal mechanism. The dashed line corresponds to the revenue curve (in this example of the uniform distribution) and the shaded region corresponds to the optimal revenue. The shaded region in the right figure depicts the revenue of a single badge mechanism with quantile threshold κ .

7.2 Increasing Marginal Status

In this section we analyze the revenue of badge mechanisms for settings where status has an increasing marginal behavior. The prototypical such status function that we will analyze is the case where the status of a player is inversely proportional to the proportion of players ranked at least as high:

$$f(t) = \frac{1}{t} \quad (32)$$

In this case, we show that if the mechanism designer sets $\log(N) + 1$ badges in a particular way, then the expected total output produced by agents is at least half of the total output in the optimal mechanism. Our analysis makes use of the *revenue equivalence principle*: if two mechanisms induce the same interim allocation function, then their revenue is the same. We achieve our main result by constructing a set of $\log(N) + 1$ badges such that each player's status value is at least half of their status value in the optimal mechanism. Combining this property with the revenue equivalence principles yields that our badge mechanism is a 2-approximation to the optimal mechanism.

Theorem 6. *Let $f(t) = \frac{1}{t}$ and define κ^* as the quantile of the monopoly reserve η , i.e. $\phi(v(\kappa^*)) = 0$. The badge mechanism defined by setting quantile thresholds*

$$\kappa = \left(\kappa^*, \frac{\kappa^*}{2}, \dots, \frac{\kappa^*}{2^{\log n}} \right) \quad (33)$$

achieves at least half of the revenue of the optimal ranking mechanism.

Proof. We first show that, for any player, the interim allocation in this badge mechanism is at least half of the interim allocation in the optimal mechanism. Thus by revenue equivalence, this badge mechanism achieves at least half of the revenue of the optimal mechanism.

Consider any quantile q_i . From corollary 3, we know the interim allocation function of the optimal mechanism is equation 17. Rewriting the interim allocation in quantile space yields that the interim allocation for a player of quantile q_i is:

$$\hat{x}(q_i) = \begin{cases} \frac{1-(1-q_i)^n}{q_i} & \text{if } q_i \leq \kappa^* \\ 0 & \text{otherwise} \end{cases} \quad (34)$$

Now we examine the interim allocation of the badge mechanism under quantile thresholds $\kappa = (\kappa^*, \frac{\kappa^*}{2}, \dots, \frac{\kappa^*}{2^{\log n + 1}})$. Denote this interim allocation with $\hat{x}^\kappa(q_i)$. If $q_i > \kappa^*$ then $\hat{x}^\kappa(q_i) = \hat{x}(q_i) = 0$. If $q_i \in [\kappa_j, \kappa_{j+1}) = [\frac{\kappa^*}{2^{j-1}}, \frac{\kappa^*}{2^j})$ for some $j \in [m-1]$ then observe that the interim allocation under the badge mechanism is equal to the interim allocation of a player with quantile κ_j in the optimal mechanism.

$$\hat{x}^\kappa(q_i) = \hat{x}(\kappa_j) = \frac{1 - (1 - \kappa_j)^n}{\kappa_j} \quad (35)$$

Now observe that $q_i \leq \kappa_j \leq 2q_i$, which yields:

$$\hat{x}^\kappa(q_i) = \frac{1 - (1 - \kappa_j)^n}{\kappa_j} \geq \frac{1 - (1 - q_i)^n}{2 \cdot q_i} = \frac{\hat{x}(q_i)}{2} \quad (36)$$

Last consider a player i with quantile $q_i \leq \kappa_m = \frac{\kappa^*}{2^{\log(n)}} = \frac{\kappa^*}{n} \leq \frac{1}{n}$. The interim allocation of such a player under the badge mechanism is:

$$\hat{x}^\kappa(q_i) = \hat{x}(\kappa_m) \geq \hat{x}\left(\frac{1}{n}\right) = 1 - \left(1 - \frac{1}{n}\right)^n \geq \frac{1}{2} \geq \frac{1}{2} \hat{x}(q_i) \quad (37)$$

Where we used the fact that the interim allocation is non-decreasing in ability and hence non-increasing in the quantile in the first inequality and the fact that $\hat{x}(q_i) \leq 1$ in the last inequality.

Thus we showed that for any quantile q_i , the interim allocation of player i under the badge mechanism is at least half of his interim allocation under the optimal ranking mechanism.

$$REV_{badge} = \sum_i \hat{x}^\kappa(q_i) \phi(v(q_i)) \geq \frac{1}{2} \sum_i \hat{x}(q_i) \phi(v(q_i)) = \frac{1}{2} REV_{opt}$$

■

Figure 3 portrays the relation between the optimal and the badge interim allocation probabilities. Using




Figure 3: Optimal (dashed) versus badge mechanism (solid) interim allocation probabilities.

similar analysis we can also show that the transition to approximate optimality is smooth. That is if the

designer uses $m = \frac{\log n}{c} + 1$ badges then there exists a badge mechanism that achieves a 2^c approximation. This can be achieved by setting quantile thresholds of the form $(\kappa^*, \frac{\kappa^*}{t}, \frac{\kappa^*}{t^2}, \dots, \frac{\kappa^*}{t^{m-1}})$, for $t = 2^{\frac{\log(n)}{m-1}} = 2^c$.

