Winner-Take-All Crowdsourcing Contests with Stochastic Production

Ruggiero Cavallo
Yahoo! Research
111 West 40th Street
New York, NY 10018
cavallo@yahoo-inc.com

Shaili Jain
Computer Science Department
Yale University
New Haven, CT 06520
shaili.jain@yale.edu

May 29, 2012

Abstract

We study winner-take-all crowdsourcing contests in a model of costly effort and stochastic production introduced by Cavallo and Jain [5]. In this contest format, the principal strategically selects a prize value P, agents simultaneously and strategically exert costly effort towards production yielding *stochastic* quality results, and the agent who produces the highest quality good is paid P by the principal. We derive some general characteristics of the Nash equilibria of such contests, and then give a characterization of pure strategy equilibria when qualities are distributed uniformly or exponentially as a function of effort. Given the equilibrium characterization, we evaluate expected efficiency. We find that the winner-take-all contest paradigm will achieve the efficient outcome in equilibrium for a large range of instantiations of the principal's value-for-quality, though for other instantiations it cannot.

1 Introduction

Crowdsourcing is an increasingly popular model of procurement in today's online marketplaces. In this model, a principal seeks completion of some task, posts an open call for submissions, and allows multiple agents to simultaneously submit solutions, awarding a prize to the participant with the best solution. The number and size of online crowdsourcing marketplaces has grown markedly in recent years, with some notable examples including Taskcn, Topcoder, 99designs and CrowdFlower. Crowdsourcing tasks resemble contests in that many agents simultaneously exert effort in order to win a prize, where the winner is determined based on relative performance. The process of selecting a winner based on submission quality and awarding her a lump sum prize constitutes a mechanism class, which we term *winner-take-all*.

In this paper, we initiate a study of winner-take-all contest mechanisms in the model of costly effort and stochastic production introduced in [5]. We are motivated by the following questions: What level of effort towards production is induced by winner-take-all contest schemes? How does this depend on the magnitude of the prize on offer? If the principal sets the prize to maximize his value for the produced good minus the prize-price he must pay, what prize will he choose? If the principal and agents have quasilinear (and thus interpersonally-comparable) utility, what level of social welfare (efficiency) results?

Since we take agents to be self-interested and strategic, the answer to these questions requires a game-theoretic approach—we will analyze what happens *in equilibrium*. We start by deriving some general properties of equilibria in a stochastic model of crowdsourcing, and then fully work out the implications in two specific models of stochastic production: one where quality is uniformly distributed, but increasing as a function of effort, and one where quality is exponentially distributed, again increasing with effort. We focus on pure strategy equilibria, and find that under a natural restriction on the quality distribution (encompassing the uniform and exponential cases as well as many others), each agent always has a unique best response to the strategies of others, which simplifies our analysis. This allows us to identify a unique pure-strategy equilibrium for the exponential distribution and uniform distribution, with two agents. We compare the efficiency (social welfare) of this equilibrium to that in the social-welfare maximizing outcome. This is a particularly relevant comparison in light of [5], which shows how—for arbitrary quality distributions—the efficient outcome can be achieved in equilibrium in a *non*-winner-take-all scheme. Thus this efficiency gap can be considered the price one pays by adopting the simple winner-take-all format as opposed to the more complex, but efficient, mechanism of [5]. We find that while there are instances of the problem for which winner-take-all contests are inefficient, they are perfectly efficient in Nash equilibrium for a range of cases.

2 Related Work

There has recently emerged a new line of work providing a theory of crowdsourcing contests [7, 2, 6, 5]. Most of this work focuses on the case where agents have private skill information and choose a private effort level. DiPalantino and Vojnovic [7] make the connection to all-pay auctions and model a market with multiple contests, considering the principal's optimization problem in the limit-case as the number of agents and contests goes to infinity. Archak and Sundararajan [2] and Chawla, Hartline, and Sivan [6] focus on the design of a single contest; they seek to determine how many prizes should be awarded, and of what value. Chawla et al. [6] make the connection between crowdsourcing contests and optimal auction design, finding that the optimal crowdsourcing contest—from the perspective of the principal seeking quality submissions—is a virtual valuation maximizer.

While the current paper follows the line of work of Cavallo and Jain [5], who initiate a study of stochastic production models for crowdsourcing and focus on optimizing for efficiency, most of these other previous works consider a *deterministic* model of quality as a function of effort and skill. In a deterministic model of production, if the principal's value is proportional to the maximum quality good produced, the crowdsourcing paradigm is not well-motivated from an efficiency perspective. Perhaps naturally, then, outside of [5] a principal-centric viewpoint has most frequently been adopted.

Much of the previous design-motivated work has focused on how to maximize submission quality, given a prize budget, whether it be the highest quality submission [18, 6] or the total sum of submission qualities [17, 18, 16]. Still other work has focused on maximizing the sum of the top k submissions minus the prize awarded [2]. Ghosh and Hummel [9] consider a more general class of utility functions to optimize for in a setting with virtual "points" (a type of currency); namely, they show that there is a *best contribution mechanism* that can implement the principal-optimal outcome (in their model, the principal does not experience disutility for the prize awarded). Another line of work in the economics literature uses contests to extract effort under a hidden action [14, 10, 20]. Similar to our work, the output is a stochastic function of the unobservable effort, but the setting is different in that the principal obtains value from the cumulative effort of the agents, rather than just the maximum result.

These works are related to a long literature in economics studying all-pay auctions. Prior work has focused on providing equilibrium characterizations for a complete information all-pay auction [23, 19, 3], complete information all-pay auctions with reserve prices [4], and incomplete information all-pay auctions [24, 11, 13, 1]. Still other works study sequential all-pay auctions [12, 21, 15]. There has also been work on multi-stage research tournaments that award a single prize [22, 8].

Like the current paper, most of this previous related work examines outcomes with strategic agents and thus consists largely of equilibrium analysis. However, a critical differentiator of the current paper is the model: we adopt the model of [5], where quality is a *stochastic* function of effort and skill. From an efficiency perspective, this stochasticity (combined with a deadline under which procurement is required) is the most natural way to motivate the crowdsourcing paradigm—the redundant costly-effort of simultaneous production is justified by the principal's value for high quality and inability to ensure receipt before the deadline, if production is instead ordered sequentially. Cavallo and Jain [5] introduce this model and design efficient mechanisms for the problem of crowdsourcing. Here we analyze winner-take-all mechanisms, the prevailing crowdsourcing payment scheme seen in practice, under this stochastic model of production. Our goal in doing so is to determine the effectiveness of winner-take-all payment schemes in implementing the efficient outcome. We aim to quantify the efficiency gap, or the tradeoff between implementing the more complicated yet efficient mechanisms from [5], versus the simpler yet inefficient winner-take-all mechanisms. To the best of our knowledge, no previous work has addressed the efficiency of winner-take-all contests or studied them in a model of stochastic production.

3 Model

In the crowdsourcing paradigm, multiple units of a good are simultaneously produced and submitted to a principal. There is a set of agents $I = \{1, \dots, n\}$ capable of producing goods, where each $i \in I$ makes a *privately observed* choice of effort δ_i to expend on production. We identify the effort level δ_i with the dollar value in costs ascribed to it by agent i. We assume that $\delta_i \in [0, \delta_{max}]$, $\forall i \in I$, for some $\delta_{max} \in \Re^+$. Agents are homogenous aside from their private effort choice.[1] If an agent attempts production with effort δ_i, a good is produced with quality that is a priori uncertain but is a function of δ_i.

Quality is identified with value to the principal in dollar-terms, and can be thought of as the output of a non-deterministic function mapping effort to $\Re^+$. The probability distribution over *relative* quality, given any effort level, is publicly known; and the principal has value $v \in \Re^+$, a scale factor that maps relative quality levels to absolute quality (dollar-value).[2] The principal obtains value commensurate with the maximum quality good produced. For instance, if agents i and j expend effort δ_i and δ_j and produce goods with relative qualities q_i and q_j, respectively, the principal's utility from obtaining the goods equals the maximum *asbolute* quality produced, which is $\max\{vq_i, vq_j\}$.

For any non-zero effort level $\delta_i \in (0, \delta_{max}]$, we denote the p.d.f. and c.d.f. over resulting relative quality as f_{δ_i} and F_{δ_i}, respectively, where the support of f_{δ_i} is in $\Re^+$. We assume symmetry across bidders in the sense that the private effort choice is the only differentiating factor; i.e., for two agents making the same choice of effort level, the distribution over quality they will produce is the same (though there is no presumed correlation so the resulting quality may differ). We assume the probability density over quality, evaluated at any particular quality level, is differentiable with respect to effort δ_i, for all $\delta_i \in (0, \delta_{max}]$. We assume that effort $\delta_i = 0$ deterministically yields quality 0 (no effort yields no production), and then for notational convenience we let $F_0(x) = 1$, $\forall x \geq 0$. We also assume that $\forall x > 0$, $F_{\delta_i}(x)$ converges to 1 as δ_i goes to 0.[3] Finally, we make the natural assumption that more effort has first-order stochastic dominance over less effort with respect to quality, i.e.:

$$\forall 0 \leq \delta_i < \delta_i' \leq \delta_{max},\ \forall x \in [0, v],\ F_{\delta_i}(x) \geq F_{\delta_i'}(x) \tag{1}$$

In this paper we analyze a particular family of mechanisms, winner-take-all mechanisms, where the principal posts a fixed prize P (in dollars), allows the agents to compete for the prize, and awards the lump sum prize P to the agent who submits the maximum quality good. We seek to understand the effectiveness of winner-take-all mechanisms in implementing the social welfare maximizing outcome. We adopt a quasilinear utility model and assume all players are risk-neutral. Given our identification of the quality of the good (times the principal's scale-factor value v) with the dollar value ascribed to it by the principal, a rational principal will seek to maximize his value for the maximum quality good minus the prize he pays. Likewise, given our identification of effort level δ_i with the dollar value in costs ascribed to it by agent i, a rational agent will seek to maximize his expected prize value received minus the effort he expends towards production.

The winner-take-all paradigm defines a game wherein the principal chooses a prize value P and each agent $i \in I$ chooses an effort level δ_i. We will consider equilibrium outcomes, and pay particular attention to the *efficiency* they yield. Letting $Q_i(v, \delta_i)$ be a random variable

[1] In our previous paper [5] we considered the richer model where agents were also differentiated by private skill information (see also [6]).

[2] Note that in [5], the principal's value v was *private*. Here, because the principal is essentially setting the mechanism by choosing a prize value, whether v is private or public is of no consequence to the equilibrium outcomes that result.

[3] This entails that $\forall \delta_{-i} \in [0, \delta_{max}]^{n-1}$ such that $\max_{j \in I \setminus \{i\}} \delta_j > 0$, $\int_0^\infty f_{\delta_i}(x) \prod_{j \in I \setminus \{i\}} F_{\delta_j}(x) dx$ converges to 0 as δ_i goes to 0. In words: as long as some agent other than i is exerting non-zero effort, i's probability of winning the contest (and, with it, his expected utility) goes to 0 as his effort goes to 0.

representing the absolute quality level produced by $i \in I$ who expends effort δ_i, the expected efficiency of an equilibrium in which each $i \in I$ exerts effort δ_i is:

$$\mathbb{E}[\max_{i\in I} Q_i(v,\delta_i)] - \sum_{i\in I} \delta_i \tag{2}$$

(The prize value P does not factor into social welfare since it is subtracted from the principle's utility and added to the winning agent's.) An *efficient* effort vector is a vector of effort levels that maximizes Eq. (2) given v.[4] We will ultimately compare the expected social welfare of the equilibrium outcomes in winner-take-all contests to that of the efficient outcome.

Our results regarding equilibrium analysis for the agents will be parameterized by the prize value P—we describe equilibrium agent effort levels for the principal's choice of prize value. We start with some distribution-independent facts about equilibria (Section 4), and then move on to detailed analysis of special cases—*exponentially distributed quality*, where $f_{\delta_i}(x) = \frac{1}{\delta_i} e^{-\frac{x}{\delta_i}}$, $\forall x \in \Re^+$ (Section 5), and *uniformly distributed quality*, where $f_{\delta_i}(x) = \frac{1}{\delta_i}$, $\forall x \in [0, \delta_i]$ and $f_{\delta_i}(x) = 0$ otherwise (Section 6). This equilibrium analysis will then allow us to derive the principal's equilibrium prize-setting behavior (in Sections 5.2 and 6.3, for the exponential distribution and uniform distribution, respectively), which pins down a complete equilibrium characterization and allows us to evaluate equilibrium efficiency (in Sections 5.3 and 6.4).

4 General best-response functions

Taking the principal's choice of top-submission prize value P as given (for now), we determine the best-response functions for an agent i in terms of her private effort choice δ_i.

Remark 4.1. *The expected utility for agent i $\in I$ who exerts effort $\delta_i > 0$, when other agents exert efforts δ_{-i}, is given by:* $P \int_0^\infty f_{\delta_i}(x) \prod_{j\in I\setminus\{i\}} F_{\delta_j}(x)dx - \delta_i$.

Considering the first order condition for the best-response function, we can see that any critical point satisfies:

$$\frac{\partial}{\partial \delta_i}\Big(\int_0^\infty f_{\delta_i}(x) \prod_{j\in I\setminus\{i\}} F_{\delta_j}(x)dx\Big) - \frac{1}{P} = 0 \tag{3}$$

We identify a general condition below in the following remark that allows us to identify all potential maxima of the best response function. We use this remark in the analysis of exponentially distributed quality and uniformly distributed quality.

Remark 4.2. *Fixing any δ_{-i}, if $\frac{\partial}{\partial \delta_i}[\int_0^\infty f_{\delta_i}(x) \prod_{j\in I\setminus\{i\}} F_{\delta_j}(x)dx] - \frac{1}{P}$ is monotonically decreasing in δ_i over the interval $(0, \delta_{max}]$ and Eq. (3) has a zero $\hat{\delta}_i \in (0, \delta_{max})$, then the best-response effort choice for i is $\hat{\delta}_i$; if Eq. (3) has no zero in $(0, \delta_{max})$ then the best-response is either non-participation or maximum-effort participation.*

[4] That is, effort vector δ^* is efficient if and only if $\delta^* \in \arg\max_\delta \big(\mathbb{E}[\max_{i\in I} Q_i(v,\delta_i)] - \sum_{i\in I} \delta_i\big)$.

If $\frac{\partial}{\partial \delta_i}[\int_0^\infty f_{\delta_i}(x) \prod_{j \in I \setminus \{i\}} F_{\delta_j}(x) dx]$ is *strictly* monotonically decreasing in δ_i, there is at most one value of δ_i that satisfies the above equation and if the zero exists, the value of the derivative switches from positive to negative, so we have a maximum. For the case of exponentially distributed quality, we show that in fact this expression is strictly monotonically decreasing.

An equilibrium is a profile of strategies ($\delta_1, \ldots, \delta_n$; in our case, fixing P) wherein every agent is best-responding. Thus these basic properties will help point us to the equilibria of winner-take-all contests with prize P. But before going on, without assuming anything more about the environment we note that no equilibrium will ever involve cumulative agent effort exceeding the prize value, and also no equilibrium will involve only a single agent exerting non-zero effort. These facts will be useful in deriving the equilibrium characterizations to come in the following sections.

Proposition 4.3. *For arbitrary quality distributions, for arbitrary values of $n \geq 2$ and $P \in \Re^+$, any equilibrium strategy profile has agents collectively exerting at most P units of effort.*

Proof. Consider arbitrary equilibrium effort profile $\delta_1, \ldots, \delta_n$. Each agent's utility is her expected prize reward (which we'll denote R_i) minus any effort exerted, i.e., $u_i = \mathbb{E}[R_i(\delta_1, \ldots, \delta_n)] - \delta_i$. In equilibrium each agent must have an expected utility of at least 0 since an agent can exert 0 effort and receive an expected reward of 0. Therefore, $\sum_{i=1}^n u_i \geq 0$, or $\sum_{i=1}^n (\mathbb{E}[R_i(\delta_1, \ldots, \delta_n)] - \delta_i) \geq 0$. We observe that $\sum_{i=1}^n \mathbb{E}[R_i(\delta_1, \ldots, \delta_n)] = P$, so therefore, $\sum_{i=1}^n \delta_i \leq P$. □

Proposition 4.4. *For arbitrary quality distributions and arbitrary values of $n \geq 2$, $P \in \Re^+$, and $\delta_{max} > 0$, no pure-strategy equilibrium has exactly one agent exerting non-zero effort.*

Proof. Consider a candidate equilibrium profile with one agent exerting x units of effort, where $0 < x \leq \delta_{max}$. This agent's expected utility is $P - x$. If an agent reduces his effort to $\frac{x}{2}$, his expected payoff increases to $P - \frac{x}{2}$, a profitable deviation. □

5 Exponentially distributed quality

We now derive an equilibrium characterization of the case where output depends exponentially on the choice of effort. Specifically, we define the mean of the quality distribution, when $\delta_i \in (0, \delta_{max}]$, to be $\frac{1}{\lambda} = \delta_i$, so $f_{\delta_i}(x) = \frac{1}{\delta_i} e^{-\frac{x}{\delta_i}}$ and $F_{\delta_i}(x) = 1 - e^{-\frac{x}{\delta_i}}$. Our results for this section are limited to the two-agent case. The proofs of all Lemmas are deferred to the Appendix.

5.1 Equilibrium analysis with $n = 2$ agents (exponential)

We first use the structure provided in Remark 4.2 to show that in the case of exponentially distributed quality, the best-response function has at most one critical point.

Lemma 5.1. *In the exponentially distributed quality case with $n = 2$, $\forall \delta_2 \in [0, \delta_{max}]$, the function $\frac{\partial}{\partial \delta_1} \int_0^\infty [f_{\delta_1}(x) F_{\delta_2}(x) dx] - \frac{1}{P}$ has at most one zero with respect to δ_1 over the interval $(0, \delta_{max}]$, for any value of δ_2.*

From this lemma, we can conclude that in the exponentially distributed quality case with two agents, each agent always has a unique best-response, for any choice of effort level by the other agent. We now see that if the prize value P is large enough, both agents participate with maximal effort.

Proposition 5.2. *In the exponentially distributed quality case with $n = 2$, when $P \geq 4\delta_{max}$, both agents will exert effort δ_{max} in equilibrium and this is the only equilibrium.*

Proof. If both agents exert effort δ_{max}, then they both have an expected payoff of $\frac{P}{2} - \delta_{max}$, which is strictly greater than 0. Consider a deviation for an agent where she decides to exert less effort, $\delta_1 < \delta_{max}$; her expected utility becomes $P\int_0^\infty f_{\delta_1}(x)F_{\delta_2}(x)dx - \delta_1$. The derivative of this with respect to δ_i equals $P\frac{\delta_{max}}{(\delta_1+\delta_{max})^2} - 1 > P\frac{\delta_{max}}{(2\delta_{max})^2} - 1 = \frac{P-4\delta_{max}}{4\delta_{max}}$. So when $P \geq 4\delta_{max}$, the derivate is positive everywhere and thus δ_{max} is the best-response.

Now suppose that agent 2 exerts effort $\delta_2 < \delta_{max}$, we consider the best-response to this. Consider agent 1 exerting effort δ_1. We claim that if agent 1 is best-responding, agent 1 should exert effort strictly greater than δ_2. The best-response function has a critical point at $\delta_1 = \sqrt{\delta_2}(\sqrt{P} - \sqrt{\delta_2})$ (applying Remark 4.2). Since $\delta_2 < \delta_{max}$ and $P \geq 4\delta_{max}$, then $\sqrt{P} - \sqrt{\delta_2} > \sqrt{\delta_2}$, so the derivative has a zero at $\delta_1 > \delta_2$. Note that if this zero occurs at a value of $\delta_1 > \delta_{max}$, the derivative is positive over the interval $\delta_1 \in [\delta_2, \delta_{max}]$. Therefore, if a player exerts $\delta_2 < \delta_{max}$ effort, the best-response to this is exerting strictly greater effort. Therefore, we know for any strategy profile (x, y), where $x < y$, the agent exerting lesser effort is not best-responding. Similarly, any strategy profile (x, x), where $x < \delta_{max}$ has both players not best responding, which gives us the desired result. □

We now see, through a set of three lemmas leading to a proposition, that if $P \leq 4\delta_{max}$ the sole pure strategy equilibrium is interior.

Lemma 5.3. *In the exponentially distributed quality case with $n = 2$, if agent 2 exerts δ_2 units of effort, where $0 < \delta_2 < \frac{P}{4}$, then agent 1's unique best-response δ_1 is in the interval $(\delta_2, \frac{P}{4})$.*

Lemma 5.4. *In the exponentially distributed quality case with $n = 2$, if agent 2 exerts $\frac{P}{4}$ units of effort, agent 1's unique best-response is $\delta_1 = \frac{P}{4}$.*

Lemma 5.5. *In the exponentially distributed quality case with $n = 2$, if agent 2 exerts δ_2 units of effort, where $\delta_2 > \frac{P}{4}$, then agent 1's unique best response δ_1 is in the interval $[0, \frac{P}{4})$.*

Proposition 5.6. *In the exponentially distributed quality case with $n = 2$, when $P < 4\delta_{max}$, the only pure-strategy Nash equilibrium for the winner-take-all contest has each player exert effort $\frac{P}{4}$.*

Proof. From Proposition 4.4, we know that any equilibrium strategy profile must have both players exerting non-zero effort. Suppose there exists an equilibrium profile with player 1 exerting x units of effort and player 2 exerting y units of effort, with $x, y > 0$. Without loss of generality, we can assume that $x \leq y$. First consider the case that $0 < x < \frac{P}{4}$. If $y < \frac{P}{4}$, Lemma 5.3 tells us that player 1's best-response to y is to exert effort strictly greater than y, so this cannot be part of a Nash equilibrium. If $y = \frac{P}{4}$, then player 2 is not responding to player 1 according to Lemma 5.3. If $y > \frac{P}{4}$, Lemma 5.3 tells us player 2 is not best-responding. Now consider the case that $x > \frac{P}{4}$ (which implies $y > \frac{P}{4}$). Lemma 5.5 tells us that neither player 1 nor player 2 is best-responding. Now consider the case that $x = \frac{P}{4}$. If $y > \frac{P}{4}$, Lemma 5.5 tells us that player 1 is not best-responding. If $y = \frac{P}{4}$, we know from Lemma 5.4 that both players are best-responding to each other. Thus $(\frac{P}{4}, \frac{P}{4})$ is the one pure strategy Nash equilibrium. □

Propositions 5.2 and 5.6 give us a complete characterization of the pure-strategy Nash equilibrium in winner-take-all contests with two agents where quality is exponentially distributed, as a function of effort:

Theorem 5.7. *In the exponentially distributed quality case with $n = 2$: in the winner-take-all contest with $P < 4\delta_{max}$, the only pure-strategy Nash equilibrium has both agents exert effort $\frac{P}{4}$; in the winner-take-all contest with $P \geq 4\delta_{max}$, the only pure-strategy Nash equilibrium has both agents exert effort δ_{max}.*

One major motivator for analyzing the equilibria of winner-take-all contests is to quantify the inefficiency of using such a scheme compared to an alternate scheme that always achieves efficient effort levels in equilibrium. The equilibrium characterizations above provide most of the picture, but we still need to consider the principal's strategic choice of prize value P, given her value v. This will give us a complete picture of the equilibrium behavior, allowing us—in Section 5.3—to compare the efficiency of a winner-take-all contest with the optimal solution, which is achieved by the non-winner-take-all mechanism of [5].

5.2 Equilibrium prize value with $n = 2$ agents (exponential)

Given the equilibrium characterization above, we can now examine the principal's utility-maximizing choice of P as a function of her value v.

Proposition 5.8. *In the exponentially distributed quality case with $n = 2$, in equilibrium:*

- *When $v > \frac{8}{3}$, the principal chooses $P = 4\delta_{max}$, with each agent exerting effort δ_{max}.*
- *When $v < \frac{8}{3}$, the principal chooses $P = 0$.*
- *When $v = \frac{8}{3}$, the principal is indifferent between all prizes in the interval $P \in [0, 4\delta_{max}]$.*

Proof. We note that the principal need never offer a prize of greater than $4\delta_{max}$ since a prize of $4\delta_{max}$ induces an equilibrium where both players exert effort δ_{max}. If the prize is P where $0 \leq P \leq 4\delta_{max}$, each agent will exert effort $\frac{P}{4}$ and the expected value to the principal of the highest quality good will be $\frac{vP}{4}\frac{3}{2}$. The principal's expected utility of setting a prize of P is then $\frac{vP}{4}\frac{3}{2} - P = P(\frac{v}{4}\frac{3}{2} - 1)$. This is positive if and only if $v > \frac{8}{3}$. When $v > \frac{8}{3}$, the principal's utility is maximized by making P as large as possible without exceeding $4\delta_{max}$, i.e., $4\delta_{max}$. When $v = \frac{8}{3}$, the principal's utility is 0 for any prize value $P \in [0, 4\delta_{max}]$. When $v < \frac{8}{3}$, the principal's utility is negative for all $P > 0$, so he does not order production. □

5.3 Efficiency of winner-take-all contests (exponential)

Given the equilibrium characterizations of the agents and the principal, we compare the outcome of the game under the winner-take-all mechanism as compared to the efficient outcome. First, we need to determine what the efficient outcome for the exponential distribution looks like. We show that the exponential distribution satisfies the extreme effort condition from [5], which means that the efficient outcome has all agents either exert zero effort or full effort. This will allow us to clearly compare the equilibrium efficiency of a winner-take-all contest with the efficiency of the optimal effort policy.

Lemma 5.9. *For the exponentially distributed quality case, there is an efficient policy in which each agent i exerts either no effort $\delta_i = 0$ or full effort $\delta_i = 1$.*

Given that the exponential distribution satisfies the extreme effort condition, determining the efficient solution is equivalent to simply determining the number of agents to exert full effort in order to maximize the social welfare.

In what follows, let $v(i, i+1)$ denote the value of v at which the social welfare of having i agents exert full effort is the same as the social welfare of having $i+1$ agents exert full effort, under the exponential distribution.

Lemma 5.10. *For the exponentially distributed quality case, $\forall i \in \{1, \ldots, n\}$, $v(i-1, i) = i$. Moreover, when $v > v(i-1, i)$, the social welfare of i agents exerting full effort is strictly greater than the social welfare of $i-1$ agents exerting full effort and when $v < v(i-1, i)$, the social welfare of $i-1$ agents exerting full effort is strictly greater than the social welfare of i agents exerting full effort.*

Proof. When i agents exert full effort, the expected value to the principal is $v\delta_{max}H_i$. Therefore, the social welfare of having i agents exert full effort is: $v\delta_{max}H_i - i\delta_{max}$. To determine $v(i-1, i)$, we need $v(i-1, i)$ to satisfy $v(i-1, i)\delta_{max}H_i - i\delta_{max} = v(i-1, i)\delta_{max}H_{i-1} - (i-1)\delta_{max}$ or in other words, $v(i-1, i)(H_i - H_{i-1}) = 1$ or more simply, $v(i-1, i) = i$. When $v > v(i-1, i)$, $v\delta_{max}H_i - i\delta_{max} > v\delta_{max}H_{i-1} - (i-1)\delta_{max}$ and when $v < v(i-1, i)$, $v(i-1, i)\delta_{max}H_i - i\delta_{max} < v\delta_{max}H_{i-1} - (i-1)\delta_{max}$. □

Corollary 5.11. *For the exponentially distributed quality case, $\forall i \in \{1, \ldots, n-1\}$, $v(i-1, i) < v(i, i+1)$.*

Theorem 5.12. *In the exponentially distributed quality case, if $\exists i \in \{0, ..., n-1\}$ such that $i < v < i+1$, the efficient solution involves i agents exerting full effort and $n-i$ agents exerting 0 effort. Alternatively, if $v > n$, the efficient solution involves n agents exerting full effort. Finally, if $\exists i \in \{1, ..., n\}$ such that $v = i$, there are two classes of efficient solutions: one where i agents exert full effort and one where $i-1$ agents exert full effort (with the other $n-i$ or $n-i+1$ agents exerting 0 effort).*

Proof. Since the exponential distribution satisfies the extreme effort condition, it suffices to consider solutions where agents either exert full effort or no effort. We establish this theorem via induction. We claim that for any value of $v > v(i-1, i)$, having i agents exert full effort leads to strictly higher social welfare than having $0 \leq j < i$ agents exert full effort (if any such j exists). When $i = 0$, there is no value of j, such that $0 \leq j < i$, so the claim trivially holds for all $v \geq 0$. For arbitrary integer $i > 0$, assume that we have proven the claim for $i - 1$, which means that for $v > v(i-2, i-1)$, having $i - 1$ agents exert full effort leads to higher social welfare than having $0 \leq j < i - 1$ agents exert full effort. We know that there exists a unique point $v(i-1, i)$ at which the social welfare of having $i - 1$ agents exert full effort is the same as the social welfare of having i agents exert full effort. Moreover, we know that for all $v > v(i-1, i)$, the social welfare of having i agents exert full effort is strictly greater than the social welfare of having $i - 1$ agents exert full effort. Thus having i agents exert full effort yields greater social welfare than having $0 \leq j < i$ agents exert full effort for any value $v > v(i-1, i)$. The final step of the inductive argument gives us that having n agents exert full effort leads to greater social welfare than having $0 \leq j < n$ agents exert full effort for any value $v > v(n-1, n)$.

We can establish a similar argument via induction that at any value of $v < v(i, i+1)$, having i agents exert full effort leads to strictly higher social welfare than having $n \geq j > i$ agents exert full effort (if any such j exists). When $i = n$, there is no value of j, such that $n \geq j > i$, so the claim trivially holds for all $v \geq v(n-1, n)$. For arbitrary integer $i < n$, assume that we have proven the claim for $i+1$, which means that for $v < v(i+1, i+2)$, having $i + 1$ agents exert full effort leads to higher social welfare than having $n \geq j > i + 1$ agents exert full effort. We know that for all $v < v(i, i+1)$, having i agents exert full effort leads to greater social welfare than having $i + 1$ agents exert full effort. Thus for all $v < v(i, i+1)$, having i agents exert full effort yields greater social welfare than having $n \geq j > i$ agents exert full effort. Combining these two inductive arguments, with Corollary 5.11, gives us that having i agents exert full effort when $v(i-1, i) < v < v(i, i+1)$ maximizes social welfare for all $1 \leq i \leq n - 1$. □

Proposition 5.8 and Theorem 5.12 allow us to conclude that the winner-take-all contest can be inefficient for certain values of v.

Theorem 5.13. *For the exponentially distributed quality case with* $n = 2$*, the winner-take-all contest is inefficient when* $1 < v < \frac{8}{3}$*.*

We see from Theorem 5.12, that when $v \leq 1$, the efficient outcome is for 0 agents to exert effort, which is implemented by the winner-take-all contest. However, in the case that $1 < v \leq 2$, the efficient outcome is for 1 agent to exert full effort, but the equilibrium behavior of a winner-take-all contest leads to 0 agents exerting effort. Finally, when $2 \leq v < \frac{8}{3}$, the efficient outcome is for 2 agents to exert full effort, but the equilibrium behavior of a winner-take-all contest leads to 0 agents exerting effort. Finally, when $v \geq \frac{8}{3}$, the equilibrium behavior of a winner-take-all contest leads to 2 agents exerting full effort, which is the efficient outcome according to Theorem 5.12. This is perhaps surprising, in that, for the exponential distribution with $n = 2$, the winner-take-all contest is either perfectly efficient or perfectly inefficient.

6 Uniformly distributed quality

In this section, we provide an analysis analogous to the above for the case of uniformly distributed quality. More specifically, quality is uniformly distributed between 0 and δ_i, so this distribution has a mean of $\frac{\delta_i}{2}$. For $\delta_i \in (0, \delta_{max}]$, this gives $f_{\delta_i}(x) = \frac{1}{\delta_i}$ when $0 \le x \le \delta_i$ and 0 otherwise, and $F_{\delta_i}(x) = \frac{x}{\delta_i}$ when $0 \le x \le \delta_i$ and 1 when $x > \delta_i$. We start with the $n = 2$ case and provide a equilibrium characterization, which we generalize for larger n. Spelling out the $n = 2$ case is useful since the technical lemmas are used in the general n case (e.g. Lemma 6.11).

