Tie-Breaks and Bid-Caps in All-Pay Auctions*

Nora Szech

February 2012

Abstract

We revisit the complete information all-pay auction with bid-caps introduced by Che and Gale (1998), dropping their assumption that tie-breaking must be symmetric. Any choice of tie-breaking rule leads to a different set of Nash equilibria. Compared to the optimal bid-cap of Che and Gale we obtain that in order to maximize the sum of bids, the designer prefers to set a less restrictive bid-cap combined with a tie-breaking rule which slightly favors the weaker bidder. Moreover, the designer is better off breaking ties deterministically in favor of the weak bidder than symmetrically except when bidding costs are strongly convex.

JEL Classification: D72, D44

Keywords: All-Pay Auctions, Contests, Bid-Caps, Tie-Breaking, Lobbying

*I would like to thank Benny Moldovanu, Frank Rosar, Patrick Schmitz, Ron Siegel and Christian Westheide. Nora Szech, University of Bonn, Lennéstr. 37, 53113 Bonn, Germany; email: nszech@uni-bonn.de.

1 Introduction

At the Olympic Games of 1896, the first Olympic Games of the Modern Era, two weightlifters impressed the Athenian audience particularly: the Scotsman Launceston Elliot and the Dane Viggo Jensen. Both lifted the same highest weight. The jury decided to solve the tie in favor of Jensen for he was considered to have the better style. Unfamiliar with this tie-breaking rule, the British delegates protested against the decision. This finally led to Elliot and Jensen obtaining the permission to try again for lifting higher weights. Yet both failed. In the end, Jensen was declared the champion.

Nowadays, it is neither style nor the energy of the own country's delegates that helps to win a tie in a weightlifting-contest. Instead, when a tie occurs, it is the own bodyweight that decides whether one wins or loses: Tie-breaking is solved in favor of the lighter athlete. Behind this is the idea that a lighter athlete, though in the same weight-class, probably has to exert more effort to lift the same weight than a heavier competitor.

Clearly, sport contests are more interesting the more intensely the athletes compete. For the designer, it his hence a natural objective to maximize the sum of efforts exerted by the contestants or equivalently the average effort per contestant. Che and Gale (1998) show that handicaps can be a very effective tool to raise the effort levels in all-pay contests like sport contests. Yet they restrict their analysis to symmetric tie-breaking. We are going to allow the designer not only to set handicaps optimally, but also to choose the optimal tie-breaking rule.

It turns out that any choice of tie-breaking rule leads to a different set of Nash equilibria. We provide a complete characterization of the rich equilibrium structure. Depending on the combination of tie-breaking rule and bid-cap, the designer can enforce pure equilibria as well as mixed equilibria where either of the bidders earns zero payoff. Compared to Che and Gale's optimal bid-cap under symmetric tie-breaking, we obtain that the designer optimally sets a less restrictive bid-cap combined with a tie-breaking rule which slightly favors the ex ante weaker bidder. Whereas the unique Nash equilibrium of the unrestricted all-pay auction is in mixed strategies, both of these policies have the effect that they force bidders to play a pure strategy equilibrium where both bidders bid the bid-cap. The optimal policy exploits the fact that the weaker bidder is willing to bid more if tie-breaking is biased in his favor. If this bias is not too large this does not deter the stronger bidder from competition.

In real contests, the toss of an unfair coin may seem rather unintuitive for lay people (and more difficult to realize – solving a tie 60 : 40 requires some preparation, while a fair coin is always at hands). Indeed, in many applications ties are broken either in favor of one bidder or 50 : 50. Starting with this observation we consider the designer's problem if he is restricted to choosing between symmetric and deterministic tie-breaking rules. Even under this restriction, the designer can do better than in the optimal policy of Che and Gale (1998) by setting a bid-cap which is just small enough to influence equilibrium behavior and choosing to break ties in favor of the weak bidder. Concretely, this means that he excludes bids which are larger than the weak bidder's valuation.

Superficially, this policy seems like a minimal intervention into the game but it has important consequences. Recall that in an unrestricted all-pay auction there is a mixed equilibrium in which the weak bidder lays out with a positive probability. The designer's policy gives rise to an equilibrium which differs from the equilibrium of the unrestricted all-pay auction in only one respect: The weak bidder makes a preemptive bid (by bidding at the bid-cap) with the same probability with which he would lay out in the unrestricted auction.

Finally, we show that if σ is the designer's payoff from the unrestricted all-pay auction, then the policy of Che and Gale yields at most 2σ. The policy of favoring the weak bidder and setting a minimal bid-cap yields at most 3σ and the globally optimal policy yields at most 4σ. The advantage of all three bid-cap/tie-breaking policies over the unrestricted all-pay auction gets larger when bidders get more asymmetric.

From these results, it follows that there appear to be two promising policies for the designer: 1) He can enforce a pure equilibrium by setting a rigid bid-cap and use a tie-breaking rule that is symmetric or slightly favors the weaker bidder. 2) He can break ties deterministically in favor of the weak bidder and set a bid-cap that is just rigid enough to matter. We study the robustness of these policies in more general settings, trying to understand under which circumstances each of these policies is preferable.

Relying on and extending recent results of Siegel (2010), we compare the performance of our policies in a setting with considerably more general cost functions. The second policy is remarkably robust against details of the cost function. Notably, it always yields an improvement over the unrestricted all-pay auction. In contrast, the first strategy is better for sufficiently convex cost functions while for sufficiently concave cost functions it performs even worse than the unrestricted all-pay auction.

Moreover, the requirements on the designer's knowledge are considerably stronger for the first policy than for the second. Finally unlike the first one, the second policy leads to an improvement over the unrestricted all-pay auction regardless of whether the designer maximizes the winning bid, the sum of bids, or a mixture of the two.

Together, we believe that these results demonstrate clearly that tie-breaking rules matter in all-pay auctions with bid-caps, and explain why we often observe deterministic tie-breaking rules in favor of the weaker bidder in practice.

1.1 Related Literature

In the vast literature on all-pay auctions[1], tie-breaking rules have received comparatively little attention so far. Indeed, in many all-pay auction games, tie-breaking occurs with probability zero in any equilibrium and thus the set of Nash equilibria is invariant to the choice of tie-breaking rule. The simplest example is a complete information all-pay auction where at least two bidders have positive valuations for the object for sale.[2] In other related games, the choice of tie-breaking rule is a necessity since a Nash equilibrium exists only for certain tie-breaking rules. For the possibly simplest example consider a two-player complete information all-pay auctions where bidders have valuations $v_1 > 0$ and $v_2 = 0$. Then a Nash equilibrium (which is constituted by both bidders bidding zero) exists only if tie-breaking always favors bidder 1.

In contrast, in all-pay auctions with (binding) bid-caps the choice of tie-breaking rule is decisive since in equilibrium both bidders play the bid-cap with positive probability. Yet in the literature so far only the case of symmetric tie-breaking has

[1]See Konrad (2009) for an overview.

[2]This has been shown, among others and in increasing generality, by Hillman and Samet (1987), Hillman and Riley (1989), Baye, Kovenock and de Vries (1996), and Siegel (2009). A parallel result holds for the incomplete information case studied first by Weber (1985) and Hillman and Riley (1989).

been considered. This concerns both, the complete information case studied by Che and Gale (1998) and Persico and Sahuguet (2006)[3] and the incomplete information case studied by Gavious, Moldovanu and Sela (2003) and Sahuguet (2006).[4]

Kaplan and Wettstein (2006) initiated a discussion of "soft" bid-caps which do take the form of an artificial stronger increase in the cost function from some effort-level on (as opposed to a binding upper bound on bids). The relation to this literature is discussed in Section 6.7.

See Che and Gale (1998) for a discussion of the relation to policies other than bid-caps such as minimum-bid requirements.

1.2 Outline

The paper is structured as follows: Section 2 introduces the model. Section 3 characterizes the bidders' equilibrium behavior for all combinations of bid-caps and tie-breaking rules. Section 4 analyzes the designer's optimization problem, first for the case of arbitrary tie-breaking rules and then for the case of tie-breaking rules that are either deterministic or symmetric. Section 5 generalizes the previous analysis to general cost functions. Section 6 discusses robustness, and various extensions and implications of our analysis. Among the issues addressed are knowledge requirements of our various policies, the number of bidders, legal prescriptions of treating bidders symmetrically, different objectives of the designer, the incomplete information case, and "soft" bid-caps. We delegate all proofs to the Appendix.

2 The Model

We consider a complete information all-pay auction with two bidders 1 and 2 with valuations v_1 and v_2 for winning. Unless otherwise noted $v_1 > v_2$. Each bidder is restricted to choose his bid b from the interval $[0, m]$ at costs of b. If a bidder submits the strictly highest bid, he wins. If both bidders submit the same bid, bidder 1 wins with probability $\alpha \in [0, 1]$, otherwise bidder 2 wins. We assume that the designer chooses α and the handicap-level m in order to maximize the sum of

[3] Persico and Sahuguet (2006) embed the model of Che and Gale (1998) into a model of electoral competition where parties try to attract heterogeneous voters. In their setting, the assumption of symmetric tie-breaking takes the form of an assumption that undecided voters toss a fair coin.

[4] See Section 6.6 for further discussion of the incomplete information case.

bids or, equivalently, the average bid. This is the setting of Che and Gale (1998) with the only difference that they restrict their analysis to the case of $\alpha = \frac{1}{2}$, and hence symmetric tie-breaking.

Note that it is without loss of generality to focus on tie-breaking rules that do not change depending on the bids exerted:

Lemma 1 *Under any, possibly bid-dependent tie-breaking rule, in any equilibrium and for all* m*, no bidder sets an atom inside* $(0, m)$*. Moreover, at most one bidder places an atom on* 0*.*

Hence, in equilibrium, tie-breaking may occur with positive probability only if both bidders bid the handicap-level m. Thus it is without loss of generality to analyze tie-breaking rules which are not bid-dependent.

If the designer does not impose a handicap, i.e. $m = \infty$, the game is simply the complete information all-pay auction. It is well-known that then, in the unique equilibrium, both bidders mix uniformly over $[0, v_2)$.[5] Moreover, bidder 2 places an atom of size $1 - \frac{v_1}{v_2}$ on 0. Clearly, for $m \geq v_2$, this equilibrium survives. Yet we will see below that depending on the tie-breaking rule uniqueness of equilibrium may break down for $m = v_2$. The case $m = 0$ is trivial. Thus we assume in the following that the designer chooses (α, m) from the set

$$C = [0, 1] \times (0, v_2].$$

3 All Equilibria

We now characterize all equilibria for all pairs $(\alpha, m) \in C$.

3.1 Overview

Let us start with a general overview and an exemplary picture of how the structure of equilibria looks like. For this, we partition C into three sets, C_I, C_{II}, and C_{III}, corresponding to three qualitatively different cases I, II, and III. Figure 1 depicts these three sets for $v_1 = 2$ and $v_2 = 1$.[6] The boundaries have to be analyzed separately. Thus we denote the segment separating C_I and C_{II} by S, the upper

[5] See, e.g., Hillman and Riley (1989).

[6] For different values $v_1 > v_2$, the picture is rescaled, but it does not change qualitatively.

boundary of C_{III} by T, and the upper boundary of C_{II} by U. For future reference, we denote the line where α is fixed at value $\frac{1}{2}$ by C_G, which is the set studied by Che and Gale (1998).

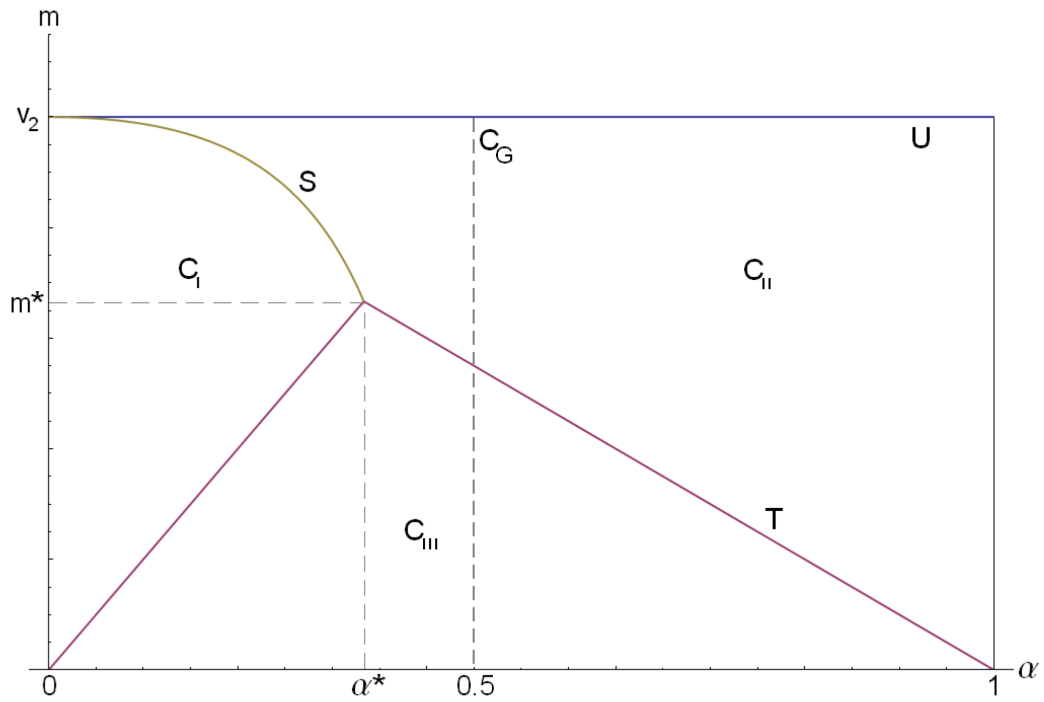

Figure 1: Partition of the designer's choice set C

The three Cases I, II, and III correspond to three types of equilibria. The equilibrium supports of the two bidders in these three cases are given schematically in Figure 2:

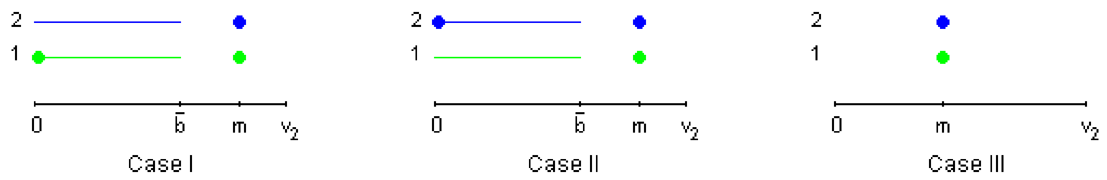

Figure 2: Equilibrium supports in the three cases

Cases C_I and C_{II}:

$(\alpha, m) \in C_I \cup C_{II}$ leads to a mixed equilibrium where both bidders mix over an interval $[0, \bar{b}]$ with $\bar{b} < m$. In addition both bidders place atoms on m and one bidder puts an atom on 0. Obviously, the bidder who sets the atom in 0 earns 0 in equilibrium. In Case I, this is the strong bidder, bidder 1, while in Case II, it is the weak bidder, bidder 2, who obtains zero payoff. On the boundary segment S, there is no atom in 0 and consequently, both bidders earn a payoff of 0. This observation alone reveals the power of the tie-breaking rule: For any $m \in (0, v_2)$, the designer can induce a positive payoff for either bidder by his choice of tie-breaking.

Since $C_I \cap C_G = \emptyset$, mixed equilibria where the weak bidder earns a positive payoff do not occur in Che and Gale's analysis. Yet such equilibria exist in our extended setting: Consider the tie-breaking rule $\alpha = 0$ which always favors bidder 2 and some $m \in (0, v_2)$. Bidder 2 can guarantee himself a positive payoff by submitting a bid of m regardless of his opponent's strategy. By the following lemma (which also proves useful for the later analysis), this implies that bidder 1 must obtain zero payoff.

Lemma 2 *Fix $(\alpha, m) \in C$ and fix an equilibrium. Denote by S_1 and S_2 the supports of the bidders' equilibrium strategies. Then*

$$\inf S_1 \cup S_2 \in \{0, m\}.$$

The lemma shows that one bidder must obtain zero payoff in equilibrium, unless we are in an equilibrium where both bidders play m. Yet for $\alpha = 0$, it is not an equilibrium that both bidders play m. (Bidder 1 would rather deviate to zero than play m in this case.) Hence it follows that bidder 1 obtains zero payoff. The set C_I is the region in C on which an extension of this reasoning goes through.

Case C_{III}:

The set $C_{III} \cup T$ corresponds to cases in which it is an equilibrium that both bidders choose m. Obviously, in such an equilibrium, the designer's payoff is given by $2m$. Thus it is easy to see that the designer can improve upon the configuration $\left(\frac{1}{2}, \frac{1}{2}v_2\right)$ which is optimal within C_G as shown by Che and Gale by choosing a smaller α and a larger m within the set $C_{III} \cup T$. Moreover, (α^*, m^*) which is the upper corner of the triangle C_{III}, explicitly

$$(\alpha^*, m^*) = \left(\frac{v_2}{v_1 + v_2}, \frac{v_1 v_2}{v_1 + v_2}\right) \in T,$$

is the optimal configuration within $C_{III} \cup T$ since it is the configuration in $C_{III} \cup T$ associated with the largest m.

3.2 Explicit Characterization

Starting with the following definition which characterizes explicitly the regions of the state space depicted in Figure 1, we now turn to a rigorous characterization of equilibrium.

Definition 1 *Define critical levels of m by*

$$m_1(\alpha) = \min(\alpha v_1, (1-\alpha)v_2) \text{ for } \alpha \in [0,1]$$

and

$$m_2(\alpha) = v_2 - \frac{\alpha^2}{1-2\alpha}(v_1 - v_2) \text{ for } \alpha \in \left[0, \frac{v_2}{v_1+v_2}\right].$$

$m_1(\alpha)$ separates Case III from Cases I and II, while $m_2(\alpha)$ separates Cases I and II. Furthermore, define

$$C_I = \left\{(\alpha, m) \in C \,\middle|\, \alpha \in \left[0, \frac{v_2}{v_1+v_2}\right) \text{ and } m_2(\alpha) > m > m_1(\alpha)\right\},$$

as well as

$$\begin{aligned} C_{II} &= \left\{(\alpha, m) \in C \,\middle|\, \alpha \in \left[0, \frac{v_2}{v_1+v_2}\right) \text{ and } v_2 > m > m_2(\alpha)\right\} \\ &\cup \left\{(\alpha, m) \in C \,\middle|\, \alpha \in \left[\frac{v_2}{v_1+v_2}, 1\right] \text{ and } v_2 > m > m_1(\alpha)\right\}, \end{aligned}$$

and

$$C_{III} = \{(\alpha, m) \in C \,|m < m_1(\alpha)\}.$$

Denote by S the segment of C separating C_I and C_{II},

$$S = \left\{(\alpha, m) \in C \,\middle|\, \alpha \in \left[0, \frac{v_2}{v_1+v_2}\right) \text{ and } m = m_2(\alpha)\right\},$$

by T the segment of C separating C_I and C_{II} from C_{III},

$$T = \{(\alpha, m) \in C \,|\alpha \in [0,1] \text{ and } m = m_1(\alpha), \}$$

and by U the upper boundary,

$$U = [0,1] \times \{v_2\}.$$

Clearly, $C = C_I \cup C_I \cup C_{II} \cup C_{III} \cup S \cup T \cup U$ and this conjunction is disjoint. Finally, denote by C_G the region in C analyzed by Che and Gale (1998),

$$C_G = \left\{ \frac{1}{2} \right\} \times (0, v_2].$$

The following proposition provides a detailed characterization of equilibrium in the cases $C_I \cup C_{II} \cup C_{III}$.

Proposition 1 *Denote by π_i the equilibrium payoff of bidder i, by c_i the atom in 0 played by bidder i and by d_i the atom bidder i places on m. Recall that $v_1 > v_2$. For each $(\alpha, m) \in C \setminus (S \cup T \cup U)$ the unique equilibrium is characterized as follows:*

(i) For $(\alpha, m) \in C_I$, we have $\pi_1 = 0$ and

$$\pi_2 = v_2 - \left(\frac{\alpha^2}{(1-\alpha)^2} v_1 + \left(1 - \frac{\alpha^2}{(1-\alpha)^2} \right) m \right).$$

Bidder 1 puts atoms of sizes

$$c_1 = \frac{\pi_2}{v_2} \qquad \text{and} \qquad d_1 = \frac{v_1 - m}{v_2} \frac{\alpha}{(1-\alpha)^2}$$

on 0 and m and mixes uniformly over $[0, \bar{b}]$ with his remaining probability mass, where

$$\bar{b} = \frac{m - \alpha v_1}{1 - \alpha}.$$

Bidder 2 puts an atom of size

$$d_2 = \frac{v_1 - m}{v_2} \frac{1}{1-\alpha}$$

on m and mixes uniformly over $[0, \bar{b}]$ with his remaining mass.

(ii) For $(\alpha, m) \in C_{II}$, we have $\pi_2 = 0$ and

$$\pi_1 = v_1 - \left(\frac{(1-\alpha)^2}{\alpha^2} v_2 + \left(1 - \frac{(1-\alpha)^2}{\alpha^2} \right) m \right).$$

Bidder 2 puts atoms of sizes

$$c_2 = \frac{\pi_1}{v_1} \qquad \text{and} \qquad d_2 = \frac{v_2 - m}{v_2} \frac{1-\alpha}{\alpha^2}$$

on 0 *and* m *and mixes uniformly over* $[0, \overline{b}]$ *with his remaining mass, where*

$$\overline{b} = \frac{m - (1-\alpha)v_2}{\alpha}.$$

Bidder 1 *sets an atom of size*

$$d_1 = \frac{v_2 - m}{v_2}\frac{1}{\alpha}$$

on m *and mixes uniformly over* $[0, \overline{b}]$ *with his remaining mass.*

(iii) For $(\alpha, m) \in C_{III}$*, both bidders play* m *with probability* 1*. Accordingly,*

$$\pi_1 = \alpha v_1 - m \qquad \textit{and} \qquad \pi_2 = (1-\alpha)v_2 - m.$$

We now discuss the boundary cases S, T, and U.[7] For $(\alpha, m) \in S$, the equilibria of Cases (i) and (ii) of the proposition coincide and accordingly the atoms in zero vanish. The resulting vector of strategies is a Nash equilibrium.

Observe that C_{III} is exactly the set where the payoffs π_1 and π_2 from Case (iii) of the proposition are positive: Then, no bidder has an incentive to deviate to a lower bid. At the upper boundary, i.e., for $(\alpha, m) \in T$, there is typically a continuum of equilibria where the bidder who earns zero payoff mixes over $\{0, m\}$ with sufficiently large but otherwise arbitrary mass on m. The only exception is the case

$$(\alpha^*, m^*) = \left(\frac{v_2}{v_1 + v_2}, \frac{v_1 v_2}{v_1 + v_2}\right) \in T$$

where both bidders earn zero payoff. In this case, it is the unique equilibrium that both bid $\frac{v_1 v_2}{v_1 + v_2}$. It can be shown that for $t = (\alpha, m) \in T$ the set of Nash equilibria is the set of convex combinations of $\underline{E}_t$ and $\overline{E}_t$ where $\underline{E}_t$ and $\overline{E}_t$ are the limits as $\varepsilon \downarrow 0$ of the unique equilibria obtained, respectively, for $(\alpha, m - \varepsilon)$ and $(\alpha, m + \varepsilon)$.[8]

For $(\alpha, m) \in U$, with $\alpha > 0$ the unique Nash equilibrium is just the unique equilibrium of the complete information all-pay auction without bid-cap. For $(\alpha, m) = (0, v_2)$ there is a continuum of Nash equilibria, where both bidders mix uniformly over $[0, v_2)$. Additionally, bidder 2 places atoms c_2 and d_2 on 0 and on $m = v_2$ with

[7] While one may find the arising non-uniquenesses of equilibrium theoretically interesting, we feel that the following informal discussion is sufficient since only boundary cases are concerned.

[8] Erroneously (and in conflict with their own uniqueness result), Che and Gale claim that a similar non-uniqueness of equilibrium applies to the case $(\frac{1}{2}, \frac{v_1}{2})$.

$c_2 + d_2 = 1 - \frac{v_2}{v_1}$. Clearly, the case $d_2 = 0$ corresponds to the well-known equilibrium of the complete information all-pay auction. This equilibrium is obviously the limit as ε goes to zero of the unique equilibrium obtained for $(\varepsilon, v_2) \in C$. At the other extreme lies the equilibrium with $c_2 = 0$. It can easily be shown that this equilibrium is the limit of the unique equilibria obtained for $(0, v_2 - \varepsilon)$.

Thus – as we will discuss in more detail below – the combination of an m slightly below v_2 and $\alpha = 0$, a restriction which on the surface minimally perturbs the game, induces an equilibrium which differs from the usual equilibrium of the complete-information all-pay auction in one important respect: The ex ante weaker bidder makes preemptive bids of m with the probability with which he lays out in the unrestricted all-pay auction. Obviously, this is highly interesting for a designer wishing to maximize bids.

We close our characterization of equilibrium by discussing one surprising finding of Che and Gale's analysis in light of the broader picture of equilibrium we have derived: Che and Gale show that in their mixed strategy case, $C_G \cap C_{II}$, the bid-cap does not affect the bidders' equilibrium payoffs. Indeed, if α is fixed at $\frac{1}{2}$, the expressions for π_1 and π_2 from Case (ii) of Proposition 1 become

$$\pi_1 = v_1 - v_2 \qquad \text{and} \qquad \pi_2 = 0.$$

These are the well-known equilibrium payoffs from the all-pay auction without bid-caps, which are of course independent of m. A similar independence of m does not arise for any other tie-breaking $\alpha \neq \frac{1}{2}$. In fact, it is easy to see that within C_{II}, π_1 is increasing in m for $\alpha > \frac{1}{2}$ and decreasing for $\alpha < \frac{1}{2}$.

4 The Design Problem

We now analyze the designer's problem of setting α and m in a way that maximizes the sum of the bidder's bids or, equivalently, the average bid. We refer to the expected sum of bids as the designer's payoff in the following. First we analyze the designer's optimization problem. Then we ask what the designer's optimal policy is if he is restricted to choosing α from $\{0, \frac{1}{2}, 1\}$. The latter analysis is motivated by the fact that in practice we often observe symmetric tie-breaking or tie-breaking completely in favor of the weaker bidder. We suppose that more complex tie-breaking rules may be more difficult to implement and more difficult to understand.

We find that in both optimization problems, the designer prefers to favor the ex ante weaker bidder in tie-breaking, though to a different extent. Moving away from symmetric tie-breaking can increase payoffs by at most 100% if the designer can choose any α and by at most 50% if the designer is restricted to $\alpha \in \{0, \frac{1}{2}, 1\}$. The larger the asymmetry between the bidders, the larger is the gain from asymmetric tie-breaking.

Following Che and Gale (1998), we asume that the equilibrium played is the one most favorable to the designer. In this equilibrium there are no atoms at 0. The selection is relevant for $(\alpha, m) \in T$ with the exception of the tip of the triangle and the upper left corner $(\alpha, m) = (0, v_2)$, compare Figure 1. This selection is the limit of the equilibria for slightly lower m. The selection is important only when α is in $\{0, \frac{1}{2}, 1\}$.

Afterwards, we give a quantitative payoff comparison for the different policies that turned out to be favorable in the different settings. In the final subsection, we have a brief look at the problem of *minimizing* expected efforts in the tie-breaking rule for a fixed bid-cap. It turns out that when ties are broken in favor of the ex-ante stronger bidder, total efforts are decreasing monotonically in the bid-cap.

4.1 Global Optimization

The global optimization problem is complicated by the fact that – mainly due to the transition between pure and mixed equilibrium – the sum of bids which we denote by σ in the following is discontinuous on T and non-differentiable on S. Thus a rather piece-wise approach to optimization is in order.

Observe first that for $(\alpha, m) \in C_{III} \cup T$ we have

$$\sigma(\alpha, m) = 2m.$$

Thus the optimal (α, m) in this region is given by the tip (α^*, m^*) of this triangle since this is associated with the largest value of m:

$$(\alpha^*, m^*) = \left(\frac{v_2}{v_1 + v_2}, \frac{v_1 v_2}{v_1 + v_2} \right)$$

The notation (α^*, m^*) is chosen in anticipation of the fact that this will turn out to be the global optimum. Now we analyze $\sigma(\alpha, m)$ in the set $C_I \cup C_{II} \cup S \cup U$. Following the notation of Proposition 1, $\sigma(\alpha, m)$ is easily seen to be given by

$$\sigma(\alpha, m) = \sum_{i=1}^{2} (1 - c_i - d_i)\frac{\bar{b}}{2} + d_i m \tag{1}$$

Our first result shows that the monotonicity behavior of $\sigma(\alpha, m)$ is surprisingly complex:

Lemma 3 *For $(\alpha, m) \in C_I$ we have $\frac{d\sigma(\alpha,m)}{dm} > 0$, for $(\alpha, m) \in C_{II}$ with $\alpha < \frac{1}{2}$ we have $\frac{d\sigma(\alpha,m)}{dm} < 0$, and for $(\alpha, m) \in C_{II}$ with $\alpha > \frac{1}{2}$ we have $\frac{d\sigma(\alpha,m)}{dm} > 0$.*

Thus increasing the bid-cap is beneficial for the designer in the region C_I and in the part of C_{II} which lies to the right of C_G. Between these regions, increasing the bid-cap has an adverse effect on the sum of bids. This result together with a few additional observations allows us to substantially restrict the region in which we have to search for a global optimum:

Corollary 1 *If (α, m) is a global maximizer of $\sigma(\alpha, m)$ then $(\alpha, m) \in \overline{S}$.*

Here, $\overline{S}$, the closure of S, is nothing other than S together with its two end points, $(0, v_2)$ and (α^*, m^*). From here it is a straightforward calculation to the main result of this section:

Proposition 2 *For all $m \geq 0$ and $\alpha \in [0, 1]$ where $(\alpha, m) \neq (\alpha^*, m^*)$ we have*

$$\sigma(\alpha^*, m^*) > \sigma(\alpha, m).$$

4.2 Simple Tie-Breaking Rules

We now examine the restricted setting where the designer can only choose between symmetric and deterministic tie-breaking rules: $\alpha \in \{0, \frac{1}{2}, 1\}$. We start with a discussion of the options the designer has under this restriction: The case $\alpha = \frac{1}{2}$ has been studied by Che and Gale. In this case, setting $m = \frac{v_2}{2}$ is optimal and it leads to a payoff of v_2 for the designer. Allowing the designer to choose $\alpha = 1$ is actually irrelevant: As shown in Lemma 3, the expected sum of bids is increasing in m for

$\alpha = 1$. For $m = v_2$, the designer obtains the payoff from the unrestricted all-pay auction which is given by

$$\frac{v_2}{2} + \frac{v_2}{v_1}\frac{v_2}{2} < v_2.$$

Now consider the case $\alpha = 0$: Again by Lemma 3, the expected sum of bids is increasing in m. Yet setting $m = v_2$ does not yield the payoff of the unrestricted all-pay auction in this case: Instead of putting an atom on zero, bidder 2 makes a preemptive bid of v_2. This leads to a payoff of

$$\frac{v_2}{2} + \frac{v_2}{v_1}\frac{v_2}{2} + (1 - \frac{v_2}{v_1})v_2 > v_2.$$

We therefore conclude:[9]

Corollary 2 *If the designer can choose $\alpha \in \{0, \frac{1}{2}, 1\}$ and $m \geq 0$ then his optimal choice is given by $(\alpha, m) = (0, v_2)$.*

Figure 3 depicts the expected sum of bids for $\alpha \in \{0, \frac{1}{2}, 1\}$ and $\alpha = \alpha^*$ as a function of m:

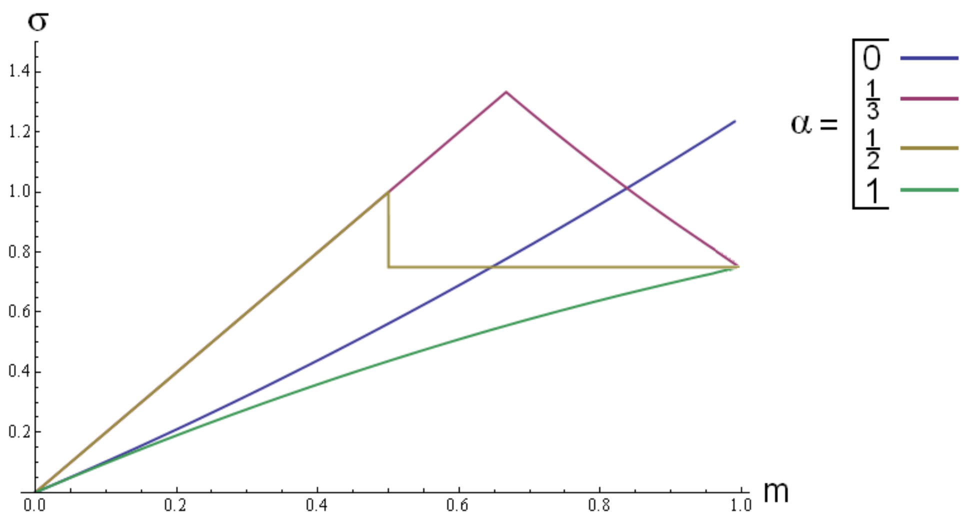

Figure 3: Expected sum of bids for $v_1 = 2$, $v_2 = 1$ and $\alpha \in \{0, \alpha^*, \frac{1}{2}, 1\}$ as a function of m where $\alpha^* = \frac{v_2}{v_1+v_2} = \frac{1}{3}$.

[9]Recall that the optimal policies for $\alpha = 0$ and $\alpha = \frac{1}{2}$ rely on equilibria which are not unique. To make the desired equilibria unique, the designer has to choose m marginally smaller than indicated.

4.3 Quantitative Comparison of Policies

To close our analysis of the designer's optimization problem under linear bidding costs, we compare the maximal differences between four policies which have turned out to be important in the previous discussion: Let P^* be the globally optimal policy of setting (α^*, m^*). Let P^R be the policy $(0, v_2)$ which was optimal in the restricted decision problem of the previous section: breaking ties in favor of the weaker bidder and setting the minimal bid-cap which has an effect. let P^{CG} be the policy $(\frac{1}{2}, \frac{v_2}{2})$ which is optimal in the setting of Che and Gale. Finally let P^0 be the policy of doing nothing and letting bidders play a complete information all-pay auction. Denote the designer's payoffs under those for policies by σ^*, σ^R, σ^{CG} and σ^0. As shown above, we have

$$\sigma^* = \frac{2v_1v_2}{v_1 + v_2}, \quad \sigma^R = \frac{v_2}{2} + \frac{v_2}{v_1}\frac{v_2}{2} + (1 - \frac{v_2}{v_1})v_2,$$

and

$$\sigma^{CG} = v_2, \quad \text{and} \quad \sigma^0 = \frac{v_2}{2} + \frac{v_2}{v_1}\frac{v_2}{2}.$$

Considering the limits $v_1 \uparrow \infty$ and $v_1 \downarrow v_2$ leads to especially clean results:

Lemma 4 *We have*

$$\lim_{v_1 \uparrow \infty} \begin{pmatrix} \sigma^* \\ \sigma^R \\ \sigma^{CG} \\ \sigma^0 \end{pmatrix} = \begin{pmatrix} 2 \\ \frac{3}{2} \\ 1 \\ \frac{1}{2} \end{pmatrix} v_2.$$

and

$$\lim_{v_1 \downarrow v_2} \sigma^* = \lim_{v_1 \downarrow v_2} \sigma^R = \lim_{v_1 \downarrow v_2} \sigma^{CG} = \lim_{v_1 \downarrow v_2} \sigma^0 = v_2$$

A straightforward calculation reveals that the cases treated in the Lemma correspond to the extreme cases, showing that the optimal policy is at most twice as good as the policy of Che and Gale and at most four times as good as the unrestricted all-pay auction. The policy P^R gives a substantial advantage over P^{CG}. Moreover, we see that policies P^* and P^R manage to exploit large values of v_1: σ^* and σ^R are increasing in v_1. In contrast, σ^{CG} is independent of v_1 and σ^0 is decreasing in v_1. The fact that the difference in the strengths of the two bidders increases overshadows the fact that - in the sum - bidders are willing to bid more. The observation about P^0 is well-known: It is the basis of the result of Baye, Kovenock and de Vries (1993) that excluding high-valuation bidders from a complete information all-pay auction may be beneficial to the designer's payoff if it reduces asymmetries. In this perspective, Che and Gale have shown that setting a bid-cap with symmetric tie-breaking is just as valuable to the designer as being able to replace the stronger bidder with a new bidder who is just as good as the weaker bidder. Our policies with asymmetric tie-breaking P^* and P^R actually exploit the fact that bidder 1 is willing to bid more than bidder 2. So to say, in the limit $v_1 \uparrow \infty$, P^* leads to the same payoffs as a symmetric all-pay auction where both bidders have valuations $2v_2$.

Finally, note that while the payoffs of the four policies coincide as $v_1 \downarrow v_2$, the policies themselves do not coincide: As $v_1 \downarrow v_2$, policies P^* and P^{CG} converge to an auction where both bidders bid $\frac{v_2}{2}$ for sure and tie-breaking is symmetric. P^R and P^0 both converge to a symmetric complete-information all-pay auction where tie-breaking is irrelevant since there are no atoms in the symmetric case. In this light, it is especially interesting to note that the payoff ranking is always $P^* \succ P^R \succ P^{CG} \succ P^0$.

4.4 Reducing the sum of bids

Che and Gale (1998)'s initial motivation for considering bid-caps in political lobbying is that bid-caps are often imposed in an attempt to reduce lobbying efforts. In the light of their analysis, it seems that bid-caps are hardly suitable for this purpose since they may in fact lead to higher effort exertion than without the cap. This leads us to the question whether bid-caps coupled with a suitable tie-breaking rule may indeed be useful for controlling effort exertion.[10]

A simple calculation shows that for a fixed bid-cap m setting $\alpha = 1$ minimizes the expected sum of bids. For any $m \in (0, v_2)$ this leads to a reduction in the sum of bids. In fact, coupled with tie-breaking in favor of the stronger bidder, bid-caps have just the effect which one might naively expect from a bid-cap: Compared to the unrestricted auction, bidding on $(0, m)$ is unaffected. The stronger bidder moves mass from (m, v_2) into an atom in m. The weaker bidder increases his atom in 0 by the mass he previously had in (m, v_2). This shows that the optimal tie-breaking strategy strongly depends on the designer's goal. Moreover, introducing bid-caps in an all-pay auction makes the tie-breaking rule a highly flexible design instrument.

[10] Note that if the designer actually wanted to minimize the sum of bids, he could simply set a bid-cap of zero and an arbitrary tie-breaking rule. Typically, caps on political lobbying are not set to zero even if they are supposed to reduce effort exertion.

5 General Cost Functions

The analysis of the previous section identified two types of strategies by which the designer could use bid-caps and tie-breaking rules to boost his payoff: Under the policy P^R, the designer broke ties in favor of the weak bidder and made a bid-cap that was just small enough to matter. This forced the weak bidder to make preemptive bids with some probability instead of laying out with the same probability. Under policies P^* and P^{CG}, the designer made a more rigid bid-cap, forcing bidders to play a pure equilibrium. In this section, we study how well these strategies perform in a more general setting where we dispense with the assumption of linear bidding costs.

We find that P^R is considerably more robust than policies that enforce a pure strategy equilibrium. Unlike those policies, P^R always yields an improvement over the unrestricted all-pay auction. It depends remarkably little on the details of the cost function in the sense that the payoff difference to the unrestricted all-pay auction only depends on the value of the cost functions at the upper boundary of the equilibrium support. For strongly concave cost functions, the relative improvement over the unrestricted all-pay auction and over the other class policies gets arbitrarily large. In contrast, for sufficiently convex cost-functions, the policies that enforce pure equilibria are better than P^R.[11]

In Section 6.1 below, we observe that under general cost-functions the knowledge requirements on a designer enforcing P^R are considerably smaller than those on a designer who wishes to enforce the optimal pure equilibrium. Together, these observations suggest that from practical point of view policy P^R is highly attractive.

For the purpose of the following discussion, we slightly redefine the policies we consider: We denote by P^* the policy which induces the designer-optimal *pure* strategy equilibrium. Analogously, we denote by P^{CG} the policy associated with designer-optimal *pure* strategy equilibrium available under symmetric tie-breaking.[12] P^R now denotes the policy of setting the minimal bid-cap which has an impact and always breaking ties in favor of the weaker bidder. P^0 denotes again the policy of playing an unrestricted all-pay auction.

[11] See the discussion below: In this context, by "convexity" we mean that cost functions take most of their increase in the bidding interval at high bids. It does not refer to a local property. The same is the case for "concavity".

[12] Note that - as we will see below - these two policies may be dominated by the unrestricted all-pay auction for some cost functions. Thus, P^* is not necessarily a global optimum anymore.

5.1 The Model with General Cost Functions

Consider the following generalized model: Bidder i has costs $c_i(b)$ from making a bid of b where c_i is a continuous, strictly increasing function with $c_i(0) = 0$ and with $c_i(b) > v_i$ for some $b < \infty$. We drop our assumption that $v_1 > v_2$ but we still assume that bidder 1 is the stronger bidder: We assume there exists a $t > 0$ such that

$$v_1 - c_1(t) = p_1 > 0 \quad \text{and} \quad v_2 - c_2(t) = 0.$$

Clearly, t is the largest bid bidder 2 might be willing to make.[13] Siegel (2010, Theorem 3) has shown that in this case the unique equilibrium of the all-pay auction without bid-cap is given by bidders 1 and 2 mixing with the distribution function

$$F_1(b) = \frac{c_2(b)}{v_2} \quad \text{and} \quad F_2(b) = \frac{p_1 + c_1(b)}{v_1} \tag{2}$$

on $[0, t]$. Note that bidder 2 places an atom of size $\frac{p_1}{v_1} = \frac{v_1 - c_1(t)}{v_1}$ on 0. Thus, bidder 2 makes a payoff of zero.

5.2 Policy P^0

We consider first P^0, the policy of letting bidders play a complete information all-pay auction without bid-cap. We denote by σ^0 the associated payoff to the designer. An easy calculation gives us a formula for the designer's equilibrium payoffs, i.e., the sum of bids in this case:

Corollary 3 *The designer's payoff in the unrestricted all-pay auction is given by*[14]

$$\sigma^0 = \sum_{i=1}^{2} \frac{t}{v_i}(c_i(t) - \bar{c}_i(t))$$

where for $b \in [0, t]$, $\bar{c}_i(t)$ *denotes the average costs on* $[0, b]$,

$$\bar{c}_i(b) = \frac{1}{b}\int_0^b c_i(x)dx.$$

[13] In the terminology of Siegel (2010), t is the so-called threshold, defined as the reach of the bidder with the smaller reach. The reach is the highest bid a bidder can make without earning a negative payoff regardless of his opponents strategy. p_1 is the so-called power of bidder 1.

[14] To our knowledge, this compact expression for equilibrium efforts is not common in the literature and may be of minor independent interest: The designers' payoff depends on the cost functions only through the differences between costs at the threshold and average costs.

5.3 Policies P^* and P^{CG}

It is clear that there is still a region C_{III} of the state space C where a pure strategy equilibrium arises in which both bidders play m. Define the functions h_1 and h_2 by

$$h_1(\alpha) = c_1^{-1}(\alpha v_1) \quad \text{and} \quad h_2(\alpha) = c_2^{-1}((1-\alpha)v_2)$$

where c_i^{-1} denotes the inverse of c_i which is well-defined since c_i is strictly increasing. Then it is easy to see that C_{III} is given by

$$C_{III} = \{(\alpha, m) \in C | \min(h_1(\alpha), h_2(\alpha)) < m\}.$$

Define again T to be the upper boundary of C_{III}, i.e.

$$T = \{(\alpha, m) \in C | \min(h_1(\alpha), h_2(\alpha)) = m\}.$$

To sum up, a pure equilibrium arises, if m is such that αv_1 is greater than $c_1(m)$ and if $(1-\alpha)v_2$ is greater than $c_2(m)$.

With general cost-functions, C_{III} does typically not have the triangular shape we saw with linear costs. However, since its boundary is the minimum of the increasing function $h_1(\alpha)$ and the decreasing function $h_2(\alpha)$, the boundary will increase up to some point and decrease from there on. Accordingly, we define (α^*, m^*) to be the unique pure strategy equilibrium associated with the maximal bid-cap m^* that is associated with a pure equilibrium. Accordingly, P^* defined above is the policy of implementing (α^*, m^*). Denote by σ^* the resulting payoff of the designer. Analogously, there is a unique maximal bid-cap m^{CG} which leads to a pure strategy equilibrium for $\alpha = \frac{1}{2}$. P^{CG} and σ^{CG} are the associated policy and designer's payoff. Clearly, we have $\sigma^* = 2m^*$ and $\sigma^{CG} = 2m^{CG}$.

Except in generic cases where $\alpha^* = \frac{1}{2}$, P^* is strictly superior to P^{CG}. In the case where $c_1 \equiv c_2$, it is easy to see that just as in the linear case

$$\alpha^* = \frac{v_2}{v_1 + v_2}.$$

Generally however, α^* may lie anywhere in the interval $[0, 1]$. More importantly, the values of m^* and m^{CG} may lie anywhere in the interval $[0, t]$: If *both* cost functions decrease slowly at low bids and increase up to $c_1(t)$ and $c_2(t)$ only on a short interval below t, then m^* and m^{CG} can lie arbitrarily close to t. Conversely, if at least one of the cost functions increases strongly already for small bids, m^* and m^{CG} can be arbitrarily close to zero.[15]

5.4 Policy P^R

Next we turn to $C \setminus (C_{III} \cup T)$, the cases where no pure equilibria exist. It is easy to check that the arguments which gave us the structural properties of the equilibrium in the linear case still go through in the present setting: There are no atoms in $(0, m)$ and both bidders mix over the same interval which goes down to zero. Additionally, at most one bidder places an atom on zero and for $\alpha \in (0, 1)$ both place an atom on m. Also for $\alpha \in (0, 1)$, there is an interval below m on which bidders are inactive. Furthermore, by choosing $\alpha \in \{0, 1\}$, the designer can force either of the two bidders to earn zero payoff, implying that there should again be a boundary segment S connecting the points $(0, t)$ and (α^*, m^*) such that - in the logic of Figure 1 - bidder 2 yields a positive payoff to the left of S and bidder 1 yields a positive payoff to the right of S.[16]

Does the designer still have a policy P^R which forces the weaker bidder 2 to exchange his atom on zero (which by (2) he plays in the unrestricted auction) against an atom on t? The following lemma shows that this is indeed the case:

Lemma 5 *For $\alpha = 0$ and $m \in (0, t)$, the unique equilibrium is given by*

$$F_1(b) = \frac{v_2 - c_2(m)}{v_2} + \frac{c_2(b)}{v_2} \quad \textit{and} \quad F_2(b) = \frac{c_1(b)}{v_1}$$

for $b \in [0, m)$ and $F_1(m) = F_2(m) = 1$. Accordingly, bidder 2's strategy has an atom of size $\frac{v_1 - c_1(m)}{v_1}$ in m while bidder 1's strategy has an atom of size $\frac{v_2 - c_2(m)}{v_2}$ in 0.

Considering the equilibrium from Lemma 5 in the limit $m \uparrow t$, gives the vector of strategies

$$F_1(b) = \frac{c_2(b)}{v_2} \quad \text{and} \quad F_2(b) = \frac{c_1(b)}{v_1}$$

for $b \in [0, t)$ and an atom of size $\frac{v_1 - c_1(t)}{v_1}$ by bidder 2 in t. It is easy to see that this is indeed an equilibrium, and that it differs from the one in (2) only in one

[15] See the example of Section 5.5.

[16] A more detailed study of mixed equilibria in this general case goes beyond the scope of this paper. We conjecture that Corollary 1 still holds in this setting, i.e., the designer's globally optimal policy must lie somewhere in the segment $\overline{S}$ such that both bidders earn zero payoffs.

respect: Bidder 2 has moved his atom from zero to t. Thus there is again a non-uniqueness of equilibrium in $(0,t)$. Moreover, we see that - just as in the case of linear costs - by setting $\alpha = 0$ and m marginally below t, the designer generates a unique equilibrium where bidder 2 uses his former drop-out probability to make preemptive bids. Consequently, we call this policy of the designer P^R and denote its payoff to the designer by σ^R. It follows easily from Corollary 3 that

$$\begin{aligned}\sigma^R &= \sigma^0 + (v_1 - c_1(t))\frac{t}{v_1} \\ &= \frac{t}{v_1}(v_1 - \overline{c}_1(t)) + \frac{t}{v_2}(v_2 - \overline{c}_2(t)).\end{aligned}$$

Unlike P^* and P^{CG}, P^R is always an improvement over the all-pay auction without bid-cap P^0. Moreover, P^R exhibits a similar robustness towards the details of the cost function as P^0: It depends only on average costs and on costs at the top. The difference between σ^R and σ^0 even depends only on the costs in t.

5.5 An Example

To illustrate the observations of this section, we consider the case $c_1(b) = c_2(b) = b^\gamma$ for some $\gamma > 0$, $v_1 > v_2 = 1$.[17] Then for all γ, we have $t = 1$. It is straightforward to obtain

$$m^* = \left(\frac{v_1}{v_1+1}\right)^{\frac{1}{\gamma}} \quad \text{and} \quad m^{CG} = \left(\frac{1}{2}\right)^{\frac{1}{\gamma}}.$$

For $\gamma \downarrow 0$, i.e., if costs become more and more concave, we have $\sigma^* = 2m^* \downarrow 0$ and $\sigma^{CG} = 2m^{CG} \downarrow 0$. Conversely, for $\gamma \uparrow \infty$ when costs become more and more convex, we have $\sigma^* = 2m^* \uparrow 2$ and $\sigma^{CG} = 2m^{CG} \uparrow 2$. Note that 2 is the maximal payoff the designer can hope for in this setting: With a bid-cap m he can never earn more than $2m$ and a bid-cap $m > t = 1$ would have no bite.

How does this compare to the unrestricted all-pay auction? From Lemma 3, we obtain that in this example

$$\sigma^0 = \left(1 + \frac{1}{v_1}\right)\frac{\gamma}{\gamma+1}$$

which converges to zero considerably more slowly than σ^* and σ^{CG} as $\gamma \downarrow 0$. In

[17] The choice $v_2 = 1$ is crucial to make this example well-behaved: It ensures that t remains fixed when we vary γ, making a comparison across different values of γ reasonable.

consequence, for small values of γ and thus a strongly concave cost function, the unrestricted all-pay auction policy P^0 outperforms P^* and P^{CG}.[18] For $\gamma \uparrow \infty$, σ^0 converges to $1 + 1/v_1 < 2$. Note that the intuition behind the example does not rely on the specific choice of cost function. We thus see that the performance of policies P^* and P^{CG} is rather sensitive to the details of the cost function.

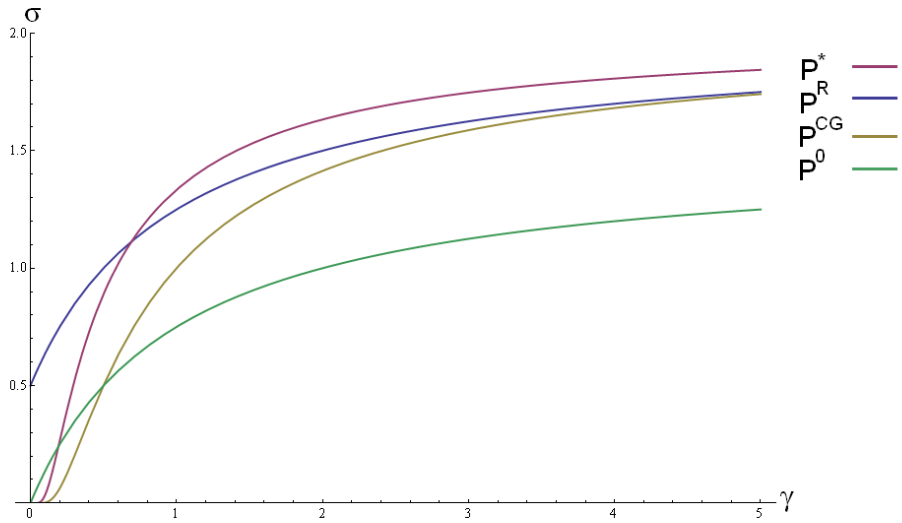

Figure 4: Expected sum of bids under the different policies for $v_1 = 2$, $v_2 = 1$ and $c_1(b) = c_2(b) = b^\gamma$ as a function of γ.

Let us now compare P^R to P^* and P^{CG}. By (3),

$$\sigma^R = \sigma^0 + 1 - \frac{1}{v_1} = 2 - \frac{v_1 + 1}{v_1} \frac{1}{1 + \gamma}.$$

For $\gamma \uparrow \infty$, σ^R converges to 2 just like σ^* and σ^{CG}. For $\gamma \downarrow 0$, we have $\sigma^R \downarrow 1 - \frac{1}{v_1}$. Thus, P^R is the only of the policies under consideration which yields a strictly positive payoff even as costs become arbitrarily concave: The size of the atom in t played by bidder 2 only depends on t, on v_1 and on $c_1(t)$ - it is independent of the shape of the cost function within $(0, t)$. It is not difficult to find examples where both P^* and P^{CG} outperform P^R.[19] Nevertheless, we see that P^R is a surprisingly robust policy. Figure 4 summarizes the results of this example.

[18] These observations may be seen as distant cousins of the results of Gavious, Moldovanu and Sela (2002), who show that bid-caps are profitable only with sufficiently convex cost functions in the symmetric incomplete-information case with symmetric tie-breaking.

[19] Take for instance the present example with $v_1 = 2$ and $\gamma = 10$.

6 Discussion

In the following, we discuss robustness, extensions and implications of our analysis. Specifically, we address the following questions: In Section 6.1, we discuss the knowledge requirements of the different policies, showing that P^R has generally less rigid knowledge requirements than the other policies. Sections 6.2 and 6.3 briefly address the symmetric case $v_1 = v_2$ and the case of more than two bidders. In Section 6.5, we discuss the problem of a designer who wants to maximize the winning bid instead of the sum of bids. We show that (unlike P^* and P^{CG}) policy P^R leads to an improvement over the unrestricted all-pay auction also under this objective. Finally, Sections 6.6 and 6.7 relate our results to the literature on the incomplete information case and to the literature on "soft" bid-caps.

6.1 Knowledge Requirements

How much does the designer need to know in order to apply the policies P^*, P^R and P^{CG} introduced above? In the case of linear costs, there is a clear ranking in the knowledge requirements: For P^{CG} the designer needs to know only the smaller valuation, for P^R the designer also needs to know whose valuation this is, for P^* he needs to know both valuations. Yet these are only small differences in knowledge requirements. Moreover, the smaller requirements of P^{CG} are offset by the fact that P^R and P^* lead to higher payoffs. Both P^R and P^{CG} lead to the payoffs of the complete information all-pay auction if the designer chooses m too large. It is easy to show that for any estimate $\hat{v}_2$ of v_2, implementing P^R with $m = \hat{v}_2$ weakly dominates implementing P^{CG}. Likewise, one can show that perturbing P^{CG} into the direction of P^* by biasing tie-breaking in favor of the weaker bidder and increasing the bid-cap works well if the bid-cap is not increased by too much.

Comparing knowledge requirements in the case of non-linear costs leads to a markedly different picture: For P^R the designer needs to know only the highest possible equilibrium bid t of the weaker bidder and the weaker bidder's identity. Note that observing a number of unrestricted all-pay auctions between bidders 1 and 2 (who do not know they are observed) should give the designer a good proxy for t (which is simply the upper bound of the equilibrium support in this case). In contrast, for implementing P^{CG} and P^*, the designer has to determine, respectively, $c_i^{-1}(\frac{1}{2}v_i)$ and $c_i^{-1}(\alpha_i v_i)$ for both bidders. Thus the designer needs to know the bidders' cost functions at certain intermediate values which is a rather strong requirement.

6.2 The Symmetric Case

One small gap left by our above analysis may deserve a few words: We only considered the case $v_1 > v_2$ but left our the symmetric case $v_1 = v_2 = v$. In this boundary case, bid-caps and asymmetric tie-breaking rules are at best without harm: There are no ex ante differences which could be alleviated in order to strengthen competition. The general structure of the equilibrium is not changed under symmetry: There are still three regions, one with a pure equilibrium and two with mixed equilibria in which either of the bidders earns a positive payoff. Denote again by (α^*, m^*) the combination of tie-breaking rule and bid-cap which leads to the most profitable pure equilibrium. It is easy to see that $(\alpha^*, m^*) = (\frac{1}{2}, \frac{v}{2})$. The main difference to the asymmetric case is that the boundary S between the two mixed cases has a different shape: It no longer connects (α^*, m^*) with one of the corners. Instead, it is symmetric and connects $(\alpha^*, m^*) = (\frac{1}{2}, \frac{v}{2})$ with $(\frac{1}{2}, v)$.

6.3 More Than Two Bidders

An obvious question is how our analysis extends to the n-bidder case. We believe that the focus on two bidders is justified for two reasons: In many situations of competition, two bidders are the most typical case. More importantly, at least under linear (or ordered[20]) costs, only the two strongest bidders actively bid in a complete information all-pay auction. While very rigid bid-caps may encourage participation of further bidders, it is rather intuitive that the designer obtains the best results by focusing on the two strongest bidders, which allows him to set a much less rigid bid-cap and hence to obtain higher outcomes. See also the discussion in Che and Gale (1998).

6.4 Asymmetric Bid-Caps

One might wonder what happens if the designer utilizes not only asymmetric tie-breaking but also asymmetric bid-caps m_1 and m_2. Such asymmetric bid-caps should not have a huge effect: Intuitively, the bidder with the less restrictive bid-cap should not have an incentive to submit bids in ranges where the other bidder cannot compete. Note that technically, unless tie-breaking always favors the bidder with the less restrictive bid-cap, an equilibrium typically does not exist: The bidder who has the possibility will always out-bid the opponent's bid-cap by an arbitrarily small

[20] See Siegel (2009).

amount to bias tie-breaking in his favor. Finally, observe that the designer can enforce policy P^R with an arbitrary tie-breaking rule α by forcing the stronger bidder 1 to make bids *strictly* smaller than v_2 while leaving bidder 2's bidding unrestricted.

6.5 Maximizing the Winning Bid

In some applications, the designer may be interested in maximizing the expected winning bid instead of the expected sum of bids. In such a situation, the designer can improve upon the unrestricted all-pay auction policy P^0 by setting a bid-cap and a suitable tie-breaking rule: For all cost-functions we considered, policy P^R induces higher bids from the weaker bidder without having an impact on the stronger bidder's behavior. Thus the designer can always improve upon the unrestricted all-pay auction. Therefore, P^R is also well-suited for a designer whose goals are a mixture of maximizing the sum of bids and the winning bid. In contrast, forcing bidders to play a pure equilibrium is typically detrimental to the size of the winning bid. Under linear costs of bidding, policy P^{CG} always leads to smaller winning bids than P^0. Unless v_1 is much greater than v_2, the same holds true for P^*.

6.6 Incomplete Information

A natural question left open by our analysis is whether the results carry over to the case of incomplete information. Bid-caps with symmetric tie-breaking have been studied by Gavious, Moldovanu and Sela (2003) and by Sahuguet (2006): Gavious, Moldovanu and Sela (2003) study the ex ante symmetric n-bidder case under concave or convex cost-functions while Sahuguet considers the two-bidder case under linear costs when the bidders' valuations are drawn from uniform distributions with different supports. In both cases, it turns out that (due to a bunching of types at the bid-cap) all bidders play the bid-cap with positive probability. Therefore different tie-breaking rules must lead to different equilibria. In the ex ante symmetric setting of Gavious, Moldovanu and Sela, there is little reason to expect that varying the tie-breaking rule and thus creating an ex ante asymmetry will be beneficial to the designer. Yet it seems plausible that asymmetric tie-breaking may be beneficial to the designer if there are ex ante asymmetries like in the setting of Sahuguet (2006). A further analysis is a challenging task which we leave to further research.

6.7 Soft Bid-Caps

Kaplan and Wettstein (2006) point out that in many applications the designer will not have the power to enforce a rigid bid-cap: It seems likely that - for example in political lobbying or in competition for high-profile employees - bid-caps will mainly have the consequence of forcing the contestants into less direct or semi-legal ways of exerting effort. Consequently, Kaplan and Wettstein propose that the designer's possibility of setting bid-caps should rather be modeled as a possibility to sharply increase costs of bidding from some point on. Che and Gale (2006) show that setting such "soft" bid-caps may still be beneficial to the designer.

If such a cost increase is modeled as a kink in the cost function which leaves it continuous and monotone, Theorem 3 of Siegel (2010) (given in formula (2) above) covers the equilibrium analysis of Kaplan and Wettstein (2006) and Che and Gale (2006) (both before and after the soft bid-cap is introduced) as special cases. After the introduction of such a soft bid-cap, ties still occur with probability zero in equilibrium. Hence the tie-breaking rule has no bite.

We believe that a closer look at tie-breaking is nevertheless important, even in light of the critique of Kaplan and Wettstein. For one thing, there are applications where setting a binding cap is certainly possible: In Gavious, Moldovanu and Sela's (2003) example of speed limits in motor sports, there is little reason to doubt that binding rules can be enforced. The same should be true for many applications in sports and games.

Moreover, there is little reason to assume that a bid-cap which may be circumvented e.g. by semi-legal activities corresponds to a continuous kink in the costs of exerting effort. For instance, consider the example of raising money for electoral competition given by Kaplan and Wettstein. There are obviously fixed costs involved in raising funds in illegal ways: For example, there is a risk of image-loss and legal complications and there is a need for specialized additional knowledge. These costs cannot be thought of as merely proportional to the amount raised illegally. Therefore it seems more realistic to assume that the costs of exerting effort make an upwards jump even at a bid-cap that may be circumvented. If this jump is strong enough, i.e., if the stakes are not high enough to make an illegal solution attractive, we are essentially back in the situation of a binding bid-cap. The case of a smaller jump in bid-costs needs a separate analysis, but there are good reasons to expect that tie-breaking still plays a role.

Finally, note that the analysis of Section 5 provides an easy starting point for discussing soft bid-caps under considerably weaker assumptions on the cost functions than in Kaplan and Wettstein (2006) and in Che and Gale (2006): Consider bidders with valuations v_i and strictly increasing, continuous cost functions c_i. Like in Section 5, denote by t the highest bid at which both bidders can yield a non-negative payoff. Denote by $\overline{c}_i(b)$ bidder i's average costs on the interval $[0, b]$. In Corollary 3 we showed that the expected sum of bids is then given by

$$\sum_i \frac{t}{v_i}(c_i(t) - \overline{c}_i(t)).$$

If the designer implements a soft bid-cap which leads to a kink in the bidder's cost functions at some $m < t$ (leaving it continuous), this has two effects on the sum of bids: There is a negative effect as t is decreased but there is also a positive effect as the difference between costs at the cap and average costs is increased through the kink. We do not pursue this analysis further here, but clearly the designer's task is to choose the bid-cap in a way that the latter effect dominates the former.

A Proofs

Proof of Lemma 1

The proof is by contradiction: Assume bidder i placed an atom in $e \in (0, m)$. If the other bidder j was active right below e, then j would prefer to bid slightly more than e thus substantially raising his probability of winning through a marginal increase in bidding costs. If bidder j was inactive right below e, then i would prefer to shift his atom downwards a little. Thus there are no atoms in $(0, m)$. Now assume both bidders played zero with positive probability. Then any bidder who is not always favored by the tie-breaking rule prefers to shift his atom slightly upwards. □

Proof of Lemma 2

The proof is by contradiction: Assume that $e \in (0, m)$ is the smallest bid in one of the bidders' equilibrium supports. Since there are no atoms in $(0, m)$, any bidder who plays e in equilibrium must make a strictly negative equilibrium payoff there since he has positive costs of bidding and a winning probability of zero at e. Since bidders can always secure a payoff of zero from bidding zero this is a contradiction. □

Proof of Proposition 1

The basic message about the structure of equilibrium in our game is contained already in Lemma 2: The lower bound of bidders' supports is either zero or m. Since both bidders bidding zero is not a Nash equilibrium, this implies that there is either a mixed equilibrium where at least one bidder mixes down to zero (and where accordingly at least one bidder earns zero payoff by Lemma 1), or a pure equilibrium where both bidders bid m.

We start with case (iii), i.e., $(\alpha, m) \in C_{III}$. In this case, both bidders can secure themselves a strictly positive payoff by bidding m regardless of the opponents strategy. Thus there cannot by a mixed equilibrium (where one bidder would earn zero) and the unique equilibrium is given by both bidders bidding m.

We now consider the remaining cases, i.e., $(\alpha, m) \in C_I \cup C_{II}$. Write α_i and α_j for the tie-breaking proportions of the two bidders, i.e. $\alpha_1 = \alpha$ and $\alpha_2 = 1 - \alpha$. We assume until further notice that $\alpha \in (0, 1)$. The case $\alpha \in \{0, m\}$ is treated at the end.

For $(\alpha, m) \in C_I \cup C_{II}$, both bidders bidding m is not a Nash equilibrium. Accordingly, there must be a mixed equilibrium where at least one bidder's support goes down to zero. From Lemma 1 we know that there may be at most three atoms, two in m and one in 0. Furthermore, we can easily show the following: $S_i \cap (0, m) = S_j \cap (0, m)$ and there is a $\bar{b} \in [0, m)$ such that $S_i \cap (0, m) = (0, \bar{b})$: In words, this means that the non-atomic parts of the bidders strategies' must have the same support which is an interval. To see this, note that when a bidder mixes alone over an interval where the opponent is inactive, the bidder prefers to concentrate mass at the lower end of the interval. If both bidders are inactive over an interval but (at least) one of them is active on higher values in $(0, m)$, then he prefers to deviate from right above the empty interval into the interval.[21]

Moreover note that - just like in a complete information all-pay auction - since bidder i's mixing must make bidder j indifferent between all bids in $(0, \bar{b})$, bidder i must mix in a way that j's changes in bidding costs must be offset by the changes in the probability of winning. Thus i must mix with a uniform density proportional to $1/v_j$ when playing bids in $(0, m)$.

Thus we can restrict our search for equilibria in the case $(\alpha, m) \in C_I \cup C_{II}$, $\alpha \in (0, 1)$, to strategy pairs which can be described as follows: Bidder i sets atoms c_i and d_i in 0 and m. Moreover, he mixes uniformly over $(0, \bar{b})$ with the remaining mass. We now have to find out which combinations $(c_1, c_2, d_1, d_2, \bar{b})$ correspond to Nash equilibria. Observe that from that fact that each bidder must earn the same payoff in 0, $\bar{b}$ and m we get the following set of linear equations: For $i \neq j$,

$$(\alpha_i d_j + (1 - d_j))v_i - m = v_i(1 - d_j) - \bar{b} \quad \text{and} \quad v_i(1 - d_j) - \bar{b} = c_j v_i. \qquad (3)$$

This gives us four constraints, however we have five parameters. Now recall that at most one bidder sets an atom in zero, i.e., $c_1 = 0$ or $c_2 = 0$. Assuming either of these two constraints, we can thus easily solve the system of equations for unique vectors $(c_1, c_2, d_1, d_2, \bar{b})$. This leaves us with two equilibrium candidates (which are stated explicitly in cases (i) and (ii) of the proposition). It remains to be explored for which values (α, m) each of these is an equilibrium. For this purpose we study, which values (α, m) lead to non-negative payoffs for both bidders in each of the two equilibrium candidates. As we will see, this will give us the distinction between cases C_I and C_{II}: There are disjoint regions on which one of the two candidates yields non-negative payoffs while the other does not and vice-versa.

[21]Observe that this does not preclude that there is an interval below m where both bidders are inactive. This is because both bidders may play an atom in m.

We begin with some preliminary calculations: Solving the two conditions in (3) for, respectively, m and $\bar{b}$ and then equating them for $\{i, j\}$ and $\{j, i\}$ gives

$$v_i(1 - c_j - (1-\alpha_i)d_j) = v_j(1 - c_i - (1-\alpha_j)d_i) \text{ and } v_i(1 - c_j - d_j) = v_j(1 - c_i - d_i).$$

Subtracting these two equations from each other we obtain the following result on the relation of d_i and d_j:

$$d_i = d_j \frac{\alpha_i}{\alpha_j} \frac{v_i}{v_j} \tag{4}$$

Now assume $c_j = 0$. Then the two conditions in (3) imply that

$$d_j = \frac{v_i - m}{v_i(1-\alpha_i)} \text{ and } \bar{b} = v_i - \frac{v_i - m}{1-\alpha_i}. \tag{5}$$

Clearly, $c_j = 0$ implies $\pi_i = 0$ so we have to study under which conditions $\pi_j \geq 0$. Using (4), (5) and the fact that $\alpha_i = 1 - \alpha_j$, one easily calculates that

$$\pi_j = v_j(1 - d_i) - \bar{b} = v_j - v_i + \left(\frac{1}{\alpha_j} - \frac{\alpha_i}{\alpha_j^2} \right)(v_i - m).$$

Further rewriting the right hand side, we can thus conclude that $\pi_j \geq 0$ is equivalent to

$$\pi_j = v_j - v_i + \frac{1}{\alpha_j}\left(2 - \frac{1}{\alpha_j}\right)(v_i - m) \geq 0.$$

Viewed as a function of $z = 1/\alpha_j$ this is a second-degree polynomial with negative leading coefficient which is non-negative between its two zeros (if those exist). Solving this quadratic equation we thus have that $\pi_j \geq 0$ is equivalent to

$$1 - \sqrt{1 + \frac{v_j - v_i}{v_i - m}} \leq \frac{1}{\alpha_j} \leq 1 + \sqrt{1 + \frac{v_j - v_i}{v_i - m}} \tag{6}$$

which can be written as

$$-\sqrt{\frac{v_j - m}{v_i - m}} \leq \frac{1 - \alpha_j}{\alpha_j} \leq \sqrt{\frac{v_j - m}{v_i - m}}.$$

This expression shows clearly that the term under the square-root in (6) is always non-negative. Since $\alpha_j \in (0, 1)$ we have thus shown that the equilibrium candidate we found under the condition $c_j = 0$ is reasonable only if

$$\frac{\alpha_i}{\alpha_j} \leq \sqrt{\frac{v_j - m}{v_i - m}}. \tag{7}$$

Yet now note that if we exchange the roles of i and j and rearrange terms, we obtain that the equilibrium candidate we find in the other case $c_i = 0$ is reasonable only if

$$\frac{\alpha_i}{\alpha_j} \geq \sqrt{\frac{v_j - m}{v_i - m}}.$$

Hence, except in the boundary case where (7) holds with equality, we have identified a unique equilibrium candidate for all $(\alpha, m) \in C_I \cup C_{II}$. It is straightforward to check that the two candidates coincide in this boundary case and that they indeed constitute equilibria. We still have to write the condition that separates C_I and C_{II} in the (more revealing) form given in the proposition. Set $i = 1$, $j = 2$, $\alpha_i = \alpha$ and recall that $v_1 \geq v_2$. Consider the boundary case which we have to solve for m:

$$\frac{\alpha}{1-\alpha} = \sqrt{\frac{v_2 - m}{v_1 - m}}$$

Since the left hand side is positive, this is equivalent to

$$m\left(1 - \left(\frac{\alpha}{1-\alpha}\right)^2\right) = v_2 - \left(\frac{\alpha}{1-\alpha}\right)^2 v_1.$$

Since the equation cannot be fulfilled for $\alpha = \frac{1}{2}$, we can exclude that case in the following and write

$$m = v_2 - \frac{\alpha^2}{1-2\alpha}(v_1 - v_2). \tag{8}$$

Clearly, at the boundary for $\alpha \downarrow 0$, (8) is solved by $m = v_2$, showing that bidder 1 earns zero payoff in the upper left corner of C. For $\alpha^* = v_2/(v_1 + v_2)$, (8) is solved by $m^* = v_1 v_2/(v_1 + v_2)$. The resulting point (α^*, m^*) lies on the boundary of C_{III}. It is easy to see that larger values of α are irrelevant here since they do not lead to solutions in $C_I \cup C_{II}$.

To conclude the proof we have to discuss the case $\alpha_i = 1$ for $i \in \{1, 2\}$. In this case, bidder i can secure himself a positive payoff by bidding m. Since one bidder must earn zero payoff in equilibrium, it must hence be bidder j, $j \neq i$. Thus we obtain $c_j = 0$ and can solve the system of equations (3) for unique values of c_1, c_2,

d_1, d_2 and $\overline{b}$. The resulting equilibrium has some qualitative differences to the case $\alpha_i \in (0,1)$ since $\overline{b} = m$ and $d_j = 0$: There is no gap between the interval where bidders mix and m. Moreover, each bidder places only one atom, the strong one at the top, the weak one at the bottom. Nevertheless, not surprisingly, this solution is easily seen to be a boundary case of the solution for $\alpha_i \in (0,1)$. □

Proof of Lemma 3

Recall that it holds that

$$\sigma(\alpha, m) = \sum_{i=1}^{2} (1 - c_i - d_i) \frac{\overline{b}}{2} + d_i m.$$

Using the explicit expressions for c_i, d_i and $\overline{b}$ given in Cases (i) and (ii) of Proposition 1 it is tedious but straightforward to calculate that for $(\alpha, m) \in C_I$ we have

$$\frac{d\sigma(\alpha, m)}{dm} = \frac{(1-2\alpha)(m(v_1 - v_2) + v_1 v_2)}{(1-\alpha)^2 v_1 v_2}.$$

This expression is positive since $\alpha < \frac{1}{2}$ for $(\alpha, m) \in C_I$ and $v_1 > v_2$. For $(\alpha, m) \in C_{II}$ we obtain

$$\frac{d\sigma(\alpha, m)}{dm} = \frac{(2\alpha - 1)(v_1(v_2 - m) + m v_2)}{\alpha^2 v_1 v_2}.$$

Since $v_2 > m$, this expression is positive for $\alpha > \frac{1}{2}$ and negative for $\alpha < \frac{1}{2}$. □

Proof of Corollary 1

Recall that (α^*, m^*) is optimal for the designer within $T \cup C_{III}$. Furthermore, (α^*, m^*) is better than the expected payoff from the unrestricted all-pay auction and thus $\sigma(\alpha^*, m^*) > \sigma(\alpha, m)$ for all $(\alpha, m) \in U$ with $\alpha > 0$. Finally recall that the equilibrium non-uniqueness at $(\alpha, m) \in T \setminus (\alpha^*, m^*)$ is such that the equilibria in (α, m) are convex combinations of the limits as $\varepsilon \downarrow 0$ of the unique equilibria obtained for $(\alpha, m+\varepsilon)$ and $(\alpha, m-\varepsilon)$. As $\varepsilon \downarrow 0$ the unique equilibria for $(\alpha, m+\varepsilon)$ and $(\alpha, m-\varepsilon)$ differ in that for $(\alpha, m+\varepsilon)$ one of the bidder has an atom in 0 which he shifts to m for $(\alpha, m-\varepsilon)$. Thus when keeping $\alpha \in (0,1)$ fixed and increasing m, the function $\sigma_(\alpha, m)$ makes a downwards jump when passing the line T except in the case $\alpha = \alpha^*$.

Now keep α fixed and increase m from 0 to v_2. We have to distinguish three cases: $\alpha > \frac{1}{2}$, $\alpha \in [\alpha^* \frac{1}{2}]$ and $\alpha < \alpha^*$. Consider first the case $\alpha > \frac{1}{2}$. Then by Lemma 3, $\sigma(\alpha, m)$ is increasing until the boundary of T is reached. Then it makes a downwards jump and then it increases again until $m = v_2$. In the light of these considerations, there are two local maxima and thus two potential candidates for optimal values of $\sigma(\alpha, \cdot)$: The unique (α, m) with $(\alpha, m) \in T$ and $(\alpha, v_2) \in U$. However, as discussed above, both points are dominated by (α^*, m^*).

Now consider a fixed $\alpha \in [\alpha^*, \frac{1}{2}]$. In this case $\sigma(\alpha, m)$ is increasing in m until the boundary of T is reached. Then there is a downwards jump and then $\sigma(\alpha, m)$ decreases until U is reached. In this case, there is a unique maximizer m of $\sigma_(\alpha, \cdot)$ and this m must be such that $(\alpha, m) \in T$. But this implies that (α^*, m^*) must be better.

Finally consider the case of $\alpha < \alpha_*$. Similar to the first case, there are two local maximizer of $\sigma(\alpha, \cdot)$, on with $(\alpha, m) \in T$ and one with $(\alpha, m) \in S$. For the first one it is clear that it is dominated by (α^*, m^*). Thus global maximizers must lie in $\overline{S}$.

□

Proof of Proposition 2

All we have to do is maximize $\sigma(\alpha, m_2(\alpha))$ in $\alpha \in [0, \alpha^*]$ since $m_2(\alpha)$ as defined in Definition 1 is the function for which $(\alpha, m) = (\alpha, m_2(\alpha))$ is equivalent to $(\alpha, m) \in \overline{S}$. A straightforward calculation reveals that

$$\frac{d\sigma(\alpha, m_2(\alpha))}{d\alpha} = \frac{\alpha(1-\alpha)(v_1 - v_2)^3}{(1-2\alpha)v_1 v_2} > 0.$$

Thus it is optimal to choose α as large as possible: $(\alpha^*, m_2(\alpha^*)) = (\alpha^*, m^*)$ is a global optimizer. □

Proof of Corollary 3

Denote by σ_1^0 and σ_2^0 the contributions to σ^0 of bidders 1 and 2, i.e., the expected equilibrium bids of these bidders. Since bidder 1's strategy F_1 is atomless, we obtain by (2)

$$\sigma_1^0 = \int_0^t 1 - F_1(b)db = \int_0^t 1 - \frac{c_2(b)}{v_2}db = \frac{t}{v_2}(c_2(t) - \overline{c}_2(t)),$$

where in the last step we use that $c_2(t) = v_2$. In the case of bidder 2, we must take into account the atom he places in zero. It is easy to verify that we can write σ_2^0 as

$$\sigma_2^0 = (1 - F_2(0)) \int_0^t 1 - \frac{F_2(b) - F_2(0)}{1 - F_2(0)}db.$$

Inserting $F_2(0) = 1 - \frac{c_1(t)}{v_1}$ and simplifying analogously to the case of σ_1^0 yields

$$\sigma_2^0 = \frac{t}{v_1}(c_1(t) - \bar{c}_1(t)),$$

Calculating $\sigma^0 = \sigma_1^0 + \sigma_2^0$ completes the proof. □

Proof of Lemma 5

Clearly, bidder 2 makes a positive payoff and thus if there is an atom in zero, it is played by bidder 1. Bidder 1 earns zero payoff. Moreover, bidder 2 can secure himself a payoff of $v_2 - c_2(m)$. He cannot earn more than this because a higher payoff would only be possible if supports ended at some $\bar{b} < m$. But then, bidder 1 could overbid bidder 2 and earn a positive payoff himself. This implies that bidder 1 must play an atom of size $(v_2 - c_2(m))/v_2$ in zero. Thus to make bidder 2 indifferent, bidder 1 has to mix with

$$F_1(b) = \frac{v_2 - c_2(m)}{v_2} + \frac{c_2(b)}{v_2}$$

on some interval $[0, \bar{b}]$ and then put the remaining mass in an atom in m. However, since tie-breaking favors bidder 2, bidder 1 would move such an atom down to $\bar{b}$ if $\bar{b} < m$. This implies $\bar{b} = m$ and accordingly, bidder 1 does not have mass left for an atom in m. Thus bidder 2 must mix over $[0, m)$ as well and put the remaining mass into an atom in m. Since he has to make bidder 1 indifferent, the function F_2 given in the lemma is his only choice. □

References

[1] **Baye, Michael, Dan Kovenock, and Casper de Vries.** 1993. "Rigging the lobbying process: An application of the all-pay-auction." *American Economic Review*, 83(1): 289-94.

[2] **Baye, Michael, Dan Kovenock, and Casper de Vries.** 1996. "The all-pay-auction with complete information." *Economic Theory*, 8(2): 91-305.

[3] **Clark, Derek J., and Christian Riis.** 1998. "Competition Over More Than One Prize." *American Economic Review*, 88(1): 276-89.

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

[5] **Che, Yeon-Koo, and Ian Gale.** 2006. "Caps on Political Lobbying: Reply." *American Economic Review*, 96(4): 1355-60.

[6] **Gavious, Arieh, Benny Moldovanu, and Aner Sela.** 2003. "Bid Costs and Endogenous Bid Caps." *RAND Journal of Economics*, 33(4): 709-22.

[7] **Hillman, Arye L., and John G. Riley.** 1989. "Politically Contestable Rents and Transfers." *Economics and Politics*, 1(1): 17-39, 1989.

[8] **Hillman, Arye L., and Dov Samet.** 1987. "Dissipation of Contestable Rents by Small Numbers of Contenders." *Public Choice*, 54(1): 63-82.

[9] **Kaplan, Todd, and David Wettstein.** 2006. "Caps on Political Lobbying: Comment." *American Economic Review*, 96(4): 1351-54.

[10] **Konrad, Kai A.** 2009. *Strategy and Dynamics in Contests.* Oxford: Oxford University Press.

[11] **Persico, Nicola, and Nicolas Sahuguet.** 2006. "Campaign Spending Regulation in a Model of Redistributive Politics." *Economic Theory*, 28(1): 95-124.

[12] **Sahuguet, Nicolas.** 2006. "Caps in asymmetric all-pay auctions with incomplete information." *Economics Bulletin*, 3(9): 1-8.

[13] **Siegel, Ron.** 2009. "All-Pay Contests." *Econometrica*, 77(1): 71-92.

[14] **Siegel, Ron.** 2010. "Asymmetric Contests with Conditional Investments." *American Economic Review*, 100(5): 2230-60.

[15] **Weber, Robert J.** 1985. "Auctions and Competitive Bidding." In *Fair Allocation*, ed. H. Peyton Young, 143-70. Providence RI: American Mathematical Society, 1985.