Example. (Tightness of Analysis) We give an example portraying that for the case of $f(t) = 1/t$ the optimal mechanism with m badges, cannot achieve more than a $O(\log(n)/m)$ fraction of the optimal revenue. Therefore, showing that $O(\log(n))$ badges are required to achieve a constant factor of the optimal mechanism.

Consider the case where ability is distributed uniformly in $[0, 1]$. The revenue function is then $R(q) = q(1 - q)$, with derivative $R'(q) = 1 - 2 \cdot q$ and the monopoly quantile is $1/2$. Thereby the optimal revenue is:

$$\begin{aligned} \frac{OPT}{n} &= \int_0^{1/2} (1 - 2 \cdot q) \hat{x}(q) dq = n \int_0^{1/2} (1 - 2 \cdot q) \frac{1 - (1 - q)^n}{q} dq \\ &\geq \int_0^{1/2} (1 - 2q) \frac{1 - e^{-nq}}{q} dq \\ &= \log(n) + \frac{2(1 - e^{-n/2})}{n} + \Gamma\left(0, \frac{n}{2}\right) + \gamma - 1 - \log(2) \\ &\geq \log(n) + \gamma - 1 - \log(2) = \Theta(\log(n)) \end{aligned}$$

where $\gamma \approx .57$ is the Euler-Mascheroni constant and $\Gamma\left(0, \frac{n}{2}\right) = \int_{n/2}^{\infty} \frac{e^{-t}}{t} dt \rightarrow 0$.

On the other hand we note that the maximum social welfare achievable by any mechanism that uses m badges is at most $n \cdot m$. For any bid profile b , the social welfare from any badge mechanism with m badges is simply:

$$\sum_{t=1}^m \frac{|i : r_i(b) = t|}{|i : r_i(b) \geq t|/n} \leq \sum_{t=1}^m n = n \cdot m \quad (38)$$

Therefore, it trivially follows that the optimal revenue achievable with m badges is at most $n \cdot m$. Thus as $n \rightarrow \infty$ the fraction of the optimal revenue achievable with m badges converges to $\log(n)/m$. ■

The above theorem can also be generalized for arbitrary social status functions, as follows.

Theorem 7. *The optimal badge mechanism with $\log\left(\frac{f(1/n)}{f(1)}\right)$ badges, achieves at least half of the optimal ranking mechanism.*

Another status function that we can apply the above theorem is the following:

$$f(t) = \frac{1}{n \cdot t} + \frac{1}{n \cdot t + 1} + \dots + \frac{1}{n} = H_n - H_{n \cdot t - 1} \quad (39)$$

The intuition behind the above social status function is the following: each social rank can be thought of as unlocking privileges or features of a web service, or in more general contest settings as unlocking some resource. Assume that each resource can be thought as a unit of value that is shared among the players to whom the resource is unlocked. This exact intuition is portrayed in the above function.

Observe that this function satisfies the property that $\frac{f(1/n)}{f(1)} = n \cdot H_n$. Thus applying Theorem 7 we get that $\log(n) + \log \log(n) = O(\log(n))$ badges are sufficient to get a constant approximation to the optimal revenue.

8 Discussion

We've studied a model where agents seek badges on a website as a means of achieving social status within a community. Designers can use these badges as a way to incentivize user contributions to their site. We have found that mechanisms with a small number of badges can provide good approximations but nevertheless find that mechanisms that employ an extremely large number of badges or implement ranking of all agents are optimal within these settings. One major open question is why should designers use small badge systems when ranking systems or systems with large number of badges are optimal? Although one could justify this decision from several view points (perhaps it is costly to implement new badges or it is simply desirable to have a low system complexity), we are currently investigating game-theoretic reasons for why having a small number of badges can be a good property for an incentive system.

References

- [1] Ashton Anderson, Daniel Huttenlocher, Jon Kleinberg, and Jure Leskovec. Steering user behavior with badges. In *Proceedings of the 22nd international conference on World Wide Web*, pages 95–106. International World Wide Web Conferences Steering Committee, 2013.
- [2] Jeremy Bulow and John Roberts. The simple economics of optimal auctions. *The Journal of Political Economy*, pages 1060–1090, 1989.
- [3] Ruggiero Cavallo and Shaili Jain. Efficient crowdsourcing contests. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 677–686. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [4] Shuchi Chawla, Jason D Hartline, and Balasubramanian Sivan. Optimal crowdsourcing contests. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 856–868. SIAM, 2012.
- [5] Pradeep Dubey and John Geanakoplos. Grading exams: 100, 99, 98, or a, b, c? *Games and Economic Behavior*, 69(1):72–94, 2010.
- [6] David Easley and Arpita Ghosh. Incentives, gamification, and game theory: An economic approach to badge design. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 359–376. ACM, 2013.
- [7] Arpita Ghosh and Patrick Hummel. Implementing optimal outcomes in social computing: a game-theoretic approach. In *Proceedings of the 21st international conference on World Wide Web*, pages 539–548. ACM, 2012.
- [8] Ed Hopkins and Tatiana Kornienko. Running to keep in the same place: Consumer choice as a game of status. *American Economic Review*, pages 1085–1107, 2004.
- [9] Shaili Jain, Yiling Chen, and David C Parkes. Designing incentives for online question-and-answer forums. *Games and Economic Behavior*, 2012.
- [10] Benny Moldovanu, Aner Sela, and Xianwen Shi. Contests for status. *Journal of Political Economy*, 115(2):338–363, 2007.