Agent Failures in All-Pay Auctions

Yoad Lewenberg[1] yoadlew@cs.huji.ac.il
Omer Lev[1] omerl@cs.huji.ac.il
Yoram Bachrach[2] yobach@microsoft.com
Jeffrey S. Rosenschein[1] jeff@cs.huji.ac.il

[1]The Hebrew University of Jerusalem, Israel
[2]Microsoft Research, Cambridge, United Kingdom

Abstract

All-pay auctions, a common mechanism for various human and agent interactions, suffers, like many other mechanisms, from the possibility of players' failure to participate in the auction. We model such failures and show how they affect the equilibrium state, revealing various properties, such as the lack of influence of the most-likely-to-participate player on the behavior of the other players. We perform this analysis with two scenarios: the sum-profit model, where the auctioneer obtains the sum of all submitted bids, and the max-profit model of crowdsourcing contests, where the auctioneer can only use the best submissions and thus obtains only the winning bid.

Furthermore, we examine various methods of influencing the probability of participation such as the effects of misreporting one's own probability of participating, and how influencing another player's participation chances (e.g., sabotage) changes the player's strategy.

1 Introduction

Auctions have been the focus of much research in economics, mathematics and computer science, and have received attention in the AI and multi-agent communities as a significant tool for resource and task allocation. Beyond explicit auctions, as are performed on the web (e.g., eBay) and in auction houses, auctions also *model* various real-life situations, in which people (and machines) interact and compete for some valuable item. Companies advertising during the U.S. Superbowl are, in effect, bidding to be remembered by the viewer, and are thus putting in tremendous amounts of money in order to create a memorable and unique event for the viewer, overshadowing the other advertisers.

A particularly suitable auction for modeling various scenarios in the real world is the *all-pay auction*. In this type of auction, all participants announce bids, and all of them pay those bids, while only the highest bid wins the product. Candidates applying for a job are, in a sense, participating in such a bidding process, as they put in time and effort preparing for the job interview, while only one of them is selected for the job. This is a "max-profit" auction, as the auctioneer (employer, in this case), receives only the top bid. In comparison, a workplace with an "employee of the month" competition is a "sum-profit" auctioneer, as it enjoys the fruits of all employees' labour, regardless of who won the competition.

The explosion in mass usage of the web has enabled many more all-pay auction-like interactions, including some involving an extremely large number of participants. For example, various crowdsourcing contests, such as the Netflix challenge [Bennett and Lanning, 2007], involve many participants putting in effort, with only the one that performs best winning a prize. Similar efforts can be seen throughout the web, in TopCoder.com, Amazon Mechanical Turk, Bitcoin mining, and other frameworks.

However, despite the research done on all-pay auctions in the past few years [DiPalantino and Vojnović, 2009; Chawla *et al.*, 2012; Lev *et al.*, 2013], some basic questions about all-pay auctions remain — in a full information setting, the only symmetric equilibrium has bidders' expected profit at 0, raising, naturally, the question of why bidders would participate. Allowing bidders to cooperate does not alleviate this problem (at best, it enables a bidder to have a positive expected profit, at the expense of others' negative profit [Lev *et al.*, 2013]).

This paper enriches the all-pay auction model by allowing the possibility of bidders' failure. This means that there is a probability that a bidder will not be able to participate in the auction. As most large-scale all-pay auction mechanisms have variable participation, we believe this helps capture a large family of scenarios, particularly for web-based situations and the uncertainty they contain. We propose a symmetric equilibrium for this situation, and show its various properties. Somewhat surprisingly, allowing failures makes the expected profit for bidders positive, justifying their participation.

We first examine the case where each bidder has a different failure probability, and the potential manipulations possible in this case, such as announcing a false probability (e.g., saying that you will put all your time into a TopCoder.com project) and changing the probability of others (e.g., sabotaging their car). As calculations in this general case are complex, we examine situations where bidders have the same failure probability (as is possible when weather, for example, is the main determinant of participation), enabling us to detail more information about the equilibrium in this state. We note that due to the lengthiness of many of the calculations, we omit most steps, and show the mathematical results directly.

2 Related Work

Initial research on all-pay auction was in the political sciences, modeling lobbying activities [Hillman and Riley, 1989; Baye *et al.*, 1993], but since then, much analysis (especially that dealing with the Revenue Equivalence Theorem) has been done on game-theoretic auction theory [Krishna, 2002; Klemperer, 2004]. When bidders have the same value distribution for the item, [Maskin and Riley, 2003] showed that there is a symmetric equilibrium in auctions where the winner is the bidder with the highest bid. A significant analysis of all-pay auctions in full information settings was [Baye *et al.*, 1996], showing (aided by [Hillman and Riley, 1989]) the equilibrium states in various cases of all-pay auctions, and noting that most valuations (apart from the top two), are not relevant to the winner's strategies.

More recent work has extended the basic model. [Lev *et al.*, 2013] addressed issues of mergers and collusions, while several others directly addressed crowdsourcing models. [DiPalantino and Vojnović, 2009] detailed the issues stemming from needing to choose one auction from several, and [Chawla *et al.*, 2012] dealt with optimal mechanisms for crowdsourcing. Using both theoretical and empirical tools, [Gao *et al.*, 2012] examined whether several stages were better for crowdsourcing, while [Archak and Sundararajan, 2009] addressed issues related to designing the award of the crowdsourcing contest.

The early major work on failures in auctions was [McAfee and McMillan, 1987], followed soon after by [Matthews, 1987], which introduced bidders who are not certain of how many bidders there will actually be at the auction. Their analysis showed that in first-price auctions (like our all-pay auction), risk averse bidders prefer to know the numbers, while it is the auctioneer's interest to hide that number. In the case of neutral bidders (such as ours), their model claimed that bidders were unaffected by the numerical knowledge. [Dyer *et al.*, 1989] claimed that experiments that allowed "contingent" bids (i.e., one submits several bids, depending on the number of actual participants) supported these results. [Menezes and Monteiro, 2000] presented a model where auction participants know the maximal number of bidders, but not how many will ultimately participate. However, the decision in their case was endogenous to the bidder, and therefore a reserve price has a significant effect in their model (though ultimately without change in expected revenue, in comparison to full-knowledge models). In contrast to that, our model,[1] which assumes a little more information is available to the bidders (they know the maximal number of bidders, and the probability of failure), finds that in such a scenario, bidders are better off not having everyone show up, rather than knowing the real number of contestants appearing. Empirical work done on actual auctions [Lu and Yang, 2003] seems to support some of our theoretical findings (though not specifically in all-pay auction settings).

[1]We use a framework similar to the one in [Meir *et al.*, 2012], albeit there it was used in congestion games.

3 Model

We consider an all-pay auction with a single auctioned item that is commonly valued by all the participants. This is a restricted case of the model in [Baye *et al.*, 1996], where players' item valuation could be different.

Formally, we assume that each of the n bidders issues a bid of b_i, $i = 1, \dots, n$, and all bidders value the item at 1. The highest bidders win the item and divide it among themselves, while the rest lose their bid. Thus, bidder i's utility from a combination of bids $(b_1, \dots, b_n)$ is given by:

$$\pi_i(b_1,\dots,b_n) = \begin{cases} \frac{1}{|\arg\max_j b_j|} - b_i & i \in \arg\max_j b_j \\ -b_i & i \notin \arg\max_j b_j \end{cases}$$

We are interested in a symmetric equilibrium, which in this case, without possibility of failure, is unique [Baye *et al.*, 1996; Maskin and Riley, 2003]. It is a mixed equilibrium with full support of $[0,1]$, so that each bidder's bid is distributed in $[0,1]$ according to the same cumulative distribution function F, with the density function f (since it is non-atomic, tie-breaking is not an issue). As we compare this case to that of no-failures, this is a case similar to that presented in [Baye *et al.*, 1996; Gao *et al.*, 2012; DiPalantino and Vojnović, 2009], where various results on the behavior of non-cooperative bidders have been provided. We briefly give an overview of the results without failures in Subsection 3.1.

When we allow bidders to fail, we assume that each of them has a probability of participating — $p_i \in [0,1]$. As a matter of convenience, we shall order the bidders according to their probabilities, so $0 \le p_1 \le p_2 \le \dots \le p_n \le 1$. If a bidder fails to participate, its utility is 0.[2]

3.1 Auctions without Failures

The expected utility of any participant with a bid b is:

$$\pi(b) = (1-b)\cdot Pr(winning|b) + (-b)\cdot Pr(losing|b)$$

where $Pr(winning|b)$ and $Pr(losing|b)$ are the probabilities of winning or losing the item when bidding b, respectively. In a symmetric equilibrium with n players, each of the bidders chooses his bid from a single bid distribution with a probability density function $f_n(x)$ and a cumulative distribution function $F_n(x)$. A player who bids b can only win if all the other $n-1$ players bid at most b, which occurs with probability $F_n^{n-1}(b)$. Thus, the expected utility of a player bidding b is given by:

$$\pi(b) = (1-b)\, F_n^{n-1}(b) - b\left(1 - F_n^{n-1}(b)\right) = F_n^{n-1}(b) - b$$

In a mixed Nash equilibrium, all points in the support yield the same expected utility to a player; thus, as the equilibrium has full support, $\pi(0) = \pi(x)$ for all $x \in [0,1]$. Since $\pi(0) = 0$, this means that for all bids, $F_n^{n-1}(b) = b$. Thus,

[2]The calculations utilize an average probability, so if every bidder has a distribution function on its probability to fail, these results still stand.

Variable	No Failures
Expected bid [Variance]	$\frac{1}{n}$ $\left[\frac{1}{2n-1}-\frac{1}{n^2}\right]$
Bidder utility [Variance]	0 $\left[\frac{n-1}{n(2n-1)}\right]$
Sum-profit principal utility [Variance]	1 $\left[\frac{n}{2n-1}-\frac{1}{n}\right]$
Max-profit principal utility [Variance]	$\frac{n}{2n-1}$ $\left[\frac{n(n-1)^2}{(3n-2)(2n-1)^2}\right]$

Table 1: The values, in expectation, of some of the variables when there is no possibility of failure

$\left(\int_0^b f_n(x)\,\mathrm{d}x\right)^{n-1} = b$, so $F_n(x) = x^{\frac{1}{n-1}}$ and $f_n(x) = \frac{x^{\frac{2-n}{n-1}}}{n-1}$. The various properties of an auction without failures can be found in Table 1.

4 Every Bidder with Own Failure Probability

We assume that each bidder has its own probability of participating in the auction, with $0 \le p_1 \le \ldots \le p_n \le 1$. We shall now present a symmetric equilibrium for this case, with a positive expected utility for the bidders.

We begin by defining a few helpful functions. First, we define $\lambda = \prod_{j=1}^{n-1}(1-p_j)$, and we define the following expressions for all $1 \le k \le n-1$:

$$H_k(x) = \begin{cases} \left(\frac{\lambda+x}{\prod_{j=1}^{k-1}(1-p_j)}\right)^{\frac{1}{n-k}} & k>1 \\ (\lambda+x)^{\frac{1}{n-1}} & k=1 \end{cases}$$

$$\underline{s}_k = \begin{cases} (1-p_k)^{n-k}\prod_{j=1}^{k-1}(1-p_j)-\lambda & k>1 \\ (1-p_1)^{n-1}-\lambda & k=1 \end{cases}$$

For the virtual "0" index, we use $\underline{s}_0 = 1-\lambda$. Note that because the p_is are ordered, so are the $\underline{s}_i$s: $1 \ge \underline{s}_0 \ge \underline{s}_1 \ge \ldots \ge \underline{s}_{n-1} = 0$.

4.1 Equilibrium

We are now ready to define the c.d.f.s for our equilibrium, for every player $1 \le i \le n-1$:

$$F_i(x) = \begin{cases} 1 & x \ge \underline{s}_0 \\ \frac{H_1(x)+p_i-1}{p_i} & x \in [\underline{s}_1, \underline{s}_0) \\ \vdots & \vdots \\ \frac{H_k(x)+p_i-1}{p_i} & x \in [\underline{s}_k, \underline{s}_{k-1}) \\ \vdots & \vdots \\ \frac{H_i(x)+p_i-1}{p_i} & x \in [\underline{s}_i, \underline{s}_{i-1}) \\ 0 & x < \underline{s}_i \end{cases}$$

F_n, uniquely, while it is very similar to F_{n-1} in its piecewise composition, has an atomic point in the distribution at 0 of $1-\frac{p_{n-1}}{p_n}$, so:

$$F_n(x) = \begin{cases} 1 & x \ge \underline{s}_0 \\ \frac{H_1(x)+p_n-1}{p_n} & x \in [\underline{s}_1, \underline{s}_0) \\ \vdots & \vdots \\ \frac{H_k(x)+p_n-1}{p_n} & x \in [\underline{s}_k, \underline{s}_{k-1}) \\ \vdots & \vdots \\ \frac{H_{n-1}(x)+p_n-1}{p_n} & x \in (\underline{s}_{n-1}, \underline{s}_{n-2}) \\ 1-\frac{p_{n-1}}{p_n} & x=0 \\ 0 & x<0 \end{cases}$$

Note that all c.d.f.s are continuous and piecewise differentiable,[3] and when $p_i = p_j$ it follows that $F_i = F_j$; therefore, this is a symmetric equilibrium. In the course of proving this is indeed an equilibrium, we shall calculate the expected utility of the bidders when they participate.

The logic behind this equilibrium is that bidders that participate rarely will usually bid high, while those that frequently participate in auctions with less competition would more commonly bid low.

When bidder i bids according to this distribution, i.e., $x \in [\underline{s}_k, \underline{s}_{k-1})$ for $1 \le k \le i$:

$$\begin{aligned} \pi_i(x) &= (1-x)\prod_{j=1;j\neq i}^{n}(p_jF_j(x)+1-p_j) - \\ &\quad - x\left(1-\prod_{j=1;j\neq i}^{n}(p_jF_j(x)+1-p_j)\right) = \\ &= \prod_{j=1;j\neq i}^{n}(p_jF_j(x)+1-p_j) - x = \\ &= \prod_{j=1}^{k-1}(1-p_j)\prod_{j=k;j\neq i}^{n}H_k(x) - x = \\ &= \prod_{j=1}^{k-1}(1-p_j)H_k^{n-k}(x) - x = \\ &= \prod_{j=1}^{k-1}(1-p_j)\frac{\lambda+x}{\prod_{j=1}^{k-1}(1-p_j)} - x = \\ &= \lambda \end{aligned}$$

If bidder i bids outside his support, i.e., $x \in [\underline{s}_k, \underline{s}_{k-1})$ for $i+1 \le k \le n-1$, the same equation becomes:

[3] When $\prod_{j=1}^{k-1}(1-p_j) = 0$, and H_k is undefined for some k, then there is no range for which that H_k is used.

$$\pi_i(x) = \prod_{j=1;j\neq i}^{k-1} (1-p_j) \prod_{j=k}^{n} H_k(x) - x =$$

$$= \prod_{j=1;j\neq i}^{k-1} (1-p_j) \left(\frac{\lambda + x}{\prod_{j=1}^{k-1}(1-p_j)}\right)^{\frac{n-k+1}{n-k}} - x =$$

$$= \frac{\lambda + x}{1-p_i} \left(\frac{\lambda + x}{\prod_{j=1}^{k-1}(1-p_j)}\right)^{\frac{1}{n-k}} - x$$

Since $x \in [\underline{s}_k, \underline{s}_{k-1})$, then $x < \underline{s}_{k-1} = (1-p_{k-1})^{n-k+1} \prod_{j=1}^{k-2}(1-p_j) - \lambda = (1-p_{k-1})^{n-k} \prod_{j=1}^{k-1}(1-p_j) - \lambda$, and hence $\lambda + x < (1-p_{k-1})^{n-k} \prod_{j=1}^{k-1}(1-p_j)$. Therefore:

$$\pi_i(x) < \frac{\lambda + x}{1-p_i} \left(\frac{(1-p_{k-1})^{n-k} \prod_{j=1}^{k-1}(1-p_j)}{\prod_{j=1}^{k-1}(1-p_j)}\right)^{\frac{1}{n-k}} - x =$$

$$= \frac{\lambda + x}{1-p_i}(1-p_{k-1}) - x = (\lambda + x)\frac{p_i - p_{k-1}}{1-p_i} + \lambda$$

Finally, as $i+1 \leq k$, $p_i \leq p_{k-1}$, hence $p_i - p_{k-1} \leq 0$, and $\pi_i(x) < \lambda$.

4.2 Profits

When a bidder actually participates his expected utility is λ, and therefore the overall expected utility for bidder i is $p_i\lambda$ (which, naturally, decreases with n). Notice that, as is to be expected, a bidder's profit rises the less reliable his fellow bidders are, or the fewer participants the auction has. However, the most reliable of the bidders does not affect the profits of the rest. If a bidder can set his own participation rate, if there is no bidder with $p_j = 1$, that is the best strategy; otherwise, his optimal probability should be $\frac{1}{2}$, as that maximizes $p_i(1-p_i)\prod_{j=1;j\neq i}^{n-1}(1-p_j)$.

In order to calculate the auctioneer's profit in a sum-profit model, we need to calculate the expected bid by each bidder, and for that we need to calculate the bidders' p.d.f. For $1 \leq i \leq n-1$:

$$f_i(x) = \begin{cases} 0 & x \geq \underline{s}_0 \\ \frac{(\lambda+x)^{\frac{2-n}{n-1}}}{p_i(n-1)} & x \in [\underline{s}_1, \underline{s}_0) \\ \vdots & \vdots \\ \frac{(\lambda+x)^{\frac{k+1-n}{n-k}}}{p_i(n-k)(\prod_{j=1}^{k-1}(1-p_j))^{\frac{1}{n-k}}} & x \in [\underline{s}_k, \underline{s}_{k-1}) \\ \vdots & \vdots \\ \frac{(\lambda+x)^{\frac{i+1-n}{n-i}}}{p_i(n-i)(\prod_{j=1}^{i-1}(1-p_j))^{\frac{1}{n-i}}} & x \in [\underline{s}_i, \underline{s}_{i-1}) \\ 0 & x < \underline{s}_i \end{cases}$$

and $f_n(x) = \frac{p_{n-1}}{p_n} f_{n-1}(x)$.

The expected bid by each bidder, for $1 \leq i \leq n-1$:

$$\mathbb{E}[bid_i] = \sum_{k=1}^{i} \int_{\underline{s}_k}^{\underline{s}_{k-1}} x f_i(x)\,\mathrm{d}x =$$

$$= \frac{1}{p_i}\left(\frac{1}{n} + \sum_{k=1}^{i} \frac{(1-p_k)^{n-k}\prod_{j=1}^{k}(1-p_j)}{(n-k)(n-k+1)} - \right.$$
$$\left. - \frac{(1-p_i)^{n-i}\prod_{j=1}^{i}(1-p_j)}{n-i} - p_i\lambda\right)$$

and $\mathbb{E}[bid_n] = \frac{p_{n-1}}{p_n}\mathbb{E}[bid_{n-1}]$.

This expression decreases with n, indicating, as in the no-failure model, that as more bidders participate, the chance of losing increases, causing bidders to lower their exposure. Surprisingly, when summing over all bidders, we receive a much simpler expression, and the sum-profit auctioneer's profits are:

$$\sum_{i=1}^{n} p_i \mathbb{E}[bid_i] = 1 - \lambda\left(1 + \sum_{i=1}^{n-1} p_i\right)$$

In this case, growth with n is monotonic, and hence, any addition to n is a net positive for the sum-profit auctioneer.

To calculate a max-profit auctioneer's profits, we need to first define the auctioneer's c.d.f.:

$$G(x) = \prod_{i=1}^{n} (p_i F_i(x) + 1 - p_i)$$

That is,

$$G(x) = \begin{cases} 1 & x \geq \underline{s}_0 \\ (\lambda + x)^{\frac{n}{n-1}} & x \in [\underline{s}_1, \underline{s}_0) \\ \vdots & \vdots \\ \frac{(\lambda+x)^{\frac{n-k+1}{n-k}}}{\left(\prod_{j=1}^{k-1}(1-p_j)\right)^{\frac{1}{n-k}}} & x \in [\underline{s}_k, \underline{s}_{k-1}) \\ \vdots & \vdots \\ \frac{(\lambda+x)^2}{\prod_{j=1}^{n-2}(1-p_j)} & x \in [\underline{s}_{n-1}, \underline{s}_{n-2}) \\ 0 & x < 0 \end{cases}$$

This is differentiable, and hence we can find $g(x) = \frac{\mathrm{d}}{\mathrm{d}x}G(x)$; looking for the expected profit, we have:

$$\int_{\underline{s}_{n-1}}^{\underline{s}_0} x g(x)\,\mathrm{d}x = \frac{n}{2n-1} - \lambda +$$
$$+ \sum_{k=1}^{n-1}\left(\frac{(1-p_k)^{2n-2k-1}\prod_{j=1}^{k}(1-p_j)^2}{4(n-k)^2 - 1}\right)$$

This means that the max-profit auctioneer would prefer to have two reliable players ($p_n = p_{n-1} = 1$), and the other $n-2$ bidders as unreliable as possible.

Example 4.1. *Consider how four bidders interact. Our bidders have participation probability of* $p_1 = \frac{1}{3}$, $p_2 = \frac{1}{2}$, $p_3 = \frac{3}{4}$ *and* $p_4 = 1$. *Let us look at each bidder's c.d.f.s:*

$$F_1(x) = \begin{cases} 1 & x \geq \frac{11}{12} \\ 3\left(\frac{1}{12}+x\right)^{\frac{1}{3}} - 2 & x \in \left[\frac{23}{108}, \frac{11}{12}\right) \\ 0 & x < \frac{23}{108} \end{cases}$$

$$F_2(x) = \begin{cases} 1 & x \geq \frac{11}{12} \\ 2\left(\frac{1}{12}+x\right)^{\frac{1}{3}} - 1 & x \in \left[\frac{23}{108}, \frac{11}{12}\right) \\ 2\left(\frac{3\left(\frac{1}{12}+x\right)}{2}\right)^{\frac{1}{2}} - 1 & x \in \left[\frac{1}{12}, \frac{23}{108}\right) \\ 0 & x < \frac{1}{12} \end{cases}$$

$$F_3(x) = \begin{cases} 1 & x \geq \frac{11}{12} \\ \frac{4}{3}\left(\frac{1}{12}+x\right)^{\frac{1}{3}} - \frac{1}{3} & x \in \left[\frac{23}{108}, \frac{11}{12}\right) \\ \frac{4}{3}\left(\frac{3\left(\frac{1}{12}+x\right)}{2}\right)^{\frac{1}{2}} - \frac{1}{3} & x \in \left[\frac{1}{12}, \frac{23}{108}\right) \\ 4\left(\frac{1}{12}+x\right) - \frac{1}{3} & x \in \left[0, \frac{1}{12}\right) \\ 0 & x < 0 \end{cases}$$

$$F_4(x) = \begin{cases} 1 & x \geq \frac{11}{12} \\ \left(\frac{1}{12}+x\right)^{\frac{1}{3}} & x \in \left[\frac{23}{108}, \frac{11}{12}\right) \\ \left(\frac{3\left(\frac{1}{12}+x\right)}{2}\right)^{\frac{1}{2}} & x \in \left[\frac{1}{12}, \frac{23}{108}\right) \\ 3\left(\frac{1}{12}+x\right) & x \in \left(0, \frac{1}{12}\right) \\ \frac{1}{4} & x = 0 \\ 0 & x < 0 \end{cases}$$

The expected utility for bidder 1 is $\frac{1}{36}$, *for expected bid of* $\frac{14}{27}$; *for bidder 2,* $\frac{1}{24}$ *for expected bid of* 0.394; *for bidder 3,* $\frac{1}{16}$ *for expected bid of* 0.277; *and for the last bidder,* $\frac{1}{12}$ *for expected bid of* 0.207.

A sum-profit auctioneer will see an expected profit of $\frac{113}{144}$, *while a max-profit one will get, in expectation,* 0.490.

As a comparison, in the case where we do not allow failures, the c.d.f. of the bidders is $x^{\frac{1}{3}}$ *with expected bid of* $\frac{1}{4}$ *and expected utility of* 0. *The expected profit of the sum-profit auctioneer is* 1, *while the expected profit of the max-profit auctioneer is* $\frac{4}{7}$.

5 False Identity and Sabotage

Suppose our bidder can influence others' perceptions, and create a false sense of its participation probability. What would its best strategy be, and how should the participation probability be altered? Any bid beyond $\underline{s}_0$ is sure to win, but as that would give profit of less than λ, which is less than the expected profit for non-manipulators, it is not worthwhile. So our bidder will bid in its support, with the expected profit being λ. Thus, our bidder will strive to increase λ, and would do so by trying to portray its participation probability as being as low as possible, thus lulling the other bidders with a false sense of security. Of course, this reduces the payment to auctioneers of any type, and therefore, they would try to expose such manipulation.

More interesting is the possibility of a player's changing another player's participation probability by using sabotage; thus our bidder would be the only bidder knowing the real participation probability. Our bidder, i, sabotages bidder r, with a perceived participation probability of p_r, changing his real participation probability to p'_r. Bidder i's expected profit with bid x is:

$$\pi_i(x) = \prod_{j=1;j\neq i,r}^{n} \left(p_j F_j(x) + 1 - p_j\right)\left(p'_r F_r(x) + 1 - p'_r\right)$$

The values of this function change according to the relation between r, i and x. To find the optimal strategy for a player, we will examine all the options. When x is in bidder i's support and not in bidder r's support, i.e., there is a $k < i$ for which $x \in [\underline{s}_k, \underline{s}_{k-1})$ and $r < k$:

$$\pi_i(x) = \frac{1-p'_r}{1-p_r}(\lambda + x) - x = \frac{p_r - p'_r}{1-p_r}(\lambda + x) + \lambda$$

In this case, this is larger than λ only if $p_r > p'_r$, and grows with the bid, though the maximal bid (due to the fact that $r < k$) is $\underline{s}_r$.

If $k < r$ (i.e., k is in i and r's support):

$$\pi_i(x) = \frac{p_r - p'_r}{p_r}(\lambda + x)\left(\left(\frac{\prod_{j=1}^{k-1}(1-p_j)}{\lambda + x}\right)^{\frac{1}{n-k}} - 1\right) + \lambda$$

Since $x < \underline{s}_{k-1}$, this means that $\lambda + x < (1 - p_{k-1})^{n-k}\prod_{j=1}^{k-1}(1-p_j)$, hence $\frac{\prod_{j=1}^{k-1}(1-p_j)}{\lambda+x}^{\frac{1}{n-k}} > 1$; again, if $p_r > p'_r$, this is an increase over λ — the position without sabotaging.

Due to space constraints we do not detail here the calculations for the other cases for i and r, as they collapse to the above equations. The saboteur's optimal bid is dependent on the specific makeup of the auction, and hence, so is the value of the sabotage action. For $x \in \left[\underline{s}_k, \underline{s}_{k-1}\right]$, if $\frac{1}{n-k}$ is in the range (p_k, p_{k-1}), the optimal bid is $\left(1 - \frac{1}{n-k}\right)^{n-k}\prod_{j=1}^{k-1}(1-p_j) - \lambda$, while if $\frac{1}{n-k} \leq p_{k-1}$, the optimal bid is $\underline{s}_{k-1}$, and if $\frac{1}{n-k} \geq p_k$, $\underline{s}_k$ is optimal. However, one must go through every range in player i's support to decide what is the optimal bid for a particular auction. Note that the same equations are true for every bidder that is not r (i.e., not just the manipulator i). Since at every point in i's support, it will be better than before the manipulation, that is true also for the other bidders; therefore, the manipulation has benefited not only the saboteur (though it will be able to optimize its strategy towards it), but also all other bidders other than bidder r.

6 Uniform Failure Probabilities

If we allow our bidders to have the same probability of failure (e.g., when failures stem from weather conditions), many of the calculations become more tractable, and we are able to further understand the scenario.

Variable	No failures	Uniform participation probability	Individual participation probability
Expected bid	$\frac{1}{n}$	$\frac{1}{np}\left(1-(1-p)^{n-1}\left(1+p\left(n-1\right)\right)\right)$	$\frac{1}{p_i}\left(\frac{1}{n}+\sum_{k=1}^{i}\frac{(1-p_k)^{n-k}\prod_{j=1}^{k}(1-p_j)}{(n-k)(n-k+1)}-\right.$ $\left.-\frac{(1-p_i)^{n-i}\prod_{j=1}^{i}(1-p_j)}{n-i}-p_i\lambda\right)$, $\mathbb{E}\left[bid_n\right]=\frac{p_{n-1}}{p_n}\mathbb{E}\left[bid_{n-1}\right]$
Bidder utility	0	$p(1-p)^{n-1}$	$p_i\prod_{j=1}^{n-1}(1-p_j)$
Sum-profit principal utility	1	$1-(1-p)^{n-1}\left(1+p\left(n-1\right)\right)$	$1-\lambda\left(1+\sum_{i=1}^{n-1}p_i\right)$
Max-profit principal utility	$\frac{n}{2n-1}$	$\frac{n}{2n-1}+\frac{n-1}{2n-1}\left(1-p\right)^{2n-1}-\left(1-p\right)^{n-1}$	$\frac{n}{2n-1}-\lambda+\sum_{k=1}^{n-1}\left(\frac{(1-p_k)^{2n-2k-1}\prod_{j=1}^{k}(1-p_j)^2}{4(n-k)^2-1}\right)$

Table 2: The values, in expectation, of some of the variables in a no-failure setting, when all bidders have the same participation probability, and when each member has their own participation probability.

6.1 Profit Variance

As this case is a particular instance of the general case presented above, we know the expected utilities. However, the simplification of the identical probabilities allows us to examine the variance as well. For each bidder, the expected utility is $p(1-p)^{n-1}$, monotonically decreasing in n. Using the c.d.f. calculated in the general case, we can also calculate the expected utility squared, and we use it to calculate the utility variation:

$$\frac{n-1}{n\left(2n-1\right)}-\frac{\left(1-p\right)^{n}}{n}+\left(p+\frac{1}{2n-1}\right)\left(1-p\right)^{2n-1}$$

The variance increases with p.

The expected bid is $\frac{1}{np}\left(1-(1-p)^{n-1}\left(1+p\left(n-1\right)\right)\right)$, which is neither monotonic in n nor in p. Hence, the expected profit of the *sum-profit auctioneer* is $1-(1-p)^{n-1}\left(1+p\left(n-1\right)\right)$ which is monotonically increasing in p and in n. The profit variance is $\frac{np\left(1-(1-p)^{2n-1}\right)}{2n-1}-\frac{\left(1-(1-p)^n\right)^2}{n}$. Note that as n grows, the auctioneer's expected revenue approaches that of the no-failure case.

In the case of the *max-profit auctioneer*, the expected profit is $\frac{n}{2n-1}+\frac{n-1}{2n-1}\left(1-p\right)^{2n-1}-\left(1-p\right)^{n-1}$, which is monotonically increasing in p; while not monotonic in n, for large enough n it approaches the expected revenue in the no-failure case. The variance is $\left(1-p\right)^{2n-2}-\frac{2n(1-p)^{n-1}}{2n-1}+\frac{n}{3n-2}-\frac{2(n-1)^2(1-p)^{3n-2}}{(3n-2)(2n-1)}-\left(\frac{n}{2n-1}+\frac{n-1}{2n-1}\left(1-p\right)^{2n-1}-\left(1-p\right)^{n-1}\right)^2$.

7 Conclusion and Discussion

Bidders failing to participate in auctions happen commonly, as people choose to apply to one job but not another, or to participate in the Netflix challenge but not a similar challenge offered by a competitor. Examining these scenarios enables us to understand certain fundamental issues in all-pay auctions. In the pure information, classical versions, bidders each have an expected revenue of 0; in a limited information scenario, such as the one we dealt with, bidders have positive expected revenue, and are strongly incentivized to participate in the auction. Auctioneers, on the other hand, mostly lose their strong control of the auction, and no longer pocket almost all revenues involved in the auction. However, by influencing participation probabilities, max-profit auctioneers can effectively increase their revenue in comparison to the no-failure model. A short summary of our results appears in Table 2.

The basic idea of the equilibrium we explored here was that frequent participants could allow themselves to bid lower, as there would be plenty of contests where they would be one of the few participants, and hence win with smaller bids. Infrequent bidders, on the other hand, would wish to maximize the few times they participate, and hence bid fairly high bids. As exists in the no-failure case as well, as more and more participants join, there is a concentration of bids at lower price points, as bidders are more afraid of the fierce competition. Hence, it is fairly easy to see in all of our results that as n approached larger numbers, the various variables were closer and closer to their no-failure brethren.

There is still much left to explore in these models — not only more techniques of manipulation by bidders and potential incentives by auctioneers, but also further enrichment of the model. Currently, participation rates are not influenced by other bidders' probability of participation, but, obviously, many scenarios in real-life have, effectively, a feedback loop in this regard. Finding a suitable model for such interactions, while an ambitious goal, might help us gain even further insight into these types of interactions.

Acknowledgments

This research was supported in part by Israel Science Foundation grant #1227/12, the Google Inter-University Center for Electronic Markets and Auctions, and the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI).

References

[Archak and Sundararajan, 2009] Nikolay Archak and Arun Sundararajan. Optimal design of crowdsourcing contests. In *Proceedings of the Thirtieth International Conference on Information Systems*, ICIS, Phoenix, Arizona, 2009.

[Baye *et al.*, 1993] Michael R. Baye, Dan Kovenock, and Casper G. de Vries. Rigging the lobbying process: An application of the all-pay auction. *The American Economic Review*, 83(1):289–294, March 1993.

[Baye *et al.*, 1996] Michael R. Baye, Dan Kovenock, and Casper G. de Vries. The all-pay auction with complete information. *Economic Theory*, 8(2):291–305, August 1996.

[Bennett and Lanning, 2007] James Bennett and Stan Lanning. The Netflix prize. In *Proceedings of KDD Cup and Workshop*, 2007.

[Chawla *et al.*, 2012] Shuchi Chawla, Jason D. Hartline, and Balasubramanian Sivan. Optimal crowdsourcing contests. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 856–868, Kyoto, Japan, 2012. SIAM.

[DiPalantino and Vojnović, 2009] Dominic DiPalantino and Milan Vojnović. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 119–128, Stanford, California, July 2009.

[Dyer *et al.*, 1989] Douglas Dyer, John H. Kagel, and Dan Levin. Resolving uncertainty about the number of bidders in independent private-value auctions: An experimental analysis. *The RAND Journal of Economics*, 20(2):268–279, 1989.

[Gao *et al.*, 2012] Xi Alice Gao, Yoram Bachrach, Peter Key, and Thore Graepel. Quality expectation-variance tradeoffs in crowdsourcing contests. In *Proceedings of the 26th National Conference on Artificial Intelligence (AAAI)*, Toronto, Canada, June 2012.

[Hillman and Riley, 1989] Arye L. Hillman and John G. Riley. Politically contestable rents and transfers. *Economics & Politics*, 1(1):17–39, March 1989.

[Klemperer, 2004] Paul Klemperer. *Auctions: Theory and Practice*. The Toulouse Lectures in Economics. Princeton University Press, 2004.

[Krishna, 2002] Vijay Krishna. *Auction Theory*. Academic Press, 2002.

[Lev *et al.*, 2013] Omer Lev, Maria Polukarov, Yoram Bachrach, and Jeffrey S. Rosenschein. Mergers and collusion in all-pay auctions and crowdsourcing contests. In *Proceedings of the 12th International Coference on Autonomous Agents and Multiagent Systems (AAMAS)*, St. Paul, Minnesota, May 2013. To appear.

[Lu and Yang, 2003] Dennis Lu and Jing Yang. Auction participation and market uncertainty: Evidence from canadian treasury auctions. In *Conference paper presented at the Canadian Economics Association 37th Annual Meeting*, 2003.

[Maskin and Riley, 2003] Eric Maskin and John Riley. Uniqueness of equilibrium in sealed high-bid auctions. *Games and Economic Behavior*, 45(2):395–409, November 2003.

[Matthews, 1987] Steven Matthews. Comparing auctions for risk averse buyers: A buyer's point of view. *Econometrica*, 55(3):633–646, May 1987.

[McAfee and McMillan, 1987] Randolph Preston McAfee and John McMillan. Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43(1):1–19, October 1987.

[Meir *et al.*, 2012] Reshef Meir, Moshe Tennenholtz, Yoram Bachrach, and Peter Key. Congestion games with agent failures. In *Proceedings of the 26th National Conference on Artificial Intelligence (AAAI)*, pages 1401–1407, Toronto, Canada, July 2012.

[Menezes and Monteiro, 2000] Flavio M. Menezes and Paulo K. Monteiro. Auctions with endogenous participation. *Review of Economic Design*, 5(1):71–89, March 2000.