6.1 Equilibrium analysis with $n = 2$ agents (uniform)

We start by interpreting the condition from Remark 4.2 for the uniform distribution, which means that for all but one special case, the best-response function has at most one critical point.

Lemma 6.1. *In the uniformly distributed quality case with* $n = 2$, $\forall \delta_2 \in [0, \delta_{max}]$, *the function* $\frac{\partial}{\partial \delta_1} \int_0^\infty [f_{\delta_1}(x) F_{\delta_2}(x) dx] - \frac{1}{P}$ *has at most one zero with respect to* δ_1 *over the interval* $(0, \delta_{max}]$, *for any value of* $\delta_2 \neq \frac{P}{2}$.

Proposition 6.2. *In the uniformly distributed quality case with* $n = 2$, *when* $P > 2\delta_{max}$, *both agents will exert* δ_{max} *in equilibrium and this is the only equilibrium.*

Proof. If both agents exert effort δ_{max}, then they both have an expected payoff of $\frac{P}{2} - \delta_{max}$, which is strictly greater than 0. Consider a deviation for an agent where she exerts less effort, e.g. $\delta_1 < \delta_{max}$. Remark 4.2 tells us that a critical point must satisfy $\frac{\partial}{\partial \delta_1}[\int_0^\infty f_{\delta_1}(x) F_{\delta_2}(x) dx] - \frac{1}{P} = 0$, or $\frac{P}{2\delta_{max}} - 1 = 0$. But $\frac{P}{2\delta_{max}} - 1 > 0$ for all $P > 2\delta_{max}$. Therefore, exerting effort δ_{max} is a best response to another player exerting δ_{max} effort.

Now suppose that agent 2 exerts effort $\delta_2 < \delta_{max}$, we consider the best-response to this. Consider agent 1 exerting effort δ_1. Remark 4.2 tells us that a critical point must satisfy $\frac{\partial}{\partial \delta_1}[\int_0^\infty f_{\delta_1}(x) F_{\delta_2}(x) dx] - \frac{1}{P} = 0$. If $\delta_1 > \delta_2$, this condition becomes $\frac{\delta_2}{2\delta_1^2} - \frac{1}{P} = 0$, with a critical point at $\delta_1 = \sqrt{\frac{P\delta_2}{2}}$. The derivative changes from positive to negative here, so we have a maximum. Therefore, if a player exerts $x < \delta_{max}$ effort, the best-response to this is exerting strictly greater effort. Note that if $\delta_1 < \sqrt{\frac{P\delta_2}{2}}$, the derivative is positive over the interval $\delta_1 \in [\delta_2, \delta_{max}]$, so the maximum occurs at δ_{max}, otherwise the max occurs at $\delta_1 = \sqrt{\frac{P\delta_2}{2}}$. Therefore we know for any strategy profile (x, y), where $x < y$, the agent exerting lesser effort is not best-responding. Similarly any strategy profile (x, x), where $x < \delta_{max}$ has both players not best-responding, which gives us the desired result. □

This proposition gives us an equilibrium characterization for the $P > 2\delta_{max}$ case. We now see, through a set of three lemmas leading to a proposition, that if $P \le 2\delta_{max}$ the sole pure strategy equilibrium is interior.

Lemma 6.3. *In the uniformly distributed quality case with* $n = 2$, *if agent 2 exerts* δ_2 *units of effort, where* $0 < \delta_2 < \frac{P}{2}$, *then agent 1's unique best-response* δ_1 *is in the interval* $(\delta_2, \frac{P}{2})$.

Lemma 6.4. *In the uniformly distributed quality case with $n = 2$, if agent 1 exerts $\frac{P}{2}$ units of effort, agent 2 has a (weak) best-response of exerting $\frac{P}{2}$ units of effort.*

Lemma 6.5. *In the uniformly distributed quality case with $n = 2$, if agent 2 exerts δ_2 units of effort, where $\delta_2 > \frac{P}{2}$, then agent 1's unique best response will be to exert 0 effort.*

Proposition 6.6. *In the uniformly distributed quality case with $n = 2$, when $P < 2\delta_{max}$, the only pure-strategy Nash equilibrium for the winner-take-all contest is the one in which each player exerts effort $\frac{P}{2}$.*

Proof. From Proposition 4.4, any equilibrium strategy profile must have both players exerting non-zero effort. Suppose there exists an equilibrium profile with player 1 exerting x units of effort and player 2 exerting y units of effort, with $x, y > 0$. We can assume that $x \leq y$. First consider the case that $0 < x < \frac{P}{2}$. If $y < \frac{P}{2}$, Lemma 6.3 tells us that player 1's best-response to y is a value strictly greater than y, so this cannot be part of a Nash equilibrium. If $y = \frac{P}{2}$, then player 2 is not best-responding to player 1 according to Lemma 6.3. If $y > \frac{P}{2}$, Lemma 6.5 tells us that player 1 is not best-responding. Now consider the case that $x > \frac{P}{2}$. Lemma 6.5 tells us that player 2 is not best-responding. Now consider the case that $x = \frac{P}{2}$. If $y > \frac{P}{2}$, Lemma 6.5 tells us that player 1 is not best-responding. If $y = \frac{P}{2}$, we know from Lemma 6.4, that this is a weak best-response. Thus $(\frac{P}{2}, \frac{P}{2})$ is the only pure strategy Nash equilibrium. □

Propositions 6.2 and 6.6 give us a complete characterization of the pure-strategy Nash equilibrium in winner-take-all contests with two agents in a model of costly effort and stochastic production, which we state here:

Theorem 6.7. *In the uniformly distributed quality case with $n = 2$: in the winner-take-all contest with $P < 2\delta_{max}$, the only pure-strategy Nash equilibrium has both agents exert effort $\frac{P}{2}$; in the winner-take-all contest with $P \geq 2\delta_{max}$, the only pure-strategy Nash equilibrium has both agents exert effort δ_{max}.*

6.2 Equilibrium analysis for general n (uniform)

We extend the equilibrium analysis above to settings with general n and give a characterization of the pure-strategy Nash equilibrium.

Lemma 6.8. *In the uniformly distributed quality case, for any value of P and δ_{max}, any equilibrium strategy profile has all agents who exert non-zero effort exerting the same amount of effort.*

Lemma 6.9. *In the uniformly distributed quality case, when $P \geq n\delta_{max}$, the only Nash equilibrium has each agent exert full effort.*

Lemma 6.10. *In the uniformly distributed quality case, when $P < n\delta_{max}$ there is no pure-strategy symmetric Nash equilibrium in a winner-take-all contest with $n \geq 3$.*

Lemma 6.11. *In the uniformly distributed quality case, when $P \leq 2\delta_{max}$ the only pure-strategy Nash equilibrium is one in which 2 players exert effort $\frac{P}{2}$ and the remaining agents exert 0 effort.*

Lemma 6.12. *In the uniformly distributed quality case, when $P \in (2\delta_{max}, n\delta_{max})$, consider $i \in \{2, \dots, n-1\}$ such that $P \in (i\delta_{max}, (i+1)\delta_{max}]$. Any effort profile in which i agents exert effort δ_{max} and $n-i$ agents exert effort 0 is a Nash equilibrium; moreover, if $P = (i+1)\delta_{max}$, any effort profile in which $i+1$ agents exert δ_{max} effort (with $n-i-1$ agents exerting 0 effort) is also a Nash equilibrium. This is an exhaustive characterization of the pure strategy Nash equilibria.*

Lemmas 6.9, 6.11 and 6.12 give us a complete pure-strategy equilibrium characterization for winner-take-all contests for general n:

Theorem 6.13. *In the case of uniformly distributed quality with $n \geq 2$ agents:*

- *When $P \leq 2\delta_{max}$, the only pure-strategy Nash equilibria are those in which exactly two agents exert effort $\frac{P}{2}$ and all others exert 0 effort.*
- *If $\exists i \in \{2, \dots, n-1\}$ such that $P \in (i\delta_{max}, (i+1)\delta_{max})$, the only pure-strategy Nash equilibria are those in which exactly i agents exert effort δ_{max} and all others exert 0 effort.*
- *If $\exists i \in \{2, \dots, n-1\}$ such that $P = (i+1)\delta_{max}$, there are two classes of pure-strategy Nash equilibria: one in which exactly i agents exert effort δ_{max} and all others exert 0 effort, and the other in which exactly $i+1$ agents exert δ_{max} and all others exert 0 effort.*
- *Finally, when $P > n\delta_{max}$, in the unique pure-strategy Nash equilibrium all n agents exert effort δ_{max}.*

Given the equilibrium characterizations above, we can now determine the principal's strategic choice of prize value P, given her value v. This will allow us, in Section 6.4, to compare the efficiency of a winner-take-all contest with the optimal solution.

6.3 Equilibrium prize value (uniform)

Since the principal's best choice of prize depends on the equilibrium strategies of the agents, it is difficult to say much about it independent of a given quality distribution. Here we address the uniformly distributed quality case, and note that other cases such as exponentially distributed quality can be handled similarly.[5] We start by analyzing the principal's problem in the $n = 2$ case and then extend the analysis to the general n case.

Lemma 6.14. *In the uniformly distributed quality case with $n = 2$, in equilibrium:*

- *When $v > 3$, the principal chooses $P = 2\delta_{max}$, with each agent exerting effort δ_{max}.*
- *When $v < 3$, the principal chooses $P = 0$.*
- *When $v = 3$, the principal is indifferent between all prizes in the interval $P \in [0, 2\delta_{max}]$.*

[5] We make the assumption that when there are multiple equilibria for a given prize value set by the principal, the principal can choose the more desirable equilibrium.

Now we consider the principal's problem in the general n case. Given the equilibrium analysis in the previous section and the fact that the principal is strategic, it suffices to restrict attention to the minimal P to implement a given outcome, which leads to the following lemma.

Lemma 6.15. *$\forall i \in \{2, \ldots, n-1\}$, $\exists v \in \Re^+$ such that the principal's expected utility in equilibrium is the same if he chooses the minimal P such i agents participate with full effort, or the minimal P' such that $i+1$ agents participate with full effort.*

In light of this lemma, we define notation $v(i, i+1)$, for arbitrary $i \in \{2, \ldots, n-1\}$, to be the value at which the principal is indifferent between choosing the minimal P that gets i agents to participate with full-effort, or choosing the minimal P' that gets $i+1$ to participate with full-effort, given that they play the unique pure strategy equilibrium (see Theorem 6.13).

Lemma 6.16. *In the uniformly distributed quality case, $\forall i \in \{2, \ldots, n-2\}$, $v(i, i+1) < v(i+1, i+2)$.*

Lemma 6.17. *In the uniformly distributed quality case, if $\exists i \in \{3, ..., n-1\}$ such that $v \in (i(i+1), (i+1)(i+2))$, the principal sets a prize of $i\delta_{max}$ to have i agents exert full effort. Alternatively if $v > n(n+1)$, the principal sets a prize of $n\delta_{max}$ to have n agents exert full effort. Finally, if $\exists i \in \{3, ..., n-1\}$ such that $v = i(i+1)$, the principal is indifferent between setting a prize of $(i-1)\delta_{max}$ and $i\delta_{max}$.*

Lemmas 6.14 and 6.17 establish the following theorem for $n \geq 3$:

Theorem 6.18. *In the uniformly distributed quality case with $n \geq 3$, in equilibrium:*

- *When $v < 3$, the principal sets $P = 0$.*
- *When $v = 3$, the principal is indifferent between all prizes in the interval $P \in [0, 2\delta_{max}]$.*
- *When $3 < v < 12$, the principal sets $P = 2\delta_{max}$ in order to have 2 agents exert full effort.*
- *When, for some $i \in \{3, \ldots, n-1\}$, $v = i(i+1)$, the principal is indifferent between setting a prize of $i\delta_{max}$ and a prize of $(i-1)\delta_{max}$.*
- *When, for some $i \in \{3, \ldots, n-1\}$, $v \in (i(i+1), (i+1)(i+2))$, the principal sets $P = i\delta_{max}$ to have i agents exert full effort.*
- *When $v > n(n+1)$, the principal sets $P = n\delta_{max}$ to have n agents exert full effort.*

6.4 Efficiency of winner-take-all contests (uniform)

We will now use this equilibrium characterization to evaluate the efficiency of winner-take-all when quality is uniformly distributed. Interestingly, we find that for a wide range of values for the principal, a winner-take-all scheme will yield the efficient effort policy in equilibrium; but for certain values, it will not.

The following lemma may look virtually identical to Lemma 6.16 above, however we note the following lemma is given with respect to *social welfare*, whereas Lemma 6.16 is given with respect to the principal's utility. In what follows $v(i, i+1)$ denotes the value of v at which the social welfare of having i agents exert full effort is the same as the social welfare of having $i+1$ agents exert full effort.

Lemma 6.19. *In the uniformly distributed quality case,* $\forall i \in \{1, \ldots, n-1\}$, $v(i-1, i) = i(i+1)$*. Moreover, when* $v > v(i-1, i)$*, the social welfare of* i *agents exerting full effort is strictly greater than the social welfare of* $i-1$ *agents exerting full effort and when* $v < v(i-1, i)$*, the social welfare of* $i-1$ *agents exerting full effort is strictly greater than the social welfare of* i *agents exerting full effort.*

Proof. When i agents exert full effort, the expected value to the principal is $\frac{i}{i+1}\delta_{max}v$. Hence, the social welfare of having i agents exert full effort is: $\frac{i}{i+1}\delta_{max}v - i\delta_{max}$. To determine $v(i-1, i)$, we need $v(i-1, i)$ to satisfy: $\frac{i}{i+1}\delta_{max}v(i-1, i) - i\delta_{max} = \frac{i-1}{i}\delta_{max}v(i-1, i) - (i-1)\delta_{max}$, or in other words, $v(i-1, i) = i(i+1)$. When $v > v(i-1, i)$, $\frac{i}{i+1}\delta_{max}v - i\delta_{max} > \frac{i-1}{i}\delta_{max}v - (i-1)\delta_{max}$ and when $v < v(i-1, i)$, $\frac{i}{i+1}\delta_{max}v - i\delta_{max} < \frac{i-1}{i}\delta_{max}v - (i-1)\delta_{max}$. $\square$

Corollary 6.20. *In the uniformly distributed quality case,* $\forall i \in \{1, \ldots, n-1\}$, $v(i-1, i) < v(i, i+1)$.

Theorem 6.21. *In the uniformly distributed quality case, if* $\exists i \in \{0, ..., n-1\}$ *such that* $i(i+1) < v < (i+1)(i+2)$*, the efficient solution involves* i *agents exerting full effort (with* $n-i$ *agents exerting 0 effort). Alternatively if* $v > n(n+1)$*, the efficient solution involves all* n *agents exerting full effort. Finally, if* $\exists i \in \{1, ..., n\}$ *such that* $v = i(i+1)$*, there are two classes of efficient solutions: one where* i *agents exert full effort and one where* $i-1$ *agents exert full effort (with the other* $n-i$ *or* $n-i+1$ *agents exerting 0 effort).*

Proof. Since the uniform distribution satisfies the extreme effort condition, it suffices to consider solutions where agents either exert full effort or no effort. We establish this theorem via induction. We claim that for any value of $v > v(i-1, i)$, having i agents exert full effort leads to strictly higher social welfare than having $0 \leq j < i$ agents exert full effort (if any such j exists). When $i = 0$, there is no value of j, such that $0 \leq j < i$, so the claim trivially holds for all $v \geq 0$. For arbitrary integer $i > 0$, assume that we have proven the claim for $i-1$, which means that for $v > v(i-2, i-1)$, having $i-1$ agents exert full effort leads to higher social welfare than having $0 \leq j < i-1$ agents exert full effort. We know that there exists a unique point $v(i-1, i)$ at which the social welfare of having $i-1$ agents exert full effort is the same as the social welfare of having i agents exert full effort. Moreover, we know that for all $v > v(i-1, i)$, the social welfare of having i agents exert full effort is strictly greater than the social welfare of having $i-1$ agents exert full effort. Thus having i agents exert full effort yields greater social welfare than having $0 \leq j < i$ agents exert full effort for any value $v > v(i-1, i)$. The final step of the inductive argument gives us that having n agents exert full effort leads to greater social welfare than having $0 \leq j < n$ agents exert full effort for any value $v > v(n-1, n)$.

We can establish a similar argument via induction that at any value of $v < v(i, i+1)$, having i agents exert full effort leads to strictly higher social welfare than having $n \geq j > i$

agents exert full effort (if any such j exists). When $i = n$, there is no value of j, such that $n \geq j > i$, so the claim trivially holds for all $v \geq v(n-1, n)$. For arbitrary integer $i < n$, assume that we have proven the claim for $i+1$, which means that for $v < v(i+1, i+2)$, having $i+1$ agents exert full effort leads to higher social welfare than having $n \geq j > i+1$ agents exert full effort. We know that for all $v < v(i, i+1)$, having i agents exert full effort leads to greater social welfare than having $i+1$ agents exert full effort. Thus for all $v < v(i, i+1)$, having i agents exert full effort yields greater social welfare than having $n \geq j > i$ agents exert full effort. Combining these two inductive arguments, with Corollary 6.20, gives us that having i agents exert full effort when $v(i-1, i) < v < v(i, i+1)$ maximizes social welfare for all $1 \leq i \leq n-1$. □

Theorems 6.18 and 6.21 give us the following result—the winner-take-all contest can be inefficient for a relevant range of principal's values v, specifically, the range over which the efficient solution requires that exactly one agent exert full effort with all ther agenst exerting 0 effort.

Theorem 6.22. *In the uniformly distributed quality case, for any* $n \geq 2$*, the winner-take-all contest is inefficient in equilibrium for* $2 \leq v \leq 6$*.*

We note from Theorem 6.21 that when $0 \leq v \leq 2$, the efficient outcome has 0 agents exerting full effort and when $2 < v < 6$, the efficient outcome has 1 agent exerting full effort. When $0 \leq v \leq 2$, the winner-take-all contest implements the efficient outcome. When $2 < v < 3$, the winner-take-all contest has 0 agents exerting effort in equilibrium and is therefore perfectly inefficient. When $3 \leq v < 6$, the winner-take-all contest has 2 agents exerting full effort in equilibrium, when the efficient outcome is to have 1 agent exert full effort in equilibrium. When $v \geq 6$, the winner-take-all contest implements the efficient outcome. It is perhaps surprising that for most values of v, the winner-take-all contest is in fact efficient.

7 Conclusion

In this paper we provided a thorough analysis of winner-take-all contest mechanisms in a model of stochastic production, deriving general characteristics of best response functions and giving pure-strategy equilibrium characterizations for the cases of uniformly and exponentially distributed quality. We found that in both the uniform case and the exponential case, the winner-take-all contest can implement the efficient outcome for a large range of values of v (the principal's value), but not all.

It is important to emphasize that even in the cases where the winner-take-all paradigm yields an efficient outcome in equilibrium, without a centralization procedure it is highly questionable whether the equilibrium would obtain, as significantly non-trivial agent coordination would be required; for instance, when the set of equilibria has 2 out of the n agents participate with full effort, and the other $n-2$ not participate, how would each agent determine whether to participate or not? Thus, perhaps a safer statement of the positive side of our results is to say that for certain principal values a *centrally coordinated* implementation of a winner-take-all contest may yield efficiency with rational strategic agents.

References

[1] E. Amann and W. Leininger. Asymmetic all-pay auctions with incomplete information: The two-player case. *Games and Economic Behavior*, pages 1 – 18, 1996.

[2] N. Archak and A. Sundararajan. Optimal design of crowdsourcing contests. In *Proceedings of the Thirtieth International Conference on Information Systems (ICIS-09)*, 2009.

[3] M. R. Baye, D. Kovenock, and C. G. de Vries. The all-pay auction with complete information. *Economic Theory*, 8:291–305, 1996.

[4] P. Bertoletti. On the reserve price in all-pay auctions with complete information and lobbying games. manuscript, 2010.

[5] R. Cavallo and S. Jain. Efficient crowdsourcing contests. In *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-12)*, 2012.

[6] S. Chawla, J. Hartline, and B. Sivan. Optimal crowdsourcing contests. In *Proceedings of the 23rd ACM-SIAM Symposium on Discrete Algorithms (SODA-12)*, pages 856–868, 2011.

[7] D. DiPalantino and M. Vojnovic. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC-09), Stanford, CA*, pages 119–128, 2009.

[8] R. Fullerton and R. P. McAfee. Auctioning entry into tournaments. *Journal of Political Economy*, 107:573 – 605, 1999.

[9] A. Ghosh and P. Hummel. Implementing optimal outcomes in social computing: A game-theoretic approach. In *Proceedings of the 21st ACM International World Wide Web Conference (WWW-12)*, 2012.

[10] J. Green and N. Stokey. A comparison of tournaments and contracts. *Journal of Political Economy*, 91(3):349–364, 1983.

[11] A. Hilman and J. Riley. Politically contestable rents and transfers. *Economics and Politics*, 1:17–39, 1989.

[12] K. A. Konrad and W. Leininger. The generalized stackelberg equilibrium of the all-pay auction with complete information. *Review of Economic Design*, 11(2):165–174, 2007.

[13] V. Krishna and J. Morgan. An analysis of the war of attrition and the all-pay auction. *Journal of Economic Theory*, 72:343 – 362, 1997.

[14] E. Lazear and S. Rosen. Rank order tournaments as optimum labor contracts. *Journal of Political Economy*, 89:841–864, 1981.

[15] T. X. Liu, J. Yang, L. A. Adamic, and Y. Chen. Crowdsourcing with all-pay auctions: a field experiment on taskcn. Workshop on Social Computing and User Generated Content, 2011.

[16] D. Minor. Increasing efforts through incentives in contests. Manuscript, 2011.

[17] B. Moldovanu and A. Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91:542–558, June 2001.

[18] B. Moldovanu and A. Sela. Contest architecture. *Journal of Economic Theory*, 126(1):70–96, January 2006.

[19] H. Moulin. *Game Theory for the Social Sciences.* New York University Press, 1986.

[20] B. Nalebuff and J. Stiglitz. Prizes and incentives: Towards a general theory of compensation and competition. *Bell Journal of Economics*, 14:21–43, 1983.

[21] E. Segev and A. Sela. Sequential all-pay auctions with head starts. manuscript, 2011.

[22] C. Taylor. Digging for golden carrots: An analysis of research tournaments. *American Economic Review*, (85):873–890, 1995.

[23] G. Tullock. Efficient rent-seeking. In J. M. Buchanan, R. D. Tollison, and G. Tullock, editors, *Towards a Theory of the Rent-Seeking Society*, pages 97–112. Texas A&M University Press, 1980.

[24] R. Weber. Auctions and competitive bidding. In H. Young, editor, *Fair Allocation*, pages 143–170. American Mathematical Society, 1985.

A Technical Lemmas

Lemma 5.1. *In the exponentially distributed quality case with* $n = 2$, $\forall \delta_2 \in [0, \delta_{max}]$, *the function* $\frac{\partial}{\partial \delta_1} \int_0^\infty [f_{\delta_1}(x) F_{\delta_2}(x) dx] - \frac{1}{P}$ *has at most one zero with respect to* δ_1 *over the interval* $(0, \delta_{max}]$, *for any value of* δ_2.

Proof. For the exponential distribution, we have:

$$\int_0^\infty f_{\delta_1} F_{\delta_2}(x) dx = \int_0^\infty \frac{1}{\delta_1} e^{-x/\delta_1}(1 - e^{-x/\delta_2}) \tag{4}$$

$$= \int_0^\infty \frac{1}{\delta_1} e^{-x/\delta_1} dx - \int_0^\infty \frac{1}{\delta_1} e^{-x(\frac{1}{\delta_1}+\frac{1}{\delta_2})} dx \tag{5}$$

$$= \frac{\delta_1}{\delta_1 + \delta_2} \tag{6}$$

Taking the derivative of this, we get: $\frac{\delta_2}{(\delta_1+\delta_2)^2}$, which is strongly monotonically decreasing in δ_1 at all values of $\delta_1 > 0$, for arbitrary δ_2. Therefore, $\frac{\delta_2}{(\delta_1+\delta_2)^2} - \frac{1}{P}$ is strongly monotonically decreasing in δ_1 and the expression has at most one zero, e.g. a value of $\delta_1 \in (0, \delta_{max}]$ that satisfies $\delta_1 = \sqrt{P\delta_2} - \delta_2$. If there is a zero, it must be a maximum, following Remark 4.2. □

Lemma 5.3. *In the exponentially distributed quality case with* $n = 2$, *if agent 2 exerts* δ_2 *units of effort, where* $0 < \delta_2 < \frac{P}{4}$, *then agent 1's unique best-response* δ_1 *is in the interval* $(\delta_2, \frac{P}{4})$.

Proof. Applying Lemma 5.1, the critical point must satisfy $\frac{\delta_2}{(\delta_1+\delta_2)^2} = \frac{1}{P}$, with a zero at $\delta_1 = \sqrt{\delta_2}(\sqrt{P} - \sqrt{\delta_2})$. Additionally from Lemma 5.1, this must be a maximum. Since $\delta_2 < \frac{P}{4}$, we know that $\sqrt{P} - \sqrt{\delta_2} > \sqrt{\delta_2}$ and the maximum occurs at $\delta_1 > \delta_2$. □

Lemma 5.4. *In the exponentially distributed quality case with* $n = 2$, *if agent 2 exerts* $\frac{P}{4}$ *units of effort, agent 1's unique best-response is* $\delta_1 = \frac{P}{4}$.

Proof. Applying Lemma 5.1, the critical point must satisfy $\frac{\delta_2}{(\delta_1+\delta_2)^2} - 1 = \frac{1}{P}$, with a zero at $\delta_1 = \sqrt{\delta_2}(\sqrt{P} - \sqrt{\delta_2})$. When $\delta_2 = \frac{P}{4}$, $\delta_1 = \frac{P}{4}$. Additionally from Lemma 5.1, this must be a maximum. □

Lemma 5.5. *In the exponentially distributed quality case with* $n = 2$, *if agent 2 exerts* δ_2 *units of effort, where* $\delta_2 > \frac{P}{4}$, *then agent 1's unique best response* δ_1 *is in the interval* $[0, \frac{P}{4})$.

Proof. From Lemma 5.1, the critical point must satisfy $P \cdot \frac{\delta_2}{(\delta_1+\delta_2)^2} - 1$, with a zero at $\delta_1 = \sqrt{\delta_2}(\sqrt{P} - \sqrt{\delta_2})$. Additionally from Lemma 5.1, this must be a maximum. Note that $\sqrt{\delta_2}(\sqrt{P} - \sqrt{\delta_2})$ is maximized when $\delta_2 = \frac{P}{4}$ and so $\sqrt{\delta_2}(\sqrt{P} - \sqrt{\delta_2}) < \frac{P}{4}$, when $\delta_2 > \frac{P}{4}$. □

Lemma 5.9. —em For the exponentially distributed quality case, there is an efficient policy in which each agent exerts either no effort $\delta_i = 0$ or full effort $\delta_i = 1$.

Proof. The expected marginal impact on efficiency from i exerting effort δ_i, given an arbitrary β representing the maximum quality that is to be realized by the production of the other agents equals:

$$\int_\beta^\infty v f_{\delta_i}(x)(x-\beta)dx - \delta_i = \int_\beta^\infty v\frac{1}{\delta_i}e^{-x/\delta_i}(x-\beta)dx - \delta_i \quad (7)$$

$$= \int_\beta^\infty \frac{v}{\delta_i}e^{-x/\delta_i}xdx - \int_\beta^\infty \frac{v\beta}{\delta_i}e^{-x/\delta_i}dx - \delta_1 \quad (8)$$

$$= \frac{\delta_i v}{e^\beta/\delta_i} - \delta_i \quad (9)$$

Taking the partial derivative with respect to δ_i, we get: $ve^{-\beta/\delta_i} + \frac{\beta}{\delta_i}ve^{-\beta/\delta_i} - 1$. The derivative is increasing in δ_i, regardless of the value of β, thus the derivative never changes from positive to negative and the maximum lies at one of the extremes. □

Lemma 6.1. *In the uniformly distributed quality case with $n = 2$, $\forall\delta_2 \in [0, \delta_{max}]$, the function $\frac{\partial}{\partial\delta_1}\int_0^\infty[f_{\delta_1}(x)F_{\delta_2}(x)dx] - \frac{1}{P}$ has at most one zero with respect to δ_1 over the interval $(0, \delta_{max}]$, for any value of $\delta_2 \neq \frac{P}{2}$.*

Proof. If $\delta_2 = 0$ the lemma holds trivially. Assume $\delta_2 > 0$. Note that if $\delta_1 \in (0, \delta_2]$, $\int_0^\infty f_{\delta_1}(x)F_{\delta_2}(x)dx = \frac{\delta_1}{2\delta_2}$. Taking the derivative of this we get $\frac{1}{2\delta_2}$, a constant with respect to δ_1. Therefore, $\frac{1}{2\delta_2} - \frac{1}{P}$ is also a constant with respect to d_1. Moreover, when $\delta_2 \neq \frac{P}{2}$, this expression is a non-zero constant. Now note that if $\delta_1 \geq \delta_2$, $\int_0^\infty f_{\delta_1}(x)F_{\delta_2}(x)dx = 1 - \frac{\delta_2}{2\delta_1}$. Taking the derivative, we get $\frac{\delta_2}{2\delta_1^2}$, which is strictly decreasing in δ_1. Therefore, $\frac{\delta_2}{2\delta_1^2} - \frac{1}{P}$ is strictly decreasing in δ_1 as well. We note that at $\delta_1 = \delta_2$, $\frac{\delta_1}{2\delta_2} = 1 - \frac{\delta_2}{2\delta_1}$ and that $\frac{1}{2\delta_2} = \frac{\delta_2}{2\delta_1^2}$, so both $\int_0^\infty f_{\delta_1}(x)F_{\delta_2}(x)dx$ and $\frac{\partial}{\partial\delta_1}[\int_0^\infty f_{\delta_1}(x)F_{\delta_2}(x)dx]$ are continuous functions. Therefore, when $\delta_2 \neq \frac{P}{2}$, there is at most one zero with respect to δ_1 over the interval $(0, \delta_{max}]$. □

Lemma 6.3. *In the uniformly distributed quality case with $n = 2$, if agent 2 exerts δ_2 units of effort, where $0 < \delta_2 < \frac{P}{2}$, then agent 1's unique best-response δ_1 is in the interval $(\delta_2, \frac{P}{2})$.*

Proof. If agent 2 exerts δ_2 units of effort and agent 1 exerts δ_1 units of effort, where $\delta_1 < \delta_2$, agent 1's expected payoff becomes $\frac{P}{2}\cdot\frac{\delta_1}{\delta_2} - \delta_1 = \frac{\delta_1}{\delta_2}(\frac{P}{2} - \delta_2)$. Since we know that $\delta_2 < \frac{P}{2}$, agent 1's utility is maximized when $\delta_1 = \delta_2$. Now suppose that agent 1 exerts $\delta_1 > \delta_2$ units of effort. Lemma 6.1 tells us that a critical point must satisfy $\frac{\delta_2}{2\delta_1^2} - \frac{1}{P} = 0$ or $\delta_1 = \sqrt{\frac{P\delta_2}{2}}$. The derivative switches from positive to negative, so we have a maximum. Agent 1's best-response will be to exert $\delta_1 = \sqrt{\frac{P\delta_2}{2}}$. Since $\delta_2 < \frac{P}{2}$, $\delta_2 < \delta_1 < \frac{P}{2}$. □

Lemma 6.4. *In the uniformly distributed quality case with $n = 2$, if agent 1 exerts $\frac{P}{2}$ units of effort, agent 2 has a (weak) best-response of exerting $\frac{P}{2}$ units of effort.*

Proof. If agent 2 exerts δ_2 units of effort and agent 1 exerts δ_1 units of effort, where $0 \leq \delta_1 \leq \delta_2$, agent 1's expected payoff becomes $\frac{P}{2}\cdot\frac{\delta_1}{\delta_2} - \delta_1$. This is equal to 0 for all $0 \leq \delta_1 \leq \delta_2$. Now suppose that agent 1 exerts $\delta_1 > \delta_2$ units of effort. The derivative of the best response function, when $\delta_1 > \delta_2$, is $\frac{\delta_2}{2\delta_1^2} = \frac{1}{P}$. This has a maximum at $\delta_1 = \sqrt{\frac{P\delta_2}{2}} = \frac{P}{2}$. For all $\delta_1 > \sqrt{\frac{P\delta_2}{2}}$, the utility function is decreasing. Therefore, player 1 is indifferent from exerting effort on the interval $[0, \frac{P}{2}]$, if player 2 exerts $\frac{P}{2}$ units of effort. □

Lemma 6.5. *In the uniformly distributed quality case with $n = 2$, if agent 2 exerts δ_2 units of effort, where $\delta_2 > \frac{P}{2}$, then agent 1's unique best response will be to exert 0 effort.*

Proof. If agent 2 exerts δ_2 units of effort and agent 1 exerts δ_1 units of effort, where $0 \leq \delta_1 \leq \delta_2$, agent 1's expected payoff becomes $\frac{P}{2} \cdot \frac{\delta_1}{\delta_2} - \delta_1 = \frac{\delta_1}{\delta_2}(\frac{P}{2} - \delta_2)$. Since we know that $\frac{P}{2} < \delta_2$, agent 1's utility is negative over the entire interval $0 < \delta_1 \leq \delta_2$ and 0 when $r = 0$. Now suppose that agent 1 exerts $\delta_1 > \delta_2$ units of effort. We know from Lemma 6.1, a critical point must satisfy $\frac{\delta_2}{2\delta_1^2} - \frac{1}{P} = 0$ with a maximum at $\delta_1 = \sqrt{\frac{P\delta_2}{2}}$. This occurs at a value of $\delta_1 < \delta_2$ and the utility function is decreasing for all $\delta_1 > \sqrt{\frac{P\delta_2}{2}}$. Therefore, the maximum of agent 1's utility function over the interval $\delta_1 \geq \delta_2$ occurs at $\delta_1 = \delta_2$, where the utility function is $\frac{P}{2} - \delta_2$, which is less than 0. Therefore agent 1 maximizes her utility by exerting 0 effort. □

Lemma 6.8. *In the uniformly distributed quality case, for any value of P and δ_{max}, any equilibrium strategy profile has all agents who exert non-zero effort exerting the same amount of effort.*

Proof. Consider an arbitrary strategy profile of efforts $(\delta_1, ..., \delta_n)$ that consists of i different effort levels $\gamma_1 > \gamma_2 > ... > \gamma_i$, with $a_1, ..., a_i$ number of agents exerting these efforts, respectively. The expected payoff to an agent who exerts e_i units of effort can be expressed as: $\frac{\gamma_i^n}{\gamma_1^{a_1} \cdot ... \cdot \gamma_i^{a_i}} \cdot \frac{P}{n} - \gamma_i$ or $\gamma_i \cdot (\frac{\gamma_i^{n-1}}{\gamma_1^{a_1} \cdot ... \cdot \gamma_i^{a_i}} \cdot \frac{P}{n} - 1)$. If this is part of an equilibrium strategy profile, this utility must be at least 0. If an agent exerting γ_i deviates and exerts effort γ_{i-1}, her payoff will be at least: $\frac{\gamma_{i-1}^n}{\gamma_1^{a_1} \cdot ... \cdot \gamma_{i-1}^{a_{i-1}+a_i}} \cdot \frac{P}{n} - \gamma_{i-1}$ or $\gamma_{i-1} \cdot (\frac{\gamma_{i-1}^{n-1}}{\gamma_1^{a_1} \cdot ... \cdot \gamma_{i-1}^{a_{i-1}+a_i}} \cdot \frac{P}{n} - 1)$. Since $\gamma_{i-1} > \gamma_i$ and $\frac{\gamma_i^n}{\gamma_1^{a_1} \cdot ... \cdot \gamma_i^{a_i}} < \frac{\gamma_{i-1}^n}{\gamma_1^{a_1} \cdot ... \cdot \gamma_{i-1}^{a_{i-1}+a_i}}$, an agent exerting γ_i units of effort has a profitable deviation. Since no agent can exert strictly less effort than another, this establishes the desired result. □

Lemma 6.9. *In the uniformly distributed quality case, when $P \geq n\delta_{max}$, the only Nash equilibrium has each agent exert full effort.*

Proof. If all agents exert effort δ_{max}, then they each have an expected payoff of $\frac{P}{n} - \delta_{max}$, which is strictly greater than 0. Consider a deviation for an agent where she exerts less effort, e.g. $\delta_i < \delta_{max}$. Her expected utility becomes: $(\frac{\delta_i}{\delta_{max}})^{n-1} \cdot \frac{P}{n} - \delta_i$, or $\frac{\delta_i}{\delta_{max}} \cdot ((\frac{\delta_i}{\delta_{max}})^{n-2} \frac{P}{n} - \delta_{max})$, which is strictly less than $\frac{P}{n} - \delta_{max}$. Therefore, exerting δ_{max} is a best-response to all other players exerting δ_{max}.

Due to Lemma 6.8, we know that all agents exerting non-zero effort, must exert the same effort, so it suffices to restrict our attention to such strategy profiles. Consider an arbitrary such strategy profile with i agents exerting effort δ_2 and $n - i$ agents exerting effort 0. Suppose an agent who exerts effort 0 decides to exert effort $\delta_1 > \delta_2$. Take the derivative of the best response, we find her utility is maximized at a point δ_1 that satisfies $\frac{\delta_2}{\delta_1^2} - \frac{\delta_2^i}{(i+1)\delta_1^2} = \frac{1}{P}$ or $\delta_1 = \sqrt{P(\delta_2 - \frac{\delta_2^i}{i+1})} > \sqrt{P\delta_2 \frac{i}{i-1}} > \sqrt{P\delta_2 \frac{1}{n}}$. The derivative changes from positive to negative so we have a maximum. Note that $\frac{P}{n} > \delta_{max} > \delta_2$, so $\delta_2 < \delta_1$. If $\delta_{max} < \sqrt{P(\delta_2 - \frac{\delta_2^i}{i+1})}$, we know the derivative is positive over the interval $\delta_1 \in [\delta_2, \delta_{max}]$, thus the maximum occurs at δ_{max}. If an agent who exerts δ_2 effort decides to exert δ_1 effort, she would have a critical point at $\delta_1 = \sqrt{P(\delta_2 - \frac{\delta_2^{i-1}}{i})} > \sqrt{P\delta_2 \frac{i-1}{i}} > \sqrt{P\delta_2 \frac{1}{n}}$. The derivative changes from positive to negative so we have a maximum. Note that $\frac{P}{n} > \delta_{max} > \delta_2$, so $\delta_2 < \delta_1$. If $\delta_{max} < \sqrt{P(\delta_2 - \frac{\delta_2^i}{i+1})}$, we know the derivative is positive over the interval $\delta_1 \in [\delta_2, \delta_{max}]$, thus the maximum occurs at δ_{max}. Since for any value of δ_2 and any value of i, there exists a profitable deviation, the only equilibrium is the one in which all agents exert full effort. □

Lemma 6.10. *In the uniformly distributed quality case, when $P < n\delta_{max}$ there is no pure-strategy symmetric Nash equilibrium in a winner-take-all contest with $n \geq 3$.*

Proof. Consider a strategy profile in which each player exerts effort $0 < \delta_2 < \delta_{max}$. (Actually we know that $x \leq \frac{P}{n}$.) The utility for each player in such a strategy profile is $\frac{P}{n} - x$. Consider a deviation for a player in which she decides to play effort $\delta_1 > \delta_2$. Taking the derivative of the best response, we have a critical point at $\delta_1 = \sqrt{P(\delta_2 - \frac{\delta_2^{n-1}}{n})} > \sqrt{P\delta_2 \frac{n-1}{n}}$. The derivative changes from positive to negative so we have a maximum. Note that $\frac{P(n-1)}{n} > \delta_{max}$, so $\delta_1 > \delta_2$. Note that if $\sqrt{P(\delta_2 - \frac{\delta_2^{n-1}}{n})} > \delta_{max}$, the derivative is positive over the interval $\delta_1 \in [\delta_2, \delta_{max}]$. Therefore, this player's utility is maximized when $\delta_1 = \min(\delta_{max}, \sqrt{P(\delta_2 - \frac{\delta_2^{n-1}}{n}})$.

Now consider the case that $x = \delta_{max}$. Each agent has an expected payoff of $\frac{P}{n} - \delta_{max} < 0$. Therefore, if an agent deviates and exerts 0, her expected utility becomes 0. Finally consider the case that $x = 0$. Each player has an expected payoff of 0. An agent who exerts effort ϵ, where $\epsilon < P$ has a profitable deviation. □

Lemma 6.11. *In the uniformly distributed quality case, when $P \leq 2\delta_{max}$ the only pure-strategy Nash equilibrium is one in which 2 players exert effort $\frac{P}{2}$ and the remaining agents exert 0 effort.*

Proof. From Proposition 4.4, no equilibrium strategy profile involves a single agent who exerts effort, so it suffices to consider strategy profiles where at least two agents exert effort. We start by considering the case that exactly two agents exert effort. From Proposition 6.6, we know that if two agents are exerting effort $\frac{P}{2}$, they are best-responding to each other. We also know from Proposition 6.6 that if two agents exert anything other than $\frac{P}{2}$, they are not best-responding to each other, thus any strategy profile with exactly two agents exerting non-zero effort must have them each exerting effort $\frac{P}{2}$, in order for this to be part of a Nash equilibrium.

Consider whether an agent exerting 0 effort would wish to deviate. An agent exerting 0 effort, has an expected utility of 0. If this agents deviates and exerts $\delta_1 < \frac{P}{2}$ units of effort, we get that her expected utility becomes $(\frac{\delta_1}{P/2})^2 \cdot \frac{P}{3} - \delta_1$, which is maximized when $\delta_1 = \frac{P}{2}$. However, when $\delta_1 = \frac{P}{2}$, the utility becomes negative. Therefore this agent does not have a profitable deviation when $\delta_1 \leq \frac{P}{2}$.

Consider $\delta_1 > \frac{P}{2}$. Taking the derivative of the best response function, a critical point must satisfy: $\delta_1 = \sqrt{P(\frac{P}{2} - \frac{(P/2)^2}{3})} > \frac{P}{\sqrt{3}} > \frac{P}{2}$. The derivative changes from positive to negative at this point so we have maximum. We notice that the derivative is positive over the entire interval $\delta_1 \in [\frac{P}{2}, \sqrt{P(\frac{P}{2} - \frac{(P/2)^2}{3})}]$ and at $\delta_1 = \sqrt{P(\frac{P}{2} - \frac{(P/2)^2}{3})}$, we find the utility is negative. If $\sqrt{P(\frac{P}{2} - \frac{(P/2)^2}{3})} \leq \delta_{max}$, this gives us that the utility function is negative for all $\delta_1 \geq \frac{P}{2}$ and thus the agent has no profitable deviation. If $\sqrt{P(\frac{P}{2} - \frac{(P/2)^2}{3})} > \delta_{max}$, we know that the derivative is positive on the entire interval $\delta_1 \in [\frac{P}{2}, \delta_{max}]$ and that the utility is negative at $\delta_1 = \frac{P}{2}$ and $\delta_1 = \delta_{max}$, thus the agent has no profitable deviation.

Finally, we must consider strategy profiles in which more than 2 agents exert non-zero effort. In such a strategy profile each agent who exerts non-zero effort must be exerting the same amount of effort (according to Lemma 6.8) and these agents collectively exert effort at most P (according to Proposition 4.3). Thus any equilibrium profile that has i agents exert non-zero effort (for $i \geq 3$), must have these agents exert δ units of effort where δ is at most $\frac{P}{i}$. From the proof of Lemma 6.10, we know that an agent who exerts effort at most $\frac{P}{i}$, would gain from deviating and exerting more effort. □

Lemma 6.12. *In the uniformly distributed quality case, when $P \in (2\delta_{max}, n\delta_{max})$, consider $i \in \{2, \ldots, n-1\}$ such that $P \in (i\delta_{max}, (i+1)\delta_{max}]$. Any effort profile in which i agents exert effort δ_{max} and $n-i$ agents exert effort 0 is a Nash equilibrium; moreover, if $P = (i+1)\delta_{max}$, any effort profile in which $i+1$ agents exert δ_{max} effort (with $n-i-1$ agents exerting 0 effort) is also a Nash equilibrium. This is an exhaustive characterization of the pure strategy Nash equilibria.*

Proof. From Proposition 4.3, Lemma 6.8, and Proposition 4.4, it suffices to consider strategy profiles where j agents exert effort δ_2 and $n-j$ agents exert effort 0, for any $\delta_2 \leq \frac{P}{j}$ and any $j \geq 2$. If $\delta_2 < \delta_{max}$, consider a possible deviation where an agent decides to exert more effort δ_1, where $\delta_2 \leq \delta_1 \leq \delta_{max}$. Taking the derivative of the best response function, any critical point must satisfy $\delta_1 = \sqrt{P(\delta_2 - \frac{\delta_2^{j-1}}{j})} > \sqrt{\frac{P(j-1)\delta_2}{j}}$. Thus we have a maximum when $\delta_1 > \delta_2$. Note that if $\sqrt{\frac{P(j-1)x}{j}} > \delta_{max}$, the derivative over the interval $\delta_1 \in [\delta_2, \delta_{max}]$ is always positive. Therefore the maximum of the utility function occurs at $\delta_1 = \min(\delta_{max}, \sqrt{P(\delta_2 - \frac{\delta_2^{j-1}}{j})})$ and an agent who exerts effort $\delta_2 < \delta_{max}$ has a profitable deviation.

Therefore, we can restrict attention to strategy profiles where j agents exert effort δ_{max} and and $n-j$ agents exert effort 0. Consider an agent who exerts δ_{max} effort. Her utility is $\frac{P}{j} - \delta_{max}$. In order for this to be an equilibrium, we need $\frac{P}{j} \geq \delta_{max}$, otherwise an agent can deviate and exert 0 effort. Finally, we consider possible deviations, where she exerts lesser, non-zero effort, e.g. $\delta_1 < \delta_{max}$. Her utility becomes $(\frac{\delta_1}{\delta_{max}})^{j-1} \cdot \frac{P}{j} - \delta_1$. We can write this as $\frac{\delta_1}{\delta_{max}}((\frac{\delta_1}{\delta_{max}})^{j-2} \cdot \frac{P}{j} - \delta_{max})$. We note that $(\frac{\delta_1}{\delta_{max}})^{j-2} \cdot \frac{P}{j} - \delta_{max} < \frac{P}{j} - \delta_{max}$, therefore exerting less effort cannot be a profitable deviation. Finally we consider an agent who exerts

0 effort; her expected utility is 0. If an agent deviates and exerts more effort, e.g. $\delta_1 \leq \delta_{max}$, her payoff becomes $(\frac{\delta_1}{\delta_{max}})^i \frac{P}{j+1} - \delta_1$. This quantity is maximized when $\delta_1 = \delta_{max}$. Thus if $\frac{P}{j+1} - \delta_{max} \leq 0$ (or $\frac{P}{j+1} \leq \delta_{max}$), this player does not have a profitable deviation and this is an equilibrium. Otherwise, if $\frac{P}{j+1} - \delta_{max} > 0$, this player has a profitable deviation (in exerting δ_{max}). □

Lemma 6.14. *In the uniformly distributed quality case with $n = 2$, in equilibrium:*

- *When $v > 3$, the principal chooses $P = 2\delta_{max}$, with each agent exerting effort δ_{max}.*
- *When $v < 3$, the principal sets a prize of $P = 0$.*
- *When $v = 3$, the principal is indifferent between all prizes in the interval $P \in [0, 2\delta_{max}]$.*

Proof. We note that the principal need never offer a prize of greater than $2\delta_{max}$, since a prize of $2\delta_{max}$ induces an equilibrium where both players exert effort δ_{max}. For $0 \leq P \leq 2\delta_{max}$, each agent will exert effort $\frac{P}{2}$ and the expected value to the principal of the highest quality good will be $\frac{vP}{3}$. Therefore, the principal's utility is: $\frac{vP}{3} - P = P(\frac{v}{3} - 1)$ for all $0 \leq P \leq 2\delta_{max}$. If $v > 3$, the principal's utility is always positive and maximized when $P = 2\delta_{max}$. When $v = 3$, the principal's utility is 0 for all $P \in [0, 2\delta_{max}]$. When $v < 3$, the principal's utility is negative for all $P > 0$. □

Lemma 6.15.$\forall i \in \{2, \ldots, n-1\}$, *$\exists v \in \Re^+$ such that the principal's expected utility in equilibrium is the same if he chooses the minimal P such i agents participate with full effort, or the minimal P' such that $i + 1$ agents participate with full effort.*

Proof. Given that the principal is strategic, it suffices to consider the minimum value of P to implement a desired outcome. Therefore, the action choices available to the principal are awarding a prize of $i\delta_{max}$ for all $i \in \{2, ..., n\}$. From our equilibrium analysis, we know that when the principal awards a prize of $i\delta_{max}$, it is an equilibrium for i agents to exert full effort. Thus when a principal awards a prize of $i\delta_{max}$, her utility becomes $\frac{i}{i+1} v\delta_{max} - i\delta_{max}$ and when the principal awards a prize of $(i + 1)\delta_{max}$, her utility becomes $\frac{i+1}{i+2} v\delta_{max} - (i + 1)\delta_{max}$. We see that these two expressions are equal when $v = (i + 1)(i + 2)$. □

Lemma 6.16. *In the uniformly distributed quality case, $\forall i \in \{2, \ldots, n-2\}$, $v(i, i+1) < v(i+1, i+2)$.*

Proof. When i agents exert full effort, the expected value to the principal is $\frac{i}{i+1}\delta_{max}v$. The minimum prize value needed to have i agents exert full effort is $i\delta_{max}$, assuming that the principal can pick between the multiple Nash equilibria that arise. Therefore, when i agents exert full effort, the principal's utility function becomes $\frac{i}{i+1}\delta_{max}v - i\delta_{max}$. To determine $v(i, i+1)$, we need $v(i, i+1)$ to satisfy: $\frac{i}{i+1}\delta_{max}v(i, i+1) - i\delta_{max} = \frac{i+1}{i+2}\delta_{max}v_(i, i+1) - (i + 1)\delta_{max}$, or in other words, $v(i, i+1) = (i + 1)(i + 2)$. □

Lemma 6.17. *In the uniformly distributed quality case, if $\exists i \in \{3, ..., n-1\}$ such that $v \in (i(i+1), (i+1)(i+2))$, the principal sets a prize of $i\delta_{max}$ to have i agents exert full effort. Alternatively if $v > n(n+1)$, the principal sets a prize of $n\delta_{max}$ to have n agents exert* *full effort. Finally, if* $\exists i \in \{3, ..., n-1\}$ *such that* $v = i(i+1)$, *the principal is indifferent between setting a prize of* $(i-1)\delta_{max}$ *and* $i\delta_{max}$.

Proof. We first show by strong induction that at any value $i(i+1) < v < (i+1)(i+2)$ for some $2 \leq i \leq n-1$, having i agents exert full effort yields the principal greater utility than having $j < i$ agents exert full effort. We know from from Lemma 6.14, that when $i = 0$, the base case trivially holds for $0 \leq v \leq 3$. Similary, we know from Lemma 6.14, that when $v > 3$, having 2 agents exert full effort yields strictly greater utility than having 0 agents exert full effort and it is not possible to set P so that one agent exert full effort. Hence the base case holds. If the claim holds for $i-1$, we know that for all $v \in ((i-1)i, i(i+1))$, having $i-1$ agents exert full effort yields strictly greater utility to the principal than having $j < i-1$ agents exert full effort. We know that there exists a unique point $v(i, i+1)$ at which the utility to the principal of having i agents exert full effort is the same as the utility to the principal of having $i+1$ agents exert full effort. Moreover, we know that for all $v > v(i, i+1)$, the utility to the principal of having $i+1$ agents exert full effort is strictly greater than the utility to the principal of having i agents exert full effort and that for all $v < v(i, i+1)$ the utility to the principal of having $i+1$ agents exert full effort is strictly less than the utility to the principal of having i agents exert full effort. Thus for any value of $v > v(i, i+1)$, having i agents exert full effort yields strictly greater utility to the principal than having $j < i$ agents exert full effort.

We can construct a similar argument by strong induction that at any value $i(i+1) < v < (i+1)(i+2)$ for some $2 \leq i \leq n-1$ having i agents exert full effort yields the the principal greater utility than having $j > i$ agents exert full effort. We start with $i = n$ as the base case and work our way backwards. Combining these two results, we get that having i agents exert full effort maximizes the principal's utility for any $i(i+1) < v < (i+1)(i+2)$ for some $2 \leq i \leq n-1$. □