Bid Caps in Contests

Wojciech Olszewski and Ron Siegel*

April 2016

Abstract

We study the effect of bid caps on contestants' aggregate equilibrium bids in variants of all-pay contests with a large number of heterogeneous contestants and prizes.

We show that when contestants' bidding costs are linear (all-pay auctions) or concave, a bid cap always decreases revenue. With convex costs, in contrast, bid caps may increase revenue. We also show that a flexible bid cap decreases revenue regardless of the curvature of the bidding costs, even when these costs vary across contestants.

Our findings contrast with the results for two-player contests, obtained by Che and Gale (1998, 2006), and support the results obtained by Gavious et al. (2002) and Kaplan and Wettstein (2006). We explain intuitively which features of contests drive these (in)consistencies.

1 Introduction

A bid cap is a design tool in contests, which limits contestants' bids, effort, or investments. In the context of political lobbying, Che and Gale (1998) showed, perhaps surprisingly, that bid caps can increase aggregate equilibrium bids. Several other authors argued, in various settings, that bid caps may or may not increase equilibrium aggregate bids. More recently,

*Dept. of Economics, Northwestern University, Evanston, IL 60208 (e-mail: wo@northwestern.edu) and Dept. of Economics, The Pennsylvania State University, University Park, PA 16801 (e-mail: rus41@psu.edu).
Financial support from the NSF (grant SES-1325968) is gratefully acknowledged.

Olszewski and Siegel (2016b) observed that in large contests, such as the competition for college admissions or NSF grants, imposing a bid cap (interpreted in their model as an effort cap) may result in a Pareto-improving outcome.

While this literature has delivered several important and intriguing results, the overall picture it provides regarding the effect of bid caps in contests is somewhat blurred. For example, it is not obvious why bid caps that are not too small always weakly, and sometime strictly, increase aggregate bids in Che and Gale's (1998) all-pay auction with complete information, but bid caps of any magnitude always weakly, and sometimes strictly, decrease aggregate bids in Gavious et al.'s (2002) all-pay auctions with incomplete information. One would like to know whether this difference is due to the different informational assumptions and, more generally, which features of a contest make it possible to increase revenue or effort by imposing a bid cap.

Similar questions arise in the context of flexible bid caps, that is, caps above which players can bid at an additional (and possibly very high) cost. Kaplan and Wettstein (2006) showed in a two-player model, in which the players have the same cost of bidding, that such caps always decrease aggregate bids, but Che and Gale (2006) argued that this is not always true when costs are heterogeneous. One would like to know whether this difference is due to the different assumptions on costs and, more generally, to identify settings in which flexible bid caps may increase revenue.

These questions are difficult or impossible to answer by analyzing all-pay contests directly, since such an analysis is often intractable. We circumvent a direct analysis by employing the mechanism-design approach to studying the equilibria of large contests, developed in Olszewski and Siegel (2016a). This allows us to conduct the analysis in a tractable mechanism-design framework, whose qualitative results apply to contests with sufficiently many players.

We show that in large all-pay auctions, in which players have linear costs, with complete or incomplete information about their prize valuation, no bid cap can increase aggregate equilibrium bids. This result generalizes to large contests in which players have concave costs. In contrast, bid caps can increase revenue in large contests in which players have convex costs. We provide some sufficient conditions under which this happens, and show that bid caps can not only increase aggregate bids, but also increase the expected utility of each contestant. These results imply that incomplete information is not necessary for bid caps to be an ineffective tool for increasing aggregate bids. We also show that in large contests flexible bid caps are always an ineffective tool for increasing revenue, regardless of the assumptions on contestants' information and costs.

The departure of these results from those of Che and Gale (1998, 2006) arises because in a small contest with asymmetric contestants different contestants face different competition in equilibrium. More precisely, each contestant faces the equilibrium bidding strategies of the other contestants, and these strategies differ across contestants. This difference allows even a small decrease in a bid cap to have a substantial effect on the aggregate bids, and allows a flexible cap to substantially affect the equilibrium allocation of the prize, thereby affecting aggregate bids. But when there are many contestants, each faces almost the same competition as the others, because they each face almost the same set of competitors. Consequently, a small change in a bid cap leads to a small change in aggregate bids, and a flexible cap cannot substantially change the equilibrium prize allocation, and consequently the aggregate bids. This intuition also hints at why the results of Gavious et al. (2002) are in line with ours. They assumed that players are ex-ante symmetric, and studied symmetric equilibria. In such equilibria all players face the same competition, very similarly to what happens in the equilibria of sufficiently large contests with possibly asymmetric players.

Related Literature

Che and Gale (1998) considered an all-pay auction with complete information and two players with different valuations, and identified a range of bid caps that lead to higher aggregate bids. Szech (2015) showed that changing the tie-breaking rule can further increase revenue.

Kaplan and Wettstein (2006) argued that in many settings bid caps are not rigid, that is, players can bid above the cap at an additional cost. They studied flexible caps, modelled as an increase in the cost of each bid, and showed that in the setting of Che and Gale (1998) a flexible cap always decreases the aggregate bids. Their observation that the effects of rigid and flexible caps can be very different is puzzling, since the latter can approximate the former. In response to Kaplan and Wettstein's (2006) finding, Che and Gale (2006) showed that if different players face different costs of the same bid, then a flexible cap can increase players' aggregate *costs*.

Gavious et al. (2002) studied the symmetric equilibrium of contests with ex-ante symmetric players and incomplete information, and showed that imposing a (rigid) cap may increase aggregate bids if: (a) sufficiently many players compete for a single prize and their bidding cost function is convex; or (b) any number of players compete for a single prize and the bidding cost function is sufficiently convex. They also showed that when the bidding cost function is weakly concave, which includes all-pay auctions, a bid cap always decreases revenue.

In a companion paper, Olszewski and Siegel (2016b) also study contests with a large number of contestants and prizes, but focus on specific applications. Among other design tools, they also consider bid caps, and describe bid caps that are Pareto-improving with respect to all contestants.[1] In the present paper, we provide an example in which a bid cap is Pareto-improving with respect to all contestants and the collector of revenue.

The rest of the paper is organized as follows. Section 2 introduces the contest model. Section 3 introduces the mechanism-design approach to studying large contests. Section 4 investigates the effect of rigid bid caps. Section 5 investigates the effect of flexible bid caps. The appendix contains proofs not given in the main text.

2 Asymmetric contests

In the contests we consider, n players compete for n prizes of known values. The prizes are denoted by $0 \leq y_1^n \leq y_2^n \leq \cdots \leq y_n^n \leq 1$. Prizes of value 0 are "no prize," so it is without loss of generality to have the same number of prizes as players. Player i's privately-known type $x_i^n \in [0,1]$ is distributed according to a cdf F_i^n, and these distributions are commonly known and independent across players.[2] In the special case of complete information, each cdf corresponds to a Dirac measure. Each player chooses a bid t (which can be interpreted

[1] In their model, bids are interpreted as effort, and contestants fully internalize the effects of their effort.

[2] All probability measures are defined on the σ-algebra of Borel sets.

as effort or performance), the player with the highest bid obtains the highest prize, y_n^n, the player with the second-highest bid obtains the second-highest prize, y_{n-1}^n, and so on. Ties are resolved by a fair lottery. The utility of a player of type x from bidding $t \geq 0$ and obtaining prize y is

$$xh(y) - c(t), \tag{1}$$

where $h(0) = c(0) = 0$ and h and c are continuous and strictly increasing. Notice that (1) can accommodate private information about ability by dividing each player's utility by x to obtain $h(y) - c(t)/x$. Since we study some limits of sequences of contests when n diverges to infinity, we refer to a contest with n players and n prizes as the "n-th contest."

A (rigid) bid cap is an exogenously specified bid M. If the bid cap M is introduced, no player can bid more than M. The cap is *binding* if it is lower than some equilibrium bid of the unconstrained contest. A flexible bid cap is a continuous and strictly increasing function C with $C(0) = 0$ and $C(t) \geq c(t)$. If the flexible cap C is introduced, all bids are allowed, but the cost of bidding increases from c to C. The flexible cap is *binding* if $C(t)$ strictly exceeds $c(t)$ for some equilibrium bid t of the unconstrained contest. Every contest, with or without a cap, has at least one mixed-strategy Bayesian Nash equilibrium.[3]

3 Mechanism-design approach to studying contests

As pointed out in the introduction, the analysis of asymmetric contests of the kind described in Section 2 is difficult or impossible. To overcome this problem, we will use the mechanism-design approach to studying the equilibria of large contests, developed in Olszewski and Siegel (2016a). We now describe this approach, which allows us to approximate the equilibrium outcomes of large contests by considering the mechanism that implements a particular allocation of a continuum of prizes to a continuum of agent types.

[3] This follows from Corollary 5.2 in Reny (1999).

3.1 Limit distributions

We first formalize a requirement that the contests in the sequence be "sufficiently similar" as the number of players n grows large. Let $F^n = \left(\sum_{i=1}^n F_i^n\right)/n$, so $F^n(x)$ is the expected percentile ranking (up to $1/n$) of any player of type x in the n-th contest given the distribution of players' types. Denote by G^n the empirical prize distribution, which assigns a mass of $1/n$ to each prize y_j^n (recall that there is no uncertainty about the prizes). We require that F^n converge in weak*-topology to a continuous and strictly increasing distribution F, and G^n converge to some (not necessarily continuous) distribution G.[4] Notice that the restriction on F does not imply a similar restriction on distributions F_i^n of players' types, so these distributions may have gaps and atoms.[5]

The convergence of F^n and G^n to limit distributions F and G accommodates as a special, extreme case sequences of complete-information contests with asymmetric players, in which each player i's type distribution F_i^n in the n-th contest is a Dirac measure. A simple way to see this is to first choose the desired limit distributions F and G and then set player i's deterministic type in the n-th contest to be $x_i^n = F^{-1}(i/n)$ and prize j in the n-th contest to be $y_j^n = G^{-1}(j/n)$, where

$$G^{-1}(z) = \inf\{y \in [0,1] : G(y) \geq z\} \text{ for } 0 \leq z \leq 1.$$

Then, the n-th contest is one of complete information, F^n converges to F, and G^n converges to G. Another special, extreme case is ex-ante symmetric players, with $F_i^n = F^n$ for some distributions F^n that converge to F.

3.2 Assortative allocation and transfers

As will be stated in the next subsection, the mechanism that approximates the equilibrium outcomes of large contests implements the assortative allocation, which assigns to each type x prize $y^A(x) = G^{-1}(F(x))$. That is, the location in the prize distribution of the prize

[4] Convergence in weak*-topology can be defined as convergence of *cdf*s at points at which the limit *cdf* is continuous (see Billingsley, 1995).

[5] The restriction on F precludes a limit mass of players that have an atom at a particular type, as is the case when there is a non-vanishing fraction of identical players in a contest with complete information.

assigned to type x is the same as the location of type x in the type distribution. It is well known (see, for example, Myerson (1981)) that the unique incentive-compatible mechanism that implements the assortative allocation and gives type $x = 0$ a utility of 0 specifies for every type x bid

$$t^A(x) = c^{-1}\left(xh\left(y^A(x)\right) - \int_0^x h\left(y^A(z)\right) dz\right). \tag{2}$$

The aggregate bids in the mechanism that implements the assortative allocation are

$$\int_0^1 t^A(x)\, dF(x). \tag{3}$$

3.3 Approximation results

Corollary 2 in Olszewski and Siegel (2016a), which we state as Theorem 1 below, shows that the equilibria of large contests are approximated by the unique mechanism that implements the assortative allocation.

Theorem 1 *For any $\varepsilon > 0$ there is an N such that for all $n \geq N$, in any equilibrium of the n-th contest without a bid cap each of a fraction of at least $1 - \varepsilon$ of the players i obtains with probability at least $1 - \varepsilon$ a prize that differs by at most ε from $y^A(x_i^n)$, and bids with probability at least $1 - \varepsilon$ within ε of $t^A(x_i^n)$.*

Theorem 1 implies that aggregate bids in large contests without a bid cap can be approximated by (3). More precisely, we define the average bid as the aggregate bids in an equilibrium of the n-th contest divided by n. We then have the following corollary of Theorem 1, which appears as Corollary 1 in Olszewski and Siegel (2016c).

Corollary 1 *For any $\varepsilon > 0$ there is an N such that for all $n \geq N$, in any equilibrium of the n-th contest without a bid cap the average bid is within ε of (3).*

Similar approximation results hold when a bid cap is imposed. To formulate the results, we say that a bid cap M is binding in the limit if it is lower than the bid $t^A(x)$ of some type x. For any such bid cap M, a minor modification of the proof from Olszewski and Siegel (2016a) yields the following approximation result.

Theorem 2 *Given a bid cap M, which is binding in the limit, there exists a type x^* with $t^A(x^*) < M$ such that for any $\varepsilon > 0$ there is an N such that for all $n \geq N$, in any equilibrium of the n-th contest with bid cap M,*

(A) each of a fraction of at least $1-\varepsilon$ of the players i with type $x_i^n < x^*$ bids with probability at least $1-\varepsilon$ within ε of $t^A(x_i^n)$, and obtains with probability at least $1-\varepsilon$ a prize that differs by at most ε from $y^A(x_i^n)$;

(B) each of a fraction of at least $1-\varepsilon$ of the players i with type $x_i^n > x^*$ bids M with probability at least $1-\varepsilon$, and obtains with probability at least $1-\varepsilon$ the prize $y^A(x)$ for a randomly chosen $x > x^*$ distributed according to F contingent on $x > x^*$.

Type x^* is indifferent between: (i) bidding $t^A(x^*)$ and obtaining prize $y^A(x^*)$; and (ii) bidding M and obtaining prize $y^A(x)$ for a randomly chosen $x > x^*$ distributed according to F contingent on $x > x^*$. This indifference can be expressed as

$$x^* \frac{\int_{x^*}^1 h(y^A(x))dF(x)}{1-F(x^*)} - c(M) = \int_0^{x^*} h\left(y^A(z)\right) dz, \tag{4}$$

as the left-hand side is the utility type x^* obtains from (ii), and by (2) the right-hand side is the utility type x^* obtains from (i).

We denote by $t^M(x)$ the approximating limit bid of type x with a bid cap M, i.e., $t^M(x) = t^A(x)$ for $x < x^*$ and $t^M(x) = M$ for $x \geq x^*$. We then obtain the equivalent of Corollary 1.

Corollary 2 *For any $\varepsilon > 0$ there is an N such that for all $n \geq N$, in any equilibrium of the n-th contest with a bid cap M, which is binding in the limit, the average bid is within ε of (3) with t^M instead of t^A.*

4 The effect of a bid cap

4.1 Linear or concave bidding cost

We first prove the following result for a general bidding cost c, which immediately implies that a bid cap cannot increase the aggregate bids in a large all-pay auction, that is, a large contest in which $c(t) = t$.

Theorem 3 *For any bid cap M, the aggregate cost of the limit bids with the cap is lower than without the cap. That is,*

$$\int_0^1 c(t^M(x))dF(x) \leq \int_0^1 c(t^A(x))dF(x).$$

The inequality is strict if the cap is binding in the limit.

An almost immediate corollary of Theorem 3 and Corollary 2 is the following result.

Corollary 3 *A bid cap weakly decreases the limit aggregate bids (3) when the bidding cost c is weakly concave. The decrease is strict if the cap is binding in the limit. Thus, a binding cap reduces aggregate bids in any equilibrium of a sufficiently large contest.*

4.2 Comparison to small all-pay auctions

Corollary 3 contrasts with the finding of Che and Gale (1998) that bid caps may increase aggregate bids in two-player all-pay auctions with complete information.[6] To see why the results are so different for small and large contests, we recall Che and Gale's (1998) result in the context of an example.

Example 1 *Consider an all-pay auction with two players and one prize. Suppose player 1 values the prize at* 2 *and player 2 values the prize at* 8. *Without a bid cap, the unique equilibrium of this all-pay auction has the following form: player 1 bids $t = 0$ with probability* $3/4$, *and with the remaining probability randomizes uniformly over the interval* $(0, 2)$. *Player 2 randomizes uniformly over the interval* $(0, 2)$. *The resulting expected aggregate bids are* $5/4$.

In this setting, a bid cap $T > 1$ has no effect on the expected aggregate bids. If, for example $T = 3/2$, then the equilibrium has the following form: player 1 bids $t = 0$ with probability $3/4$, *randomizes uniformly on the interval* $(0, 1)$ *with probability* $1/8$, *and bids* $3/2$ *with the remaining probability of* $1/8$. *Player 2 randomizes uniformly on the interval* $(0, 1)$ *with probability* $1/2$, *and bids* $3/2$ *with the remaining probability of* $1/2$. *The expected*

[6] Note that our results from the previous section hold for large contests with both complete and incomplete information and for all equilibria.

aggregate bids in this equilibrium are still $5/4$*. In contrast, for all caps* T *between* 0 *and* 1*, in equilibrium both players bid* $t = T$ *with probability* 1*, so the aggregate bids are* $2T$*. Thus, the expected aggregate bids increase for* $T \in (5/8, 1)$*.*

The intuition behind the example is as follows. When a binding bid cap T is imposed, players shift to T the mass their strategies previously assigned to bids higher than T. This makes bid T attractive compared to bids close to but lower than T, because by bidding T a player now has a higher chance of winning. Thus, players also shift to T the mass their strategies assigned to bids close to but lower than T. This in turn make bid T less attractive. The shift of mass from bids lower than T continues until the excess attractiveness of bidding T, due to the mass assigned to bids higher than T being shifted to T, is entirely eliminated. If $T > 1$, the revenue increase from the shift in mass from the lower bids is equal to the revenue loss from the shift in mass from the higher bids. But for $T < 1$, things are different. Then, even if the entire mass assigned to bids between 0 and T is shifted to T, the mass shifted to T of bids higher than T makes bidding T still relatively attractive to the player with the lower valuation. As a result, this player's atom at bid 0 is shifted in its entirety to T, and this discrete shift generates a discontinuous increase in revenue.

Such a discontinuous shift is possible because the two players face different competition: the low-valuation player faces the equilibrium bidding strategy of the high-valuation player, while the high-valuation player faces the equilibrium bidding strategy of the low-valuation player. When the number of players is large, the competition they face is similar, and coincides in the limit, because they each face almost the same set of competitors. Consequently, a small change in the value of a bid cap only has a small impact on the "average" equilibrium strategy. Thus, bid caps cannot increase aggregate bids. Indeed, it follows from (4) and the continuity of distribution F that x^* is a continuous function of M. Given this intuition, one can readily imagine that in a symmetric equilibrium, when players are ex-ante identical, incomplete information may play a similar role to that of a large number of bidders, preventing small changes in the value of a bid cap from having a large impact on equilibrium strategies. Thus, the results of Gavious et al. (2002), which also hold for small contests, are consistent with the ones we obtain for large contests, and differ from the results of Che and Gale (1998).

This intuition also indicates that the difference in the effect of bid caps between large contests and small contests with ex-ante asymmetric players (or in an asymmetric equilibrium) is not due to the assumption that distribution F is continuous. While we cannot investigate this in a rigorous manner, since the methods of Olszewski and Siegel (2016a) do not generalize easily to distributions with atoms, we intuitively see no reason for our results not to generalize to arbitrary distributions. To see why, suppose, for a moment, that the approach from Olszewski and Siegel (2016a) can be applied to any distribution F, and consider an F with an atom at the cutoff x^*. Then, for large n a fraction α of players with type x^* bids with high probability close to $t^A(x^*)$, and obtains with high probability a prize close to that assigned to agents of type x^* by the assortative allocation, and the remaining fraction of $1-\alpha$ bids M with high probability and obtains with high probability the prize $y^A(x)$ for a randomly chosen $x \geq x^*$ distributed according to F contingent on $x > x^*$ or $x = x^*$ and the agent being in the fraction $1-\alpha$ of the atom at x^*. The indifference of type x^* between bidding $t^A(x^*)$ and bidding M yields a condition analogous to (4), and the rest of the proof of Theorem 3 remains unaltered.

Thus, it seems that the impossibility of increasing aggregate bids with a bid cap is a consequence of small changes in the value of a bid cap having only a small impact on equilibrium strategies, which arises when different players face similar competition.

4.3 Convex bidding cost

As demonstrated by Gavious et al. (2002) in a single-prize setting, when the bidding cost is convex, some bid caps increase aggregate bids. Before discussing conditions under which this happens, we demonstrate this possibility and provide an intuition by an example. Our example also demonstrates that some bid caps are Pareto-improving, a result analogous to one obtained by Olszewski and Siegel (2016b) in a slightly different setting.

Example 2 *Suppose the limit distribution F is uniform. In the limit, there is a mass $1/2$ of identical prizes. That is, the limit distribution G has an atom of size $1/2$ at 0 and an atom of size $1/2$ at 1. Suppose first that the cost of bidding is $c(t) = t$ for all t, and no bid cap is imposed. Then, in the limit mechanism, each type $x > 1/2$ obtains a prize, and no type*

$x < 1/2$ obtains a prize. The former types bid $1/2$, and the latter types bid 0. The revenue (that is, the average bid) is $1/4$.

For the cap of $1/4$, we have that x^*, which is determined by equation (4), is equal to $1/3$. Types $x > 1/3$ bid $1/4$, types $x < 1/3$ bid 0, and the average bid is $(2/3)(1/4) < 1/4$.

Suppose now that $c(t) = t$ for $t < 1/4$ and $c(t) = 4t - 3/4$ for $t > 1/4$. This function is convex. Without a cap, each type $x > 1/2$ obtains a prize and no type $x < 1/2$ obtains a prize. The former types bid $5/16$, which costs them $c(5/16) = 1/2$, and the latter types bid 0. A flexible cap C does not change the structure of this equilibrium. The only difference is that each type $x > 1/2$ bids t^C such that $C(t^C) = 1/2$, which clearly reduces aggregate bids, as $c(t^C) \leq C(t^C)$. In particular, this flexible cap can be $C(t) = t$ for $t < 1/4$ and $C(t) = mt - (1/4)(m-1)$ for $t > 1/4$. For this family of flexible caps, the revenue decreases as m increases.

When m diverges to infinity, our family of flexible caps converges the the (rigid) cap of $1/4$. However, this "limit" cap increases revenue compared to the scenario without any cap, as $(2/3)(1/4) > (1/2)(5/16)$. Notice that not only is the revenue higher with the bid cap, but each contestant is weakly better off (and some are strictly better off). The payoff of a type $x \leq 1/3$ is 0, regardless of whether the cap is imposed, as these types bid 0 and obtain no prize. Types $x > 1/3$ are strictly better off as their payoff is $(3/4)x - (1/4) > 0$ with the cap, and is $\max\{0, x - 1/2\} < (3/4)x - (1/4)$ without the cap. (Notice that the probability of obtaining a prize in the scenario with the cap is $(1/2)/(2/3) = 3/4$.)

The intuition behind the example is that without a cap, agents with high valuations ($x > 1/2$) and agents with low valuations ($x < 1/2$) can be separated by a bid t such that the former prefer obtaining a prize at bid t to obtaining no prize at bid 0, but the latter have the opposite preference. This separation is possible for every continuous cost function, whether or not it involves a flexible cap, because what matters for the separation is not the bid but the cost incurred by agents. If the cost becomes higher, the bid that separates the agents becomes lower, and in the limit coincides with the rigid cap. With the rigid cap, however, agents can no longer be separated, because the cap (and any lower bid) is not a bid t such that agents with higher valuations prefer obtaining a prize at bid t to obtaining no prize at bid 0, but agents with lower valuations have the opposite preference. Thus, when the cap becomes rigid a mass of lower-valuation agents who bid 0 with a flexible cap instead bid the rigid cap. This increases the aggregate bids discretely, so if the separating bid t was sufficiently close to the rigid cap ($m = 4$ is sufficiently high for this purpose in our example), the discrete increase in the bids of the mass of lower-valuation agents overcomes the small decrease in the bids of higher-valuation agents.

The intuition for obtaining a Pareto improvement can be explained as follows. Without a cap, high-valuation contestants bid in a steep region of the cost function. Even a small reduction of their bids enables them to save a lot in terms of costs. Thus, the gain coming from the cap, which reduces their bids, is higher than the loss coming from the possibility of obtaining a lower prize. In turn, low-valuation contestants bid in a flat region of the cost function. Increases in their bids in this region do not raise their costs by much. Thus, the loss coming from making higher bids in response to the cap is lower than the gain coming from the possibility of obtaining a higher prize.

We now provide some sufficient conditions for a bid cap to increase the aggregate bids. First, consider a specific (but important) distribution of prizes, which consists of a mass $\beta < 1$ of identical prizes. That is, G has an atom of size $1 - \beta$ at 0 and an atom of size β at 1. To simplify notation, we assume that $h(1) = 1$. Then, given a bid cap M, (4) simplifies to

$$\frac{x^* \beta}{1 - F(x^*)} = c(M). \tag{5}$$

In this case, we can provide not only sufficient but also necessary conditions.

Theorem 4 *If the distribution of prizes consists of a mass $\beta < 1$ of identical prizes, then the limit aggregate bids with a bid cap M that is binding in the limit strictly exceed those without it if and only if*

$$(1 - F(x^*)) M > \beta c^{-1} \left(F^{-1}(1 - \beta) \right). \tag{6}$$

Proof. Without the cap, types $x < F^{-1}(1 - \beta)$ bid 0 and obtain no prize, and types $x > F^{-1}(1 - \beta)$ bid $c^{-1}(F^{-1}(1 - \beta))$ and obtain a prize. Thus, the aggregate bids without the cap are $\beta c^{-1}(F^{-1}(1 - \beta))$. For a cap $M < c^{-1}(F^{-1}(1 - \beta))$, types lower than x^* bid

0 and obtain no prize, and types higher than x^* bid M and obtain a prize with probability $\beta/[1-F(x^*)]$. Thus, the aggregate bids are higher with the cap than without it and only if **condition (6)** is satisfied. ■

This seemingly simple result allows to determine the exact range of caps M that increase the aggregate bids for specific parameter values. For example, if F is uniform, $\beta = 1/4$, and $c(t) = t^2$, then by (5) we have $x^* = 4M^2/(1+4M^2)$. By plugging this into (6), we obtain a quadratic inequality, which is satisfied for $M \in (\sqrt[2]{3}/6, \sqrt[2]{3}/2)$.

Theorem 4 also allows to derive more specific sufficient conditions, which are easier to interpret.

Proposition 1 *For any convex cost c and any bid cap $M > 0$ that is binding in the limit, if the mass of identical prizes $\beta > 0$ is small enough, then the limit aggregate bids with the cap M strictly exceed those without the cap.*

The next result allows for a general prize distribution G and generalizes the observation made in Example 2.

Proposition 2 *Suppose that a bid cap M is binding in the limit and the cost function is convex and differentiable at the cap M. Fixing the value of $c(t)$ for $t < M$, if the derivative of the cost function $c'(M)$ at the bid cap M is sufficiently large, then the limit aggregate bids with the cap M strictly exceed those without the cap.*

Remark 1 *Gavious et al. (2002) show that in a single-prize contest, if the bidding cost is convex, then a binding cap increases aggregate bids if (a) the number of players is sufficiently large, or (b) the cost function is sufficiently convex. Propositions 1 and 2 imply analogous results in our limit setting with multiple prizes.*

An immediate corollary of Theorem 2 and Corollary 1 is the following result.

Corollary 4 *Under the conditions of Propositions 1 and 2, the bid cap M strictly increases aggregate bids in any equilibrium of a sufficiently large contest.*

5 The effect of a flexible bid cap

In Example 2, introducing a flexible cap decreased aggregate bids. We now show that in large contests this is true for any cost function and any flexible cap. For the result, let $t^C(x)$ denote the limit bid of type x under a flexible bid cap C. It follows immediately from Theorem 1 that the limit equilibrium allocation under the cap C is the assortative allocation, and the bids $t^C(x)$ are determined by equation (2), with c replaced by C. This yields the following result. We say that a flexible cap is binding in the limit if $C\left(t^A(x)\right) > c\left(t^A(x)\right)$ for some type x.

Theorem 5 *A flexible bid cap decreases the limit bid of each type x. Therefore, it decreases the limit aggregate bids. If the flexible cap is binding in the limit, the decrease is strict.*

Proof. It follows from (2) that $c(t^A(x)) = C(t^C(x))$ for every $x \in [0,1]$. Since $C(t) \geq c(t)$ for all t, it follows that $t^C(x) \leq t^A(x)$. ■

An immediate corollary of Theorem 5 and Corollary 1 is the following result.

Corollary 5 *A flexible bid cap that is binding in the limit strictly decreases aggregate bids in any equilibrium of a sufficiently large contest.*

The intuition for this result is analogous to that given by Kaplan and Wettstein (2006) in the two-player model. It is the cost of a bid, not the bid itself, that matters for the player's payoff. Since there is a one-to-one correspondence between bids and their costs, one may interpret a contest as a game in which players choose not bids but costs. When we replace cost function c with cost function C, the structure of equilibria (in both the finite and the limit case) does not change. Players choose the same costs in both cases. And since cost C corresponding to any bid t is higher than cost c, the bids under c are higher than the bids under C.

As shown in Example 2, with a (rigid) bid cap we no longer have a one-to-one correspondence between bids and their costs. This cap, as opposed to a flexible cap, imposes an upper bound on costs, which makes high types choose the bid corresponding to this upper bound. This creates an opportunity for lower types to pool with the high types by increasing their bids (their costs) by a relatively small amount, and thus gain a chance of obtaining one of the prizes obtained by the high types in the unrestricted contest. The extra revenue created by the higher bids of these lower types may offset the loss of revenue caused by the bound imposed on the bids of the high types.

Until now, we studied the utility function $xh(y) - c(t)$, which implied that all bidders have the same cost of bidding. With heterogeneous bidding costs, however, Che and Gale (2006) demonstrated the possibility of increasing aggregate bids by imposing a flexible bid cap. The intuition behind their result is as follows. In their two-player model, if a flexible bid cap raises the cost of a lower-cost bidder by more than the cost of a higher-cost bidder (more precisely, what Che and Gale call an "equalizing shift" of costs), the competition between the two bidders becomes more equal, which provides incentives for more aggressive bidding. This strategic effect may dominate the more direct effect that a higher cost discourages higher bids.

But in large contests, flexible bid caps always decrease aggregate bids. That is, Theorem 5 and Corollary 5 hold even if the flexible caps vary across players. To see why, we note that Theorems 1 and 2 as well as their corollaries apply to this more general environment, because Olszewski and Siegel's (2016a) approximation result applies to a more general utility function $u(x, y, t)$ that satisfies a single-crossing condition. This condition is satisfied, for example, by $u(x, y, t) = v(x, y) - c(x, t)$, where v increases in x (for all $y > 0$) and y, and c increases in t and decreases in x (for all $t > 0$), with $v(x, 0) = c(x, 0) = 0$ for all x. For simplicity (although the result holds for any prize distribution), suppose that the limit prize distribution consists of a mass $\beta < 1$ of identical prizes. (So G has an atom of size $1 - \beta$ at 0 and an atom of size β at 1).

Without a cap, in the limit mechanism the mass β of the highest types obtains prizes and bids t, where t is the bid that makes type $x = F^{-1}(1 - \beta)$ indifferent between obtaining a prize at that bid and obtaining no prize at bid 0. That is,

$$v\left(F^{-1}(1-\beta), 1\right) = c(F^{-1}(1-\beta), t).$$

With a flexible cap $C(x, t) > c(x, t)$, the allocation is the same, and the payment of any type that obtains a prize is determined by the same equation with c replaced by C. This implies that the bid of each type, and therefore the aggregate bids, is lower with the cap.

Intuitively, when the contest is sufficiently large a flexible bid cap has a negligible effect on the allocation of prizes, in contrast to the small contests of Che and Gale (2006), in which what they call an equalizing shift results in a higher probability of obtaining the prize by the contestant with the lower valuation.

6 Appendix

Proof of Theorem 3. By definition of t^M, it suffices to show that

$$\int_{x^*}^{1} c(M)dF(x) \leq \int_{x^*}^{1} c(t^A(x))dF(x), \tag{7}$$

where type x^* is defined in Theorem 2. By rearranging (4) we obtain

$$(1 - F(x^*))c(M) = x^* \int_{x^*}^{1} h(y^A(x))dF(x) - (1 - F(x^*)) \int_0^{x^*} h\left(y^A(z)\right) dz. \tag{8}$$

Since

$$(1 - F(x^*))c(M) = \int_{x^*}^{1} c(M)dF(x),$$

the equality (8) shows that (7) is equivalent to

$$x^* \int_{x^*}^{1} h(y^A(x))dF(x) - (1 - F(x^*)) \int_0^{x^*} h\left(y^A(z)\right) dz \leq \int_{x^*}^{1} c(t^A(x))dF(x). \tag{9}$$

Since

$$(1 - F(x^*)) \int_0^{x^*} h\left(y^A(z)\right) dz = \int_{x^*}^{1} \left[\int_0^{x^*} h\left(y^A(z)\right) dz\right] dF(x),$$

to show (9) it suffices to show that

$$x^* h(y^A(x)) - \int_0^{x^*} h\left(y^A(z)\right) dz \leq c(t^A(x))$$

for every $x > x^*$. By substituting (2), this inequality is equivalent to

$$x^* h(y^A(x)) - \int_0^{x^*} h\left(y^A(z)\right) dz \leq x h\left(y^A(x)\right) - \int_0^{x} h\left(y^A(z)\right) dz,$$

which simplifies to

$$\int_{x^*}^{x} h\left(y^A(z)\right) dz \leq (x - x^*) h\left(y^A(x)\right).$$

This inequality holds for every $x > x^*$, since y^A is increasing, and is strict for an interval of types $x > x^*$ when the cap is binding,[7] which completes the proof.

Proof of Corollary 3. It suffices to show that

$$\int_{x^*}^{1} M dF(x) \leq \int_{x^*}^{1} t^A(x) dF(x),$$

with a strict inequality if the cap is binding in the limit.

Consider the random variable that takes value $c(t^A(x))$ when $x > x^*$ and value $c(M)$ when $x \leq x^*$. By Jensen's inequality applied to this random variable,

$$c^{-1}\left(\int_{x^*}^{1} c(t^A(x)) dF(x) + \int_{0}^{x^*} c(M) dF(x)\right) \leq \int_{x^*}^{1} c^{-1}(c(t^A(x))) dF(x) + \int_{0}^{x^*} c^{-1}(c(M)) dF(x).$$

By Theorem 3,

$$c^{-1}\left(\int_{x^*}^{1} c(M) dF(x) + \int_{0}^{x^*} c(M) dF(x)\right) \leq c^{-1}\left(\int_{x^*}^{1} c(t^A(x)) dF(x) + \int_{0}^{x^*} c(M) dF(x)\right),$$

with a strict inequality if the cap is binding in the limit. Thus,

$$\int_{x^*}^{1} M dF(x) = c^{-1}\left(\int_{x^*}^{1} c(M) dF(x) + \int_{0}^{x^*} c(M) dF(x)\right) - \int_{0}^{x^*} M dF(x)$$

cannot exceed, and if the cap is binding in the limit is strictly less than,

$$\int_{x^*}^{1} c^{-1}(c(t^A(x))) dF(x) + \int_{0}^{x^*} c^{-1}(c(M)) dF(x) - \int_{0}^{x^*} M dF(x) = \int_{x^*}^{1} t^A(x) dF(x).$$

[7] If $y^A(z)$ were the same for all $z \in (x^*, 1]$, then with no cap type x^* would be indifferent between bidding $t^A(x^*)$ and bidding $t = t^A(z)$ for $z \in (x^*, 1]$. So, with a binding cap $M < t$, types slightly below x^* would bid M, which would imply that type x^* is not indifferent between bidding $y^A(x^*)$ and M.

Proof of Proposition 1. By substituting $x^*\beta/c(M)$ for $1-F(x^*)$ (from (5)), condition (6) is equivalent to

$$\frac{x^*}{x} > \frac{c(M)}{M} / \frac{c(T)}{T}, \tag{10}$$

where $T = c^{-1}(F^{-1}(1-\beta))$ is the bid of the types who obtain prizes in the limit mechanism without cap, and $x = c(T) = F^{-1}(1-\beta)$ is the type that is indifferent between bidding T and obtaining a prize, and bidding 0 and obtaining no prize.

Since (5) implies that $x^* \to 1$ as $\beta \to 0$, (10) holds for sufficiently small $\beta > 0$ if it holds for $\beta = 0$ and $x^* = 1$, that is, if and only if

$$\frac{c(T)}{T} > \frac{c(M)}{M} \tag{11}$$

for $T = c^{-1}(1)$. The left-hand side is the average slope of c on $[0, c^{-1}(1)]$, and the right-hand side is the average slope of c on $[0, M]$. Since $M < c^{-1}(1)$ (no type bids more than $c^{-1}(1)$ and M is binding in the limit) and c is convex, (11) holds.

Proof of Proposition 2. The cap decreases the bids of types higher than x^M in the limit mechanism, where x^M is the type who provides effort M in the limit mechanism without a cap, and increases the bids of types between x^* and x^M, where x^* is the cutoff type that is indifferent between bidding $t^A(x)$ and obtaining prize $y^A(x)$, and bidding M and obtaining the prize $y^A(z)$ for a randomly chosen $z > x^*$. The aggregate bids are higher with the cap than without it if and only if

$$\int_{x^*}^{1}(M - t^A(x))dF(x) = \int_{x^*}^{x^M}(M - t^A(x))dF(x) + \int_{x^M}^{1}(M - t^A(x))dF(x) > 0. \tag{12}$$

If $c'(M) \to \infty$, then (2) implies that $t^A(x) \to M$ for all $x \in (x^M, 1]$. Then, the second component of the summation in (12) tends to 0, but the first component, which is positive, stays unaltered. Thus, the cap decreases the aggregate bids of types higher than x^M by less than it increases the aggregate bids of types between x^* and x^M.

References

[1] **Billingsley, Patrick.** 1995. “Probability and Measure.” John Wiley and Sons, Inc.

[2] **Che, Yeon-Koo, and Ian Gale.** 1998. “Caps on Political Lobbying.” *American Economic Review*, 88(3): 643-651.

[3] **Che, Yeon-Koo, and Ian Gale.** 2006. “Caps on Political Lobbying: Reply.” *American Economic Review*, 96(4): 1355-1360.

[4] **Gavious, Arieh, Moldovanu, Benny and Aner Sela.** 2002. “Bid Costs and Endogenous Bid Caps.” *RAND Journal of Economics*, 33(4): 709-722.

[5] **Kaplan, Todd R. and Wettstein, David.** 2006. “Caps on Political Lobbying: Comment.” *American Economic Review*, 96(4): 1351–1354.

[6] **Myerson, Roger B.** 1981. “Optimal Auction Design.” *Mathematics of Operations Research*, 6(1): 58-73.

[7] **Olszewski, Wojciech, and Ron Siegel.** 2016a. “Large Contests.” *Econometrica*, 84(2): 835-854.

[8] **Olszewski, Wojciech, and Ron Siegel.** 2016b. “Effort Caps in College Admissions and Grant Applications.” Mimeo.

[9] **Olszewski, Wojciech, and Ron Siegel.** 2016c. “Effort-Maximizing Contests.” Mimeo.

[10] **Reny, Philip J.** 1999. “On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games.” *Econometrica*, 67(5): 1029-1056.

[11] **Szech, Nora.** 2015. “Tie Breaks and Bid-Caps in All-Pay Auctions.” *Games and Economic Behavior*, 92: 138-149.