

Strategic Betting for Competitive Agents

Liad Wagman
Department of Economics
Duke University
Durham, NC, USA
liad.wagman@duke.edu

Vincent Conitzer
Depts. of Computer Science and Economics
Duke University
Durham, NC, USA
conitzer@cs.duke.edu

ABSTRACT

In many multiagent settings, each agent's goal is to come out ahead of the other agents on some metric, such as the currency obtained by the agent. In such settings, it is not appropriate for an agent to try to maximize its expected score on the metric; rather, the agent should maximize its expected probability of winning. In principle, given this objective, the game can be solved using game-theoretic techniques. However, most games of interest are far too large and complex to solve exactly. To get some intuition as to what an optimal strategy in such games should look like, we introduce a simplified game that captures some of their key aspects, and solve it (and several variants) exactly.

Specifically, the basic game that we study is the following: each agent i chooses a lottery over nonnegative numbers whose expectation is equal to its budget b_i . The agent with the highest realized outcome wins (and agents only care about winning). We show that there is a unique symmetric equilibrium when budgets are equal. We proceed to study and solve extensions, including settings where agents must obtain a minimum outcome to win; where agents choose their budgets (at a cost); and where budgets are private information.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems;
J.4 [Social and Behavioral Sciences]: Economics

General Terms

Economics, Theory

Keywords

Game theory, Nash equilibrium, strategic betting, contests, fair bets

1. INTRODUCTION

Agents are often evaluated according to some single-dimensional metric: for example, the amount of currency the agent has earned, the number of points the agent has scored, the number of tasks the agent has completed, *etc.* In settings with uncertainty, the design of the agent results in a probability distribution over this metric. The agent's designer must optimize this distribution in some way. Most

Cite as: Strategic Betting for Competitive Agents, Liad Wagman and Vincent Conitzer, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 847-856. Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

often, the designer will choose to maximize the expected score on the metric (expected amount of currency earned, *etc.*).

However, maximizing the expectation is not always the optimal thing to do. Most notably, agents are often in competition with each other, and the goal is to come out ahead of the other agents. For example, in the Trading Agent Competition's Supply Chain Management game [7], computer programs make decisions about managing a (simulated) supply chain. Their tasks include negotiating supply contracts, bidding for customer orders, and managing assembly and shipping. The winning agent is the one that has the most money in the bank at the end of the game. Another example is the AAI Computer Poker Tournament [19], where in the Bankroll Competition the winner is the agent with the most money in the end. Yet another example is the Penn-Lehman Automated Trading (PLAT) live competition, in which automated stock trading agents compete based on real stock market data [13]. (Unlike the other competitions, this one is no longer active, though apparently not due to lack of interest.) In games such as these, it can be tempting to try to maximize one's expected amount of money, but in fact, the only thing that matters is whether the agent made more money than the other agents. More recently, agent designers have started to take this into account (for example, [17]).

As a simple numerical example, suppose that agent 2 will certainly end up with \$50, and agent 1 has a choice between two strategies. Strategy 1 will give agent 1 \$40 with probability 100%; strategy 2 will give agent 1 \$60 with probability 50%, and \$10 with probability 50%. The expected earnings of strategy 2 are \$35, so if agent 1 aims to maximize expected earnings, it will choose strategy 1. However, if the goal is to come out ahead of agent 2, strategy 2 is the better choice, since it results in a 50% probability of winning, whereas strategy 1 results in a guaranteed loss. Situations such as these, where an agent has a choice between strategies that give roughly the same expected earnings but very different distributions over earnings, are quite common—for example, the agent may be able to place various bets in (say) a casino, which will reduce the agent's expected earnings only slightly but vastly increase the variance.

It should be noted that this is not a criticism of maximizing expected *utility*. Rather, it is a criticism of confusing earnings with utility. A sensible utility function here would give utility 1 for a win, and utility 0 for a loss. (Of course, in some settings an agent may have some residual utility for money, so that the utility function considers both whether the agent won and how much money the agent has. However, at least in the competitions described above, the predominant goