Bluffing Beyond Poker

Johannes Hörner
Northwestern University*

Nicolas Sahuguet
Université Libre de Bruxelles†

February 6, 2002

Abstract

This paper introduces a model for bluffing that is relevant when bets are sunk and only actions -not valuations- determine the winner. Predictions from poker are invalid in such nonzero-sum games. *Bluffing* (respectively *Sandbagging*) occurs when a weak (respectively strong) player seeks to deceive his opponent into thinking that he is strong (respectively weak). A player with a moderate valuation should bluff by making a high bet and dropping out if his bluff is called. A player with a high valuation should vary his bets and should sometimes sandbag by bidding low, to induce lower bets by his rival.

"Designers who put chess-men on the dust jackets of books about strategy are presumably thinking of the intellectual structure of the game, not its payoff structure; and one hopes that it is chess they do not understand, not war, in supposing that a zero-sum parlor game catches the spirit of a non-zero-sum diplomatic phenomenon." - Thomas Schelling

*Kellogg School of Management, Department of Managerial Economics and Decision Sciences, 2001 Sheridan Road, Evanston, IL 60208, USA e-mail: j-horner@kellogg.nwu.edu

†ECARES, 50 Av. Fr. Roosevelt, CP114, 1040 - Brussels Belgium. Email: Nicolas.Sahuguet@ulb.ac.be. Web page: **http://www.ssc.upenn.edu/~nicolas2**

1. Introduction

Deception and misdirection are common practice in human affairs. They are widely used and generally expected in politics, business and even in sports. In each of these activities the resources an actor invests in a contest will depend on the way he perceives his opponent's strength. Manipulating this perception can, therefore, have an important effect on the outcome. As Rosen (1986) has put it, "There are private incentives for a contestant to invest in signals aimed at misleading opponents' assessments. It is in the interest of a strong player to make rivals think his strength is greater than it truly is, to induce a rival to put in less effort. The same is true of a weak player in a weak field." During preliminary hearings in a legal battle, for instance, lawyers can use confidently presented expert reports to mislead the other side as to the strength of their client's arguments, encouraging a favorable settlement and avoiding costly litigation. In political lobbying, heavy initial investment by one side can discourage others from entering the contest at all.

Despite its importance, the use of deceptive strategies has been largely neglected by the economic literature, where formal analysis has largely relied on analogies with the game of poker. Von Neumann and Morgenstern (1944), for example, demonstrated that it is rational for poker players to manipulate their opponents' beliefs: players with strong hands have an interest in appearing weak, thereby inducing other players to raise the stakes. Conversely players with weak hands aim to create an appearance of strength, increasing the probability that their opponents will fold. Experienced players commonly use both strategies, known in the trade as "sandbagging" and "bluffing".

Although poker is undoubtedly a good source of intuitions about human affairs, we argue in this paper that the analogy **can also be** misleading. There are important features of the game of poker which many real-life contests do not share. First, poker has a particular payoff structure: all bets go into a pot which is taken by the winner. This means that poker is a zero-sum game: one player's gain is the other's loss. A second peculiarity of poker is the way the winner is decided. A game of poker ends either when everyone except the winner abandons the game, or when players show their cards, in which case the winner is the player with the strongest hand. As a model of business or war this is incorrect. First, as Schelling points out in the introductory quotation, most real life contests are not zero-sum. Resources expanded in battle are never recovered, for instance. Second, real life contests are rarely "beauty" contests in which the outcome is directly determined by an eventual comparison between some privately known attributes of the players. While the rules deciding contests may be complex they tend to depend, not so much on private information as on the public actions of the contestants.

In short, poker is an inaccurate representation of real-life contests, whether in business or in war. And these differences matter. Because poker is zero-sum, one player wins what his opponent loses; bets in early rounds change not only the players' beliefs but also the stakes. In business, what one player spends is lost for everybody; early bids **may** affect players beliefs but do not enter the stakes. In poker, a player with an unbeatable hand will always win back his bets and his opponent will lose his. In business, where winning has a cost, there is another option, namely a pyrrhic victory. In this setting the motivations and mechanisms underlying bluffing and sandbagging turn out to be very different from those we expect to see in poker.

In this paper, we present a model of bluffing and sandbagging in a stylized non-zero-sum contest. In our game two players compete for a prize. Each player knows how much the prize is worth to him. This "valuation" is private information. A player can use his early behavior to manipulate his opponent's beliefs about his valuation. The purpose of this paper is to understand the ways in which this can occur. More specifically, we consider a two period model. One player opens the game with an initial bid. His opponent can then either pass or cover the bid. If he passes, the game is over. If he covers, both players bid a second time, this time simultaneously. The prize is awarded to the highest bidder. All bids are sunk. This game form -but not its payoffs- closely mirrors the one of the standard poker models, facilitating a comparison carried out in Section 5.

When he makes his initial bid the first player can bid either high or low. Bidding high has two effects. It may deter the other player from continuing the contest, allowing the first player to win with no further bidding. This is the *deterrence effect*. But if the bid is covered it can also lead to an *escalation effect*. If the initial bid is interpreted as a sign of strength, the second player correctly infers that to have a chance of winning he has to bid aggressively in the second round. While the deterrence effect benefits the first player, escalation makes it more expensive for him to win.

Bidding low, the alternative option, has a *sandbagging effect*. This kind of bid certainly does not deter the second player. Interpreted as a sign of weakness it can, nonetheless, induce him to weaken his bid in the second round, so as not to waste resources. This reduces the costs of winning and makes sandbagging an attractive option for players with a high valuation.

In standard signalling games, the sender always tries to convince the receiver that he is strong - in our terms that he has a high valuation. This does not represent well the usual idea of deceptive tactics. McAfee and Hendricks (2001) analyze misdirection tactics in a military context. Generals choose how to allocate their forces and where to attack. Imperfect observability of actions generates inverted signalling. The allocation of forces is used to deceive the opponent about the location of the attack. As in poker, a player always wants to convince his opponent of the opposite of what he plans to do. In our game, however, the incentives to signal are more sophisticated. In particular a player with a high valuation can benefit both from being perceived as very strong and from being perceived as very weak. This leads to complex equilibrium behavior where both **direct** signalling and inverted signalling are present. Players with a weak valuation will make low opening bids; those with intermediate level valuations will "bluff" making high opening bids to achieve deterrence but withdraw from the contest if their bid is called - thereby avoiding escalation. Players with high valuations will choose randomly between high and low opening bids, enjoying both the deterrence effect of a high bid, and the sandbagging effect of a low one. The second player's decision whether or not to cover is less interesting. If the prize is worth enough to him he covers; if not he will pass.

These results shed a new light on deception strategies. In sports for instance, it explains the tactics adopted by long distance runners and cyclists. It provides insight into business strategies such as Airbus Industries and Boeing's use of delaying tactics and press announcements in their race to develop a "super" jumbo jet. Besides, deception is not confined to human affairs. In fact, the modelling paradigm for dynamic conflicts, the *war of attrition*, has first been developed by biologists (Maynard Smith (1974), Bishop & Cannings (1978)) to study animal behavior. The war of attrition allows for very limited information transfer. A large body of data, collected mainly by ethologists, shows that such transfer occurs during animal contests and that bluff plays a role (Maynard Smith (1982)). But as John Maynard Smith puts it (1982, p. 147), "The process is not well understood from a game theoretic point of view". Our model may be viewed as an extension of the war of attrition allowing for costly signalling. A player gets the opportunity to make a costly display, and if this display is not deterrent, a simultaneous all-pay auction takes the place of the ensuing war of attrition. Indeed, one of the most intriguing phenomena described in this paper corresponds exactly to Riechert's finding (1978) that winning spiders *Agelenopsis aperta* show a more varied behavior than losing ones.

In each of these cases players' tactics differ substantially from those they would adopt in poker. In poker, the purpose of sandbagging is to induce one's opponent to *raise the stakes*. Here, on the other hand, the goal is that he should *make a lower bid*. This is what psychologists would predict (See Gibson and Sachau (2000)), and agrees with informal observations of the ways players behave in a number of competitive settings. As for bluffing the goal - deterrence - is the same as in poker. But though players bluff for the same reasons, they should not bluff in the same way. In poker it is always rational to bluff with the worst possible hand (see, for instance, Newman

(1959)). In our model, on the other hand, where all bids are sunk, bluff has a cost, whether or not the bluff is called. As a result a player with a low valuation cannot afford to bluff; bluffing is recommended only for players who put a sufficient value on the prize. In our model, unlike poker, the advantage of bluff depends on its cost.

In what follows we will investigate these differences more closely. Section 2 presents the model. Section 3 identifies the main equilibrium. Section 4 describes other equilibria, and establishes uniqueness for given parameters. Section 5 discusses these results and Section 6 concludes.

2. The Bidding Game

Two risk-neutral players (1 and 2) compete for a prize. The two players' valuations, respectively v for player 1 and w for player 2, are drawn independently from the uniform distribution on $[0, 1]$. Valuations are private information. The contest is modeled as a two-stage game. In the first stage of the game player 1 makes an opening bid. This bid is either high or low. The cost of a low bid is normalized to 0; high bids have cost $K > 0$. If player 1 bids K, player 2 chooses whether or not to cover. Covering costs him K as well. If he does not cover, player 1 wins the prize. If he does cover, or if player 1's opening bid was 0, the game enters a second stage. The second stage of the game consists of a simultaneous first-price, all-pay auction - that is an auction where all bidders pay the price they have bid, regardless of whether or not they have won the auction. This second bid is unrestricted, that is, players can bid any nonnegative amount at this stage. The total cost incurred by a player is the sum of his bids. All costs are sunk. The prize is awarded to the highest bidder. In the case of a tie, the winner is chosen randomly. There is no discounting. The game tree can be depicted as follows:

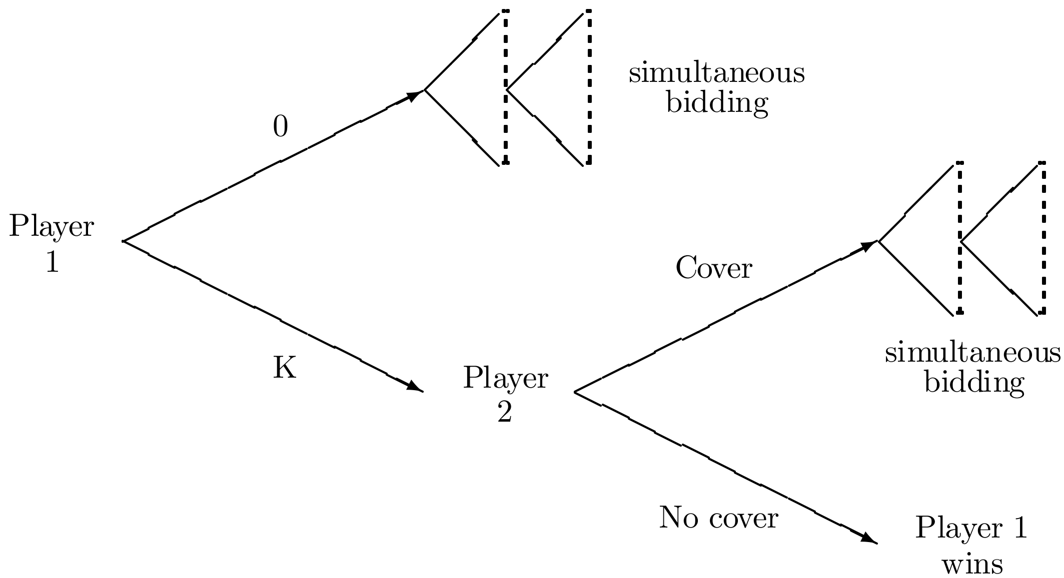

Player 1's strategy specifies his opening bid and his second bid in each of the two possible subgames as a function of his valuation v. Player 2's strategy specifies whether or not he will cover (if required) and how much he will bid in the two subgames, as a function of his valuation w. If player i's valuation is s, his bid in the subgame following an opening bid of $k = 0$ or K will be denoted $b_i^k(s)$.

A (Perfect Bayesian) equilibrium consists of strategies and beliefs for each player. The strategies are sequentially rational, that is, the bid choices maximize the expected payoffs given beliefs about the other player's valuation and strategy. Beliefs are correct and updated according to Bayes' rule.

Depending on the parameter K, different kinds of equilibria exist. We divide them up into three distinct categories, which we classify in terms of their main strategic features.

An equilibrium in which player 1 invests K with positive probability for some of his valuations, whereupon Player 2 covers for some of his valuations, is called an equilibrium with covering.

A nonrevealing equilibrium is an equilibrium in which player 1 invests 0 initially with probability one for all valuations, and thus, player 2 does not need to cover.

Finally, an equilibrium with assured deterrence is an equilibrium in which player 1 invests K with positive probability for some of his valuations, but player 2 does not cover for any of his valuations.

As will be shown, there exists for each of these three categories of equilibrium a set of values for the parameter K where the equilibrium applies. In the next section we will describe the first category of equilibrium, namely equilibrium with covering, before discussing the other equilibria and their relationship. Before doing so, however, it is useful to discuss some of the assumptions underlying the model.

This game is restricted to two periods. This simplifying assumption is not, however, a critical one. The stylized, two stage game provides interesting intuitions about strategic behavior in dynamic contests. Extending the model to a longer but finite horizon would not invalidate the qualitative insights the game can provide about bluffing and sandbagging. What is important is to maintain simultaneous bidding in the last period. Without this sandbagging does not emerge as an equilibrium behavior. The assumption that a last period exists is quite natural and fits well with the idea that in many kinds of contest there may be a deadline for allocating the prize. This issue will be discussed further in section 5.

A second assumption in the model is that opening bids are binary. Elsewhere we have examined what happens when player 1 is allowed to choose any positive bid he pleases. The analysis is more complex but yields results which are qualitatively similar to those for the simplified game.

Finally our model assumes that if player 2 does not wish to abandon the contest he has to match player 1's opening bid. In some situations, it may be more reasonable to assume that covering is unnecessary or unobservable. In such circumstances, it is easy to show that the first player would never make a high opening bid. Allowing the second player not only to match the first bid, but even to raise it is a special case of a longer but finite horizon.

3. Equilibrium with Covering

In equilibrium with covering, bids can escalate. According to circumstances player 1 may either sandbag or bluff. By definition of this kind of equilibrium, attempts by the first player to deter his opponent from competing will sometimes fail. When this happens, his opening bid is lost, the players adjust their beliefs and bidding starts afresh.

Before formally describing equilibrium with covering, it is useful to outline its main features. It follows from revealed preferences that a player's probability of winning and his expected payment are both non-decreasing functions of his valuation. This is also true of his ex ante probability of winning. In fact, for a player with a high valuation the probability of winning is a strictly increasing function of his valuation.[1] (We say that player 1 has a high valuation if he bids high in the first period and carries on bidding if his bid is covered.) An important consequence of this observation is that player 1's strategy in the first period cannot follow a simple cut-off rule, and the initial bid, therefore, cannot be monotonic in valuations. Otherwise, the highest valuation bidding low would win for sure (i.e., with probability one). But in an equilibrium with covering, in the subgame in which an initial high bid has been covered, the lowest valuation that bid high initially cannot win for sure, as he is not the highest valuation bidding so. This contradicts the fact that the probability of winning is non-decreasing in valuations.

[1] To see this, suppose that there are two high valuations $v_1 < v_2$ having the same expected probability of winning; then so should every valuation between v_1 and v_2. Therefore, there must be a positive measure of the first player's valuations making an identical bid $b > 0$ in at least one of the subgames. As the second player would not find it optimal to bid b or slightly less, these valuations would gain by slightly decreasing their bid.

Player 1*'s high valuations randomize over first period bids.*[2] This follows from the fact that winning probabilities are strictly increasing for high valuations. Thus, if a high valuation wins with a given probability p in one subgame, then he must also be the valuation winning with probability p in the other subgame.[3]

Player 2*'s covering decision is non-decreasing in his valuation.* Suppose that a given valuation covers. Since covering is costly, his probability of winning in the subgame that follows must be strictly positive. By mimicking him if necessary, larger valuations of player 2 can ensure themselves a strictly positive payoff. This implies that they cover after a high bid by player 1, as not covering yields a zero payoff.

Intermediate valuations of player 1 *bid high for sure, and nothing afterwards when covering occurs.* Indeed, as the smallest valuations of the second player who cover make arbitrarily small bids, they can only recoup the cost of covering if there is a strictly positive probability that the first player bids nothing when his high bid gets covered. Such must therefore be the behavior of some interval of valuations of the first player, the intermediate valuations. All these valuations make the same expected, total payment $K > 0$. The argument developed in footnote 1 rules then out the possibility that these valuations may also sometimes bid low initially.

Low valuations of player 1 *bid low.* Valuations smaller than the intermediate valuations have total expected payments smaller than K, and will thus, inevitably, bid low.

The following proposition should therefore come as no surprise.

Proposition 1. *For any* $K < \bar{K} \approx 0.26$*, there exists* α*,* β*,* $0 \leqq \alpha \leqq \beta \leqq 1$*,* γ*,* $p \in [0, 1]$*, such that the following strategies constitute the unique equilibrium with*

[2] This observation is meaningful because high valuations exist. To see this, suppose not. Then the first player always bids nothing whenever his early bid is covered, so that the second player's second bid must be arbitrarily small. In this case, the first player could profitably deviate by slightly increasing his bid.

[3] To see that someone must win with such probability p in each subgame, suppose not. Then there must be a gap in the probability of winning, as well as in the bids. But the bid corresponding to the gap's highest extremity is dominated by slightly smaller ones.

covering.

- There are three intervals to consider to describe player 1's first period bidding. Low valuations ($v \in [0, \alpha]$) bid low for sure. Intermediate valuations ($v \in (\alpha, \beta]$) bid high for sure. High valuations ($v \in (\beta, 1]$) bid high with probability $p \in (0, 1)$.

- After a high bid, player 2 covers if and only his valuation is high enough ($w \in (\gamma, 1]$).

- The equilibrium strategies in the subgame following a covered bid are:

$$
\begin{aligned}
b_1^K(v) &= \begin{cases} 0 \text{ for } \alpha < v \leqq \beta, \\ \frac{c_K}{1-\gamma}\left(v^{\mu} - 2K\right) \text{ for } \beta < v \leqq 1, \end{cases} \\
b_2^K(w) &= \left\{ \frac{c_K}{1-\gamma}\left(w^{\frac{\mu}{\mu-1}} - 2K\right) \text{ for } \gamma < w \leqq 1. \right.
\end{aligned}
$$

- The equilibrium strategies when the opening bid is zero are:

$$
\begin{aligned}
b_1^0(v) &= \begin{cases} c_0 v^{\frac{2-\lambda}{1-\lambda}} \text{ for } 0 \leqq v \leqq \alpha, \\ c_K\left(v^{\mu} - 2K\right) + K \text{ for } \beta < v \leqq 1, \end{cases} \\
b_2^0(w) &= \begin{cases} c_0' w^{2-\lambda} \text{ for } 0 \leqq w \leqq \gamma, \\ c_K\left(w^{\frac{\mu}{\mu-1}} - 2K\right) + K \text{ for } \gamma < w \leqq 1. \end{cases}
\end{aligned}
$$

where:

$$
\begin{aligned}
\lambda &\triangleq \beta - \alpha + p(1-\beta), \ \mu \triangleq 1 + \frac{1-p}{1-\lambda}, \\
c_K &\triangleq \frac{1-p}{2-p-\lambda}, \ c_0 \triangleq \frac{\gamma \cdot \alpha^{\frac{1}{\lambda-1}}}{2-\lambda} \text{ and } c_0' \triangleq \frac{\alpha \cdot \gamma^{\lambda-1}}{2-\lambda}.
\end{aligned}
$$

Proof

We discuss here the main steps. Details can be found in Appendix.

We first show that given first-period bidding and covering strategies, the second period bid functions constitute an equilibrium. In both subgames, beliefs are determined by Bayes' rule. In the subgame with a high opening bid, player 2's valuation is uniformly distributed over $(\gamma, 1]$. The support for player 1's valuation is $(\alpha, \beta] \cup (\beta, 1]$. The density is constant over each interval, but differs across intervals, since players with a valuation higher than β invest K with probability p. In the subgame following a zero opening bid, player 2's valuation is uniform on $[0, 1]$, while the support for player 1's valuation is $[0, \alpha] \cup (\beta, 1]$, with densities constant over but different across intervals.

To determine the equilibrium bidding functions for these subgames, we use the mapping $h(\cdot) = b_2^{-1} \circ b_1(\cdot)$, mapping player 1's valuation into player 2's valuation making the same bid. Assume for a moment that the distribution of valuations for player $i = 1, 2$, denoted, F_i. For arbitrary (differentiable) valuation distributions F_i, player 1's objective is to maximize $v \cdot F_2\left(b_2^{-1}(x)\right) - x$ over $x \in \mathbb{R}^+$ while player 2's objective is to maximize $w \cdot F_1\left(b_1^{-1}(y)\right) - y$ over $y \in \mathbb{R}^+$. First-order conditions are:

$$\begin{aligned} F_2'\left(b_2^{-1}(x)\right) \cdot \left(b_2^{-1}\right)'(x) &= \frac{1}{b_1^{-1}(x)}, \text{ and} \\ F_1'\left(b_1^{-1}(y)\right) \cdot \left(b_1^{-1}\right)'(y) &= \frac{1}{b_2^{-1}(y)}. \end{aligned}$$

These equations can be rewritten as:

$$\begin{aligned} \left(b_2^{-1}\right)'(b_1(v)) &= \frac{1}{v \cdot F_2'\left(b_2^{-1} \circ b_1(v)\right)} = \frac{1}{v \cdot F_2'(h(v))}, \\ b_1'(v) &= \frac{1}{\left(b_1^{-1}\right)'(b_1(v))} = h(v) \cdot F_1'(v), \end{aligned}$$

whenever the density is positive. Finally, since $h'(v) = \left(b_2^{-1}\right)'(b_1(v)) \cdot b_1'(v)$, we obtain:

$$h'(v) = \frac{h(v) \cdot F_1'(v)}{v \cdot F_2'(h(v))}.$$

When the distribution is uniform over intervals (as is the case in both subgames), $F_1'(v)$ and $F_2'(h(v))$ are piecewise constant. This ordinary differential equation, with unknown h, is therefore easily solved in each subgame, and the bidding functions b_i

are then obtained from $b_1'(v) = F_1'(v) \cdot h(v)$ (with the lowest type making a zero bid) and $b_2(w) = b_1(h^{-1}(w))$. The details are omitted.

We now turn to the optimality of the first period bidding and covering decisions. The equilibrium is characterized by the following four relationships:

$$\begin{aligned} \frac{\beta-\alpha}{\lambda} \cdot \gamma &= K, \qquad \frac{\alpha\gamma}{(2-\lambda)} = K, 1 \\ \frac{p}{\lambda}(1-\gamma) &= \frac{(1-p)}{(1-\lambda)}, \quad \gamma = \beta^{\mu-1}. \end{aligned} \tag{3.1}$$

Let σ_+ (σ_-) denote valuations arbitrarily close to σ from above (from below). The first equation specifies player 2's optimal covering decision. Players with valuation γ_+ find it optimal to cover. In this subgame, players with valuation γ_+ make bids arbitrarily close to zero and win only when player 1 bids 0. This occurs with probability $\frac{\beta-\alpha}{\lambda}$, whenever player 1 has an intermediate valuation and bluffed with his opening bid. Hence, this equality states that the expected profit from covering with valuation γ_+, $\frac{\beta-\alpha}{\lambda} \cdot \gamma_+$, approximately equals the cost of covering K. It follows that it is optimal for players with higher valuations to cover - they make strictly positive profits - and for players with lower valuations to pass. The second equation ensures that in the subgame with a zero opening bid, the equilibrium bid $b_1^0(\alpha)$ $(= c_0\alpha^{\frac{2-\lambda}{1-\lambda}} = \frac{\alpha\gamma}{(2-\lambda)})$, when player 1 has valuation α, equals K, ensuring that expected profits are continuous at α. Notice that this also implies continuity of profits at β. Players with valuation α and β (in fact, any player with valuation in $[\alpha, \beta]$) are thus indifferent between bidding K or 0. Next, for valuations above β to be indifferent between both bids, it must be that their expected profits are the same. This is the case if and only if a player with valuation β_+ is indifferent between both actions, and the marginal profit with respect to valuation (for valuations larger than β) is the same in both subgames. The second equality ensured the former, and the third relationship establishes the latter. Last, since there is an atom of types for player 1 bidding zero in the subgame K, there cannot be such an atom for player 2, which requires that the function h maps type

β of player 2 into type γ of Player 2, which yields the fourth equation. We show in Appendix that this system admits a solution $(\alpha, \beta, \gamma, p) \in [0,1]^4$, with $\alpha \leqq \beta$, if and only if $K < \bar{K}$, where $\bar{K} \approx 0.26$ is the unique solution to $\bar{K} e^{\frac{1}{\bar{K}-1}} = \frac{e^{-2}}{2}$. ■

The intuition behind the structure of equilibria with covering is the following. Intermediate valuations bluff. By choosing an early bid that only high valuations otherwise make, they take advantage of the deterrence effect generated by these high valuations, who are prepared to bid aggressively even if their opponent covers. As only sufficiently high valuations of player 2 cover, it is then in the bluffer's best interest to give up in such an event.

High valuations sandbag. By bidding low, they take advantage of the low valuations' behavior. Such a choice triggers less aggressive bidding by the second player, who does not want to take the chance of unnecessarily wasting resources, given the probability that this opponent is weak.

Of course, both bluffing and sandbagging are rational behaviors. They correspond to the two different ways in which a player may manipulate his rival's beliefs. The chance that bluffing succeeds may satisfy an intermediate valuation, but not a high valuation, who wants to win with high probability, achieved either through repeated large bids, or through sandbagging.

Who wins, who loses in this game? The natural benchmark is the simultaneous, static, first-price all-pay auction, in which all bids are sunk and the highest bidder wins. As intermediate valuations of the first player take advantage of high valuations, and high valuations take advantage of low valuations, it is not surprising that the intermediate and high valuations benefit from the opportunity of bidding early. Losers are found among the low valuations, who are hurt by the high valuations mimicking them, as this exerts an upward pressure on bids. The situation is reversed for the second player. High valuations may have to first reveal themselves through the cover, which is sunk. Low valuations, however, are able to better adjust their bid, avoiding to waste resources when a high bid reveals the first player to be at least of the intermediate valuation. (The straightforward calculations justifying these claims are omitted).

In some applications, one may want to maximize the revenue of the game, that is, the total expected payments of the players (lobbying - from the politician's point of view), while in others, one may want to minimize it (military conflicts). It is straightforward to show that a moderate, intermediate value of K minimizes revenue. As for revenue maximization, as the static first-price, all-pay auction is an optimal auction by the Revenue Equivalence Theorem (See Myerson (1981)), it is obviously best to set $K = 0$, so that the dynamic game essentially collapses to the static one.

4. The other Equilibria and their Structure

Equilibria with covering are certainly the most interesting ones, displaying intriguing strategic features. Two other kinds of equilibria exist, however. A standard refinement yields a striking existence and uniqueness result, where the selected equilibrium depends on the particular value of K. All the formal arguments can be found in Appendix 2.

4.1. Equilibria with Assured Deterrence

In an equilibrium with assured deterrence, the first player sometimes bid high and the second player never covers. It is straightforward to see that the first player must actually follow a 'threshold' strategy: he bids high if and only if his valuation is sufficiently high. This kind of equilibrium makes sense for relatively high values of K. Although a high bid has an assured deterrence effect, it is costly enough to be chosen only by the highest valuations of player 1. Player 2 has therefore two good reasons to give up after a high bid: covering is expensive, and the opponent is strong. Observe in particular that the second player does not cover even if his valuation is one, thus for sure larger than the first player's. This is due to the asymmetry between players. When the second player gets to cover, his beliefs are updated given the high bid, and this leads him to be pessimistic.

4.2. Non-revealing Equilibria

In a non-revealing equilibrium, the first player bids low for sure. The game is then essentially equivalent to a static, first-price all-pay auction. This equilibrium seems reasonable when K is very large, so that the high bid is unattractive.

4.3. The Structure of Equilibria

Since all these kinds of equilibria may exist, which one is more likely to emerge? Obviously, this depends on the value of the parameter K. (Perfect Bayesian) Equilibria with covering exist if and only if $K < \bar{K}$, equilibria with assured deterrence exist if and only if $K < 1/2$, and non-revealing equilibria always exist. However, the beliefs used to construct non-revealing equilibria for low K are not plausible. Such equilibria make sense for large K, when early bidding is not worthwhile, but seem pathological otherwise. The Intuitive Criterion does not have any bite in this game because, while constraining the support of beliefs that can be held after a deviation, it does not impose any restriction on the relative likelihood of those valuations belonging to this support. With a continuum of valuations, this leaves considerable leeway. One might wish to impose that, if a player has incentives to deviate for two distinct valuations, his opponent's beliefs after observing such a deviation should preserve the relative likelihood of these valuations. This is the main idea behind Perfect Sequential Equilibrium (P.S.E.), defined by Grossman and Perry (1986). While the formal definition is cumbersome and therefore relegated to Appendix 2, its spirit is straightforward. Fix a Perfect Bayesian Equilibrium and suppose that a player deviates. His opponent hypothesizes that the move was made by some subset C of the player's valuations, and revises his belief according to Bayes' rule conditional upon the player's valuation being in C. If the Perfect Bayesian Equilibrium that follows given these beliefs is preferred to the original equilibrium by precisely the valuations in C, then the original equilibrium fails to be sequential. The point is that this deviation allows the players' valuations in C to convincingly separate themselves from the other valuations, so that it is not credible for his opponent to hold any other belief after such a deviation. This eliminates equilibria based on such beliefs. This refinement is inspired by a forward induction argument. Deviations should be interpreted not as trembles, but as rational signals to influence beliefs. Proposition 2, which is proved in Appendix 2, establishes the structure of P.S.E. in this game.

Proposition 2. *For each* $K \geqq 0$*, there exists a unique P.S.E. This P.S.E. is the equilibrium with covering for* $K < \bar{K}$*, the equilibrium with assured deterrence for* $\bar{K} \leqq K < 1/2$*, and the non-revealing equilibrium otherwise.*

This result, illustrated in the following figure, is intuitive. For large K, deterrence is too costly and the first player does not take advantage of this opportunity. For intermediate K, deterrence is effective. Given the entry cost that it represents and the signal of strength that it conveys, a high bid deters for sure the second player. Finally, for low values of K, a high bid is not always deterrent and the equilibrium exhibits covering.

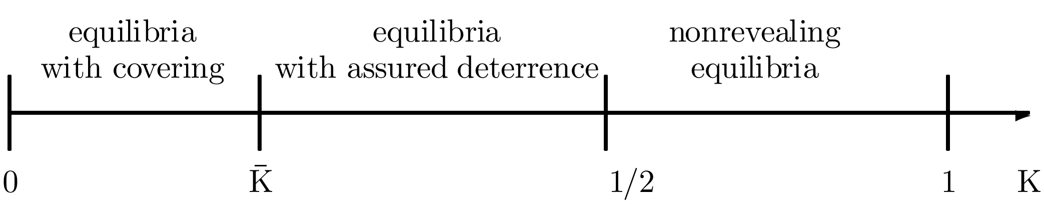

5. Discussion

To date, the main paradigm used by game theorists to discuss informational questions in a rigorous framework has been poker. Variations of this game have been studied by Borel (19xx), Von Neumann and Morgenstern (1944), Nash and Shapley (1950), to name just a few. The models developed by these authors continue to feature prominently in game theory textbooks (see Binmore (1992) for an excellent introduction). As McDonald puts it in his classic book *Strategy in Poker, Business and War*, "The theory of games originated in poker, and that game remains the ideal model of the basic strategical problem" (p. 56). As he argues, "Buyer and seller are the same as the players in a two-man poker game" (p. 73). Even "An airplane and a submarine, for example, play poker" (p. 108). Poker has undoubtedly been a wonderful guide to our intuition. However, the game we have discussed here is based on different assumptions, played using different strategies and gives different results.

Poker is a zero-sum game. In the case of a showdown, the winner is decided by the players' hands, which are beyond their control. As a former champion bridge player and authority on poker, Oswald Jacoby (quoted in McDonald, p. 25) has put it, "If an opponent holds a royal flush, he is going to win the pot, if any; but, if you are smart enough to figure out that he holds it, you do not have to call his bet." That is why poker is not about managing cards, but money. Models which place restrictions on how much can be won or lost by a player do so only to ensure tractability. Given that he is playing a zero-sum game, a poker player can win whatever his opponent spends; bets made in early rounds affect players' behavior not only by affecting their beliefs but by raising the stakes.

In our game, on the other hand, winning is about managing the trade-off between the probability and the cost of winning. This is a trade-off very familiar to economists. By the rules of the game every bid costs. It is a non-zero sum game where what a player spends is lost for everybody, and bids have no affect on the stakes. This provides a setting where the strategic manipulation of beliefs can be studied in a purer form than in poker. At the same time it becomes possible to study the cost of the contest. As a matter of fact, one of the motivations for this work is to analyze the way in which expected total bids depend on the dynamic formulation. In lobbying, politicians like to maximize this revenue. In warfare, one may want to minimize it.

To discuss and compare the results, it is helpful to recall a distinction introduced by Von Neumann and Morgenstern. In the *Theory of Games and Economic Behavior* (1944), they identify two motivations for bluffing in poker: the first one consists in bidding high or overbidding to created a (false) impression of strength, thus conceivably inducing one's opponent to pass. The second stems from the need to create uncertainty in the opponent's mind as to the correlation between bids and hands. In their own words, "The first is to give a (false) impression of weakness in (real) weakness, the second is the desire to give a (false) impression of weakness in (real) strength". Bluffing pays off because it sometimes induces the other player to believe that his opponent's hand is really strong, thereby exerting a deterrent effect, and because, at other times, it induces him to raise, hoping that the hand is terrible while it is actually really good. By bluffing, you can win a lot with a bad hand. Without, you will only win very little, even with a good hand. In models of poker similar to the model proposed by Von Neumann (see Karlin and Restrepo, Newman and Sakai for extensions), bluffing is optimal when a player's hand is really bad, and not middle-range: failing to understand this is a mistake commonly made by amateur poker players.

Maybe these amateurs are confusing poker with the game presented here, where bluffing is genuinely optimal for players with intermediate valuations. Players with low valuations do not bluff simply because it is too expensive to do so. As in poker, bluffing pays off because it may deter the opponent from competing further. Unlike poker, however, a player who bids high with a high valuation has nothing to gain by confusing his opponent's beliefs. This may lead the other player to "call", a disastrous outcome for both players. A player with a moderate valuation gains an advantage by mimicking the behavior of a high valuation player.

Sandbagging is a behavior that is relatively rare in poker. We find it in Nash and Shapley's three-player poker model (1950). In this model all hands belong to one of just two categories: high and low. The first two players may bid low, even if they have strong hands, so as to induce the third one to raise the stakes by bidding high. As with bluffing, the confusion of beliefs this generates continues to be beneficial even when their hands are actually low. Sandbagging in our model has the opposite motivation: the aim is persuade the other player to bid low, allowing the sandbagger to win at moderate cost. This strategy is particularly attractive to a player with a high valuation, for whom winning is especially important. Sandbagging exploits the behavior of players with low valuations. By copying this behavior, the sandbagger induces his opponent to bid cautiously so as to avoid potentially unnecessary expenditure. The uncertainty created by sandbagging leads the second player to maker a stronger response to an opening bid than he would if he could be sure he was facing a low valuation player. As a result sandbagging damages players with low valuations.

According to psychologists Shepperd and Socherman (1997), sandbagging involves displaying oneself as an unworthy foe for the purpose of undermining an opponent's effort or inducing an opponent to let down his or her guard. Lulling a rival into a false sense of security is a common tactic in sports such as golf, biking or wrestling. Shepperd and Socherman point out three important factors that influence the decision to sandbag. First, people should only sandbag when they are competing with another person or persons for a desired outcome. Second, the outcome should be uncertain. Finally, the duration of the game should be limited: "Sandbagging is a useful one-time strategy that is best suited for undermining an opponent's preparation for a competition or for contests in which the outcome can be determined by a single move or play". The first and the third observations are undoubtedly important and are discussed below. As for the second point they claim that "if people anticipate outperforming their opponent, then there is no reason to display oneself as disadvantaged". If their experiments are correct, then our model suggests that people do not sandbag optimally, just as amateur poker players tend to bluff with the wrong hands. Inducing your rival to reduce his effort is most useful when you expect to match your rival's efforts. Lance Armstrong, who dominated the 2001 Tour de France biking event, is no amateur in this respect. He sandbagged his rival Jan Ullrich in the tenth leg. As a journalist described it, "On the way to his first stage win, at the famed l'Alpe d'Huez in stage 10, Armstrong acted like a rider who was destined to finally succumb to the efforts of the four-time runner-up, Jan Ullrich. His face was full of faux suffering and his Telekom rival swallowed the bait. The 1997 Tour champion churned his huge gears over the preceding climbs and worked his legs into lactic overdrive. A move which allowed Armstrong to surge on the 13.9 kilometer climb which concluded stage 10."[4] A few days later, the German tried to bluff Armstrong by violently attacking in the Pyrénées mountains. His bluff was called by the Texan.

While focusing on two periods rather than an arbitrary finite horizon is an inessential simplification, the fact that our game is of finite duration is of critical importance. In related work, we have studied an infinite horizon version of this game, and have found that, unlike bluffing, sandbagging never arises in (perfect sequential) equilibrium. This is not surprising. Bluffing forces one's opponent to make an important decision immediately: to pass or to pay. Sandbagging, on the other hand, is useless without a deadline. The lack of a deadline allows the other player to simply wait and see.

It is important that players move in turn. Consider a variant of the game in which both players can decide their opening bids simultaneously and where, if the bids turn out to be different, the low bidder gets a chance to match the higher bid. In this setting players use "Colonel Blotto" strategies: they bid high if and only if their valuation exceeds some threshold. Bidding low with a high valuation is not attractive: either the other player has bid high, in which case it is then best to cover, and the deviation merely amounts to a change in the timing of the bids, or he has bid low, in which case it will cost more to win later than it would have cost to deter the player in the first place. If players in this variant are free to choose any opening bid they like (and not just a high or a low one), opening bids are the same as they would have been in a (static) first price all-pay auction, and the lower bidder never covers. This suggests that information manipulation plays an important role only in settings where agents are unlikely to take decisions at exactly the same time.

[4] The article can be found at http://www.letour.fr/2001/us/, Sunday, July 29^{th} 2001.

6. Concluding Remarks

This paper shows that the common analogy between poker and human affairs, while useful, is often misleading. In our model, as in poker, players display bluffing and sandbagging behavior. The conditions in which these strategies are applied are, however, different. Bluffing is not used by weak players, but by moderately strong ones. Sandbagging is not used to raise the stakes, but to induce an opponent to make a lower bid. Note that Rosen's informal analysis (1986) leads to different predictions. Rosen conjectures that a strong player wants his rival to think his strength is greater than it truly is, thereby inducing him to exert less effort, and that the same applies to a weak player in a weak field; a weak player in a strong field seeks, on the other hand, to give out signals showing that he is even weaker than he actually is, thereby leading his rival to slack off. But what Rosen has in mind is private knowledge about abilities, not about valuations. This distinction seems likely to play an important role in dynamic conflicts and deserves further research.

Although the model developed in this paper is highly stylized, we believe that it captures important insights about the strategic nature of bluffing and sandbagging in situations where players choose both the timing and the level of their effort. It would be interesting, from a theoretical viewpoint, to examine whether these results carry over to versions of the game with an infinite horizon version. This would lead to a model close to wars of attrition where players signal their strength by varying their effort over time.

Appendix

A. Proof of proposition 1

- **Some general facts on All-Pay auctions**

It is useful to recall a number of facts about the equilibria of first-price all-pay auctions with incomplete information and private values. Let F_1 (F_2) be the c.d.f. of an arbitrary prior on $[0,1]$ representing player 1's (2's) valuation v (w). Assume further that $F_1(0) = F_2(0) = 0$, and that F_1 (F_2) possess positive densities f_1 (f_2). Finally, let b_1 (b_2) be the equilibrium bidding strategies and G_1 (G_2) the equilibrium bidding distribution of player 1 (2). Proofs can be found in Amann and Leininger (1996).

1. Supp $(G_1) =$ Supp (G_2).
2. G_1 (G_2) is continuous on its support with $b_1(1) \leqq 1$ $(b_2(1) \leqq 1)$.
3. Let $v > v'$, and $b = b_1(v)$, $b' = b_1(v')$. Then $G_2(b) \geqq G_2(b')$.
4. Supp $(G_1) = \left[0, \max_{v\in[0,1]} b_1(v)\right]$, Supp $(G_2) = \left[0, \max_{w\in[0,1]} b_2(w)\right]$.
5. $\min\{G_1(0), G_2(0)\} = 0$.

- **Existence and Uniqueness of the Equilibrium with covering**

We now show that (3.1),the system of four equations defining an equilibrium with covering, has a solution provided that K is smaller than $\bar{K}$, where $\bar{K}$ solves $\bar{K}e^{\frac{1}{\bar{K}-1}} = e^{-2}/2$ ($\bar{K} \simeq 0.26 < 1/2$) and that this solution is unique.

The system cannot be solved explicitly, due to the fourth equation. It is possible however to express the unknowns as functions of β and K alone:

$$\alpha = \frac{\beta\left(2K(2-\beta) + \beta^2\right)}{4K(2-\beta)}, \quad \gamma = \frac{2K}{\beta}, \quad p = \frac{2K(2-\beta) - \beta^2}{4K(1-\beta)}.$$

Since these parameters must belong to the unit interval, we obtain $2K < \beta < (K^2 + 4K)^{\frac{1}{2}} - K$. This implies that $\beta < 1/2$, since this interval would otherwise be empty. This last equation can also be written in the form: $\beta^{3+\frac{1-2K}{1-\beta}-\frac{4K}{\beta}} = 4K^2$. Define $f(\beta) = \beta^{3+\frac{1-2\cdot K}{1-\beta}-\frac{4K}{\beta}} - 4K^2$, $x_- = 2K$ and $x_+ = (K^2+4K)^{\frac{1}{2}} - K$. Note that x_- is a root of f, but does not satisfy the previous constraints. An equilibrium with covering exists if and only if there exists a root of f in (x_-, x_+). Note that $f(x_-) = 0$ and $f(x_+) = 4K^2\left(3 + 8K + 2K^2 - 2(2+K)(K(4+K))^{\frac{1}{2}}\right) \geqq 0$ since $3 + 8K + 2K^2 - 2(2+K)(K(4+K))^{\frac{1}{2}}$ is decreasing in K and is equal to zero for $K = 1/2$. Finally, $f'(x_-) = 4K\left(1 + \frac{(1-K)\ln 2K}{1-2K}\right)$. We obtain $f'(x_-) < 0 \Leftrightarrow K \leqq \bar{K}$, where $\bar{K}$ solves:

$$\bar{K}e^{\frac{1}{\bar{K}-1}} = \frac{e^{-2}}{2}.$$

This is therefore a sufficient condition for the existence of a root of f lying in the admissible interval. Necessity follows from the variations of f studied in the following paragraph establishing uniqueness. It is easy to verify that $f'(\beta) = \beta^{\frac{(\beta-2\cdot K)\cdot(\beta-2)}{\beta\cdot(\beta-1)}} \cdot h(\beta)$, where h is defined by:

$$h(\beta) = \frac{(\beta-1)\left(3\beta^2 + 4K - 4\beta - 2\beta K\right) + \left(\beta^2 + 4K - 8\beta K + 2\beta^2 K\right)\ln\beta}{(1-\beta)^2}.$$

By algebra we find that that $h'(\beta)$ equals:

$$\frac{1}{\beta(\beta-1)^3}\left[(\beta-1)\left(\beta\left(4 - 5\beta + 3\beta^2\right) + 2K\left(2 - 5\beta + \beta^2\right)\right) + 2\beta^2(2K-1)\ln\beta\right].$$

Since $\ln\beta \geqq 1 - 1/\beta$ on $\beta \in (0,1)$, it follows that the expression in square bracket is less than $(1-\beta)^2\left(\beta^2 + 2\beta K - 4K + 2\beta(\beta-1)\right)$ which is less than zero. [To see this, recall that only values of β less than $(K^2+4K)^{\frac{1}{2}} - K$ can be part of an equilibrium, which is equivalent to $\beta^2 + 2\beta K - 4K \leqq 0$. Also, $2\beta(\beta-1) < 0$]. Hence, h' is positive and h is increasing over the relevant interval. Since f is decreasing at $2K$ and a root exists in the interval of interest, it must be the only root within that interval.

Let us finally show that there cannot be other equilibria with covering. In particular, we need to show that there is no equilibrium in which high valuations randomize in a non-uniform way. Suppose high types randomize between high and low opening bids with probability $p(v)$. The distributions of types in the subgame following a covered bid are:

$$\begin{aligned} F_1(v) &= (\beta-\alpha)+\int_\beta^v p(s)\,ds \text{ for } v\in[\beta,1] \\ F_2(w) &= \frac{w-\gamma}{1-\gamma} \text{ for } w\in[\gamma,1]. \end{aligned}$$

The distributions of types in the subgame after an opening bid of zero are:

$$\begin{aligned} F_1(v) &= \alpha+\int_\beta^v (1-p(s))\,ds \text{ for } v\in[\beta,1] \\ F_2(w) &= w \text{ for } w\in[0,1]. \end{aligned}$$

The O.D.E. corresponding to both subgames are:

$$\begin{aligned} h_K'(v) &= \frac{h_K(v)\cdot p(v)}{v\cdot(1-\gamma)} \\ h_0'(v) &= \frac{h_0(v)\cdot(1-p(v))}{v}. \end{aligned}$$

Since profits equal the integral of the probability of winning (by the envelope theorem), we have:

$$\begin{aligned} \Pi_0(v) &= \Pi_0(\beta)+\int_{h_0(\beta)}^{h_0(v)} ds \\ \Pi_K(v) &= \Pi_K(\beta)+(1-\gamma)\int_{h_K(\beta)}^{h_K(v)} \frac{1}{1-\gamma}ds. \end{aligned}$$

The high types are willing to randomize only if their profits are the same in both subgames. This means that the first derivative of the profits must also be equal. We get:

$$\Pi_0'(v) = h_0(v) = \Pi_K'(v) = h_K(v).$$

The mapping $h(\cdot)$ must therefore be the same accross subgames. But then from the two O.D.E. we get:

$$\frac{h_K'(v)}{h_K(v)} = \frac{p(v)}{v \cdot (1-\gamma)} = \frac{h_0'(v)}{h_0(v)} = \frac{1-p(v)}{v}.$$

Rearranging, we get:

$$\frac{p(v)}{1-p(v)} = (1-\gamma).$$

This proves that the probability of randomization $p(\cdot)$ does not depend on types and is constant on $[\beta, 1]$. It is straightforward to verify the necessity of the other equations of the system (1) ■

B. Proof of proposition 2

- **Definition of P.S.E.**

Definition 1. *A P.B.E. is a Perfect Sequential Equilibrium (P.S.E.) if, for all players j, and all their possible deviations, there exists no P.B.E. of the subgame following the deviation, with beliefs ψ_j and ψ_i immediately prior to the deviation, and beliefs ϕ_j and ϕ_i after the deviation such that:*

1. *$\phi_j(t) = \psi_j(t)$ for all $t \in T_i$,*
2. *$\phi_i(t) = 0$ if $\psi_i(t) = 0$ or if $t \in T_j$'s expected payoff in the P.B.E. of the subgame (following the deviation) is strictly smaller than his expected payoff in the original P.B.E.*[5]*,* **and** *$\phi_i(t) > 0$ if $\psi_i(t) > 0$ and $t \in T_j$'s expected payoff in the P.B.E. of the subgame is strictly larger than his expected payoff in the original P.B.E.,*

[5] *Here and in the remainder of the definition, the expected payoff in the original P.B.E. should be understood as Player 1's expected payoff,when he follows the strategy precribed in the original P.B.E, conditional on the node where the considered deviation occurs being reached.*

3. $\frac{\phi_i(t)}{\psi_i(t)} \geqq \frac{\phi_i(t')}{\psi_i(t')}$ *whenever* $\phi_i(t') > 0$ *and* $\phi_i(t) > 0$, *for* $t \in T_j$ *whose payoff in the P.B.E. of the subgame is strictly larger than his expected payoff in the original P.B.E., with equality if* $t' \in T_j$*'s payoff in the P.B.E. of the subgame is strictly larger than his expected payoff in the original P.B.E..*

Condition (1) states that the deviator should not revise his beliefs, since he has not learnt anything about his opponent. Condition (2) put restrictions on the support of the beliefs to be considered: this support should include players who are strictly better off in the P.B.E. following the subgame, given those beliefs, than in the original P.B.E., and exclude those who are strictly worse off. Condition (3) states that, except possibly for deviators that are indifferent to the deviation, whose likelihood may possibly decrease, the deviators' relative likelihood should not be altered.

- **Proof of the structure of P.S.E.**

$K \in [1/2, 1]$: the only P.S.E. involves no opening bid. The strategies of the non-revealing equilibria given in the text do indeed form a P.B.E. and since sequential optimality does not restrict these equilibria, these equilibria are P.S.E.. Any deviation from those strategies is clearly not profitable.

Consider next a P.B.E. with $K < 1/2$ where no player 1 is willing to bid K. We show that such an equilibrium cannot be a P.S.E.. Consider a deviation by all players with $v \in [\alpha, 1]$ to investment level K. Suppose, first, that, for some $\gamma \in [0, 1)$, player 2 always finds it worthwhile to cover when his valuation lies in $(\gamma, 1]$. Consider the subgame between those valuations of players, and let $h : (\gamma, 1] \rightarrow [\alpha, 1]$, $v \longmapsto h(v)$ such that $w = h(v)$ make the same bid as a player with valuation v. $\beta \triangleq h(\gamma_+)$ is necessarily larger than α, since players with valuations arbitrarily close to β have profits arbitrarily close to $\frac{\beta-\alpha}{1-\alpha} \cdot \gamma$, which must exceed K. Obviously then, $\gamma > \alpha$ and $h(v) = v^{\frac{1-\gamma}{1-\alpha}}$. We have to verify that in this subgame all players with valuations in the interval $[\alpha, 1]$ achieve higher profits than in the original P.B.E., where players with valuation v earn $v^2/2$. In the subgame following the deviation, players with $v \in [\beta, 1]$

achieve profits $\Pi(v) = \Pi(\gamma) + \int_\beta^v \frac{s^{\frac{1-\gamma}{1-\alpha}} - \gamma}{1-\gamma} ds$, while players with valuations in $[\alpha, \beta]$ make zero profits. Ex ante profits of valuations $[\beta, 1]$ must exceed $v^2/2$, that is:

$$\gamma \cdot v + \int_\beta^v \left(s^{\frac{1-\gamma}{1-\alpha}} - \gamma \right) ds - K > v^2/2.$$

Note that the derivative of the left-hand side with respect to v is $v^{\frac{1-\gamma}{1-\alpha}}$, which is larger than the corresponding derivative of the right-hand side which is v. Hence, the inequality will hold if it holds for $v = \beta$. Consider players with valuations in $[\alpha, \beta]$. Ex ante profits from deviating are $\gamma \cdot v - K$. Marginal profits, γ, once again exceed marginal profits v in the original P.B.E., since $v \leqq \beta < \gamma$. If players with valuation α are indifferent between deviating and not deviating, players with lower valuations will prefer not to deviate. This is equivalent to requiring that $\alpha \cdot \gamma - K = \alpha^2/2$. Hence, provided that there exist γ, α such that

$$\begin{cases} \alpha \cdot \gamma - K = \alpha^2/2, \\ \frac{\beta - \alpha}{1-\alpha} \cdot \gamma = K. \end{cases}$$

the equilibrium is not a P.S.E.

Suppose now that it is not worthwhile for player 2 to cover after a deviation, regardless of valuation. It follows that the original P.B.E. is not a P.S.E. if

$$\begin{cases} \alpha - K = \alpha^2/2, \\ \frac{e^{\alpha-1} - \alpha}{1-\alpha} < K. \end{cases}$$

($\frac{e^{\alpha-1} - \alpha}{1-\alpha} = \lim_{\gamma \to 1} \frac{\beta-\alpha}{1-\alpha} \cdot \gamma$). This system guarantees that it is indeed optimal for player 2 not to cover, regardless of valuation; that if player has a valuation in the interval $(\alpha, 1]$ he will strictly prefer the expected payoff from deviating to the original expected payoff, and that all players with valuations in the interval $[0, \alpha)$ will strictly prefer the expected payoff of the original P.B.E. to their expected payoff from deviating. Although it is not difficult to show which case obtains a function of K, this is not even necessary. It is enough to note that for $\alpha = 0$, expected profits from deviating are smaller than expected profits from not deviating, whereas in both cases, since expected profits from deviating are larger than $\alpha^2 - K$, they are also larger than $\alpha^2/2$, provided that α is close enough to 1. This result is based on the fact that if K is strictly less than $1/2$. γ being a continuous function of K, there does then necessarily exist, for any $K < 1/2$, an $\alpha \in (0, 1)$ satisfying one of the two systems.

We finally need to verify that there does not exist a P.S.E. providing assured deterrence outside the interval $[\bar{K}, 1/2]$. It is then easy to show that an equilibrium with assured deterrence, as specified in the text, is a P.S.E.. Recall that in an equilibrium with assured deterrence, there exists an $\alpha \in (0, 1)$ such that all player 1s with valuations strictly smaller than α make a zero opening bid, while all players with valuations strictly larger than α invest K. In this equilibrium player 2 never covers, regardless of valuation . In simultaneous bidding between players with valuation $v \in [0, \alpha]$ and players with $w \in [0, 1]$,after a zero opening bid , expected profits of player 1 with valuation $v = \alpha$ are equivalent to $\alpha^2/(1 + \alpha)$; since a player with valuation α will be indifferent between this expected profit and the expected profit following a bid of K, which is $\alpha - K$, it follows that $\alpha = K/(1 - K)$. Consider a deviation by player 2 in which he covers. More precisely, suppose that players with $w \in (\gamma, 1]$ will cover while players with valuations lower than $\gamma \in (0, 1)$ will not . Obviously, if players with valuations arbitrarily close to γ from above have expected profits from the deviation that are arbitrarily small, players with valuations strictly above γ obtain strictly positive expected profits from deviating while players with valuations strictly below achieve strictly negative profits. These two situations compare with the zero profits that player 2 achieves in the original P.B.E., regardless of valuation. In the subgame following the deviation by player 2. Player 1's with valuations $v \in (\alpha, 1]$ play against player 2s with $w \in (\gamma, 1]$. In this subgame, player 2 with valuation γ_+ (a valuation arbitrarily close to γ from above) can achieve profit K, only if

$$\frac{\gamma^{\frac{1-\alpha}{1-\gamma}} - \alpha}{1-\alpha} \cdot \gamma = K.$$

Hence, the equilibrium with assured deterrence is a P.S.E. if and only if such a $\gamma \in (0,1)$ cannot be found. Since the left-hand side is increasing in γ, it is both necessary and sufficient that

$$\lim_{\gamma \to 1} \frac{\gamma^{\frac{1-\alpha}{1-\gamma}} - \alpha}{1-\alpha} \cdot \gamma = \frac{e^{-(1-\alpha)} - \alpha}{1-\alpha} \leqq \frac{\alpha}{1+\alpha}.$$

The latter inequality can be rewritten as $1 + \frac{1-K}{1-2K} \cdot \ln 2K \geqq 0$, which precisely states that $K \geqq \bar{K}$. Finally, when $K > 1/2$, there is no equilibrium with assured deterrence, since when K meets this condition there is no value of α in the open unit interval which satisfies $\alpha = K/(1-K)$ Finally, equilibria with covering, as specified in the text for $K < \bar{K}$, are obviously P.S.E., since they are P.B.E. and every subgame is on the equilibrium path ■.

References

[1] Amann, E., and W. Leininger, Asymmetric all-pay auctions with incomplete information: the two-player case, *Games Econ. Behav.* **14** (1996), pp. 1-18.

[2] Bellman, R., and D. Blackwell, Some two-person games involving bluffing, *Proc. Natl. Acad. Sci.* **35** (1949), pp. 600-605.

[3] Binmore, K., "Fun and Games: A text on Game Theory", D.C. Heath and Co., Lexington, MA, 1992.

[4] Bishop, D.T., C. Cannings, A Generalised War of Attrition, *J. Theoret. Biol.* 70 (1978), pp. 85-124.

[5] Borel, E.,

[6] Dixit, A., Strategic Behavior in Contests, *American Economic Review*, **77** (**5**), (1987), pp. 891-8.

[7] Gibson, B., and Sachau, D., Sandbagging as a Self-Presentational Strategy: Claiming to Be Less Than You Are, Personality and Social Psychology Bulletin, **26** (**1**), (2000), pp. 56-70.

[8] Grossman, S., and M. Perry, Perfect Sequential Equilibrium, *J. Econ. Theory* **39** (1986), pp. 97-119.

[9] Karlin, S., and R. Restrepo, "Multistage poker models", Annals of mathematics studies No. 39, Edited by M. Dresher, A. W. Tucker, and P. Wolfe, Princeton Univ. Press, Princeton, 1957.

[10] Maynard Smith, J., The Theory of Games and the Evolution of Animal Conflicts, *J. Theoret. Biol.* **47** (1974), pp. 209-221.

[11] Maynard Smith, J., "Evolution and the Theory of Games", Cambridge University Press, Cambridge, U.K., 1982

[12] McAfee, P. and K. Hendricks, Feints, Mimeo, University of Texas, 2001.

[13] McDonald, J., "Strategy in Poker, Business and War", Norton, New York, NY, 1950.

[14] Myerson R., Optimal Auction Design, *Math. Oper. Res.* **6** (1981), pp. 58-73.

[15] Nash, J. and L. Shapley, A Simple Three-Person Poker Game, in Dimand M.A, Dimand R., eds. The foundations of game theory. Volume 2. Elgar Reference Collection. Cheltenham, U.K. and Lyme, N.H.: Elgar; 1997, pp. 13-24. Previously published: [1950].

[16] Newman, D.J., A model for real poker, *Operations Research* **7** (1959), pp. 557-560.

[17] Riechert, S.E., Games Spider Play: Behavioural Variability in Territorial Disputes, *Behav. Ecol. Sociobiol.* 3, (1978), pp. 135-162.

[18] Sakai, S., A model for real poker with an upper bound of assets, *J. Optim. Theory Appl.* **50** (1986), pp. 149-163.

[19] Rosen, S., Prizes and Incentives, *American Economic Review*, **76** (**4**), (1986), pp. 716-727.

[20] Shepperd, J.A. and R.E. Socherman, On the Manipulative Behavior of Low Machiavellians; Feigning Incompetence to "Sandbag" an Opponent, *Journal of Personality and Social Psychology*, **72** (1997), pp. 1448-1459.

[21] Von Neumann, J. and O. Morgenstern,"Theory of Games and Economic Behavior", Princeton Univ. Press, Princeton, 1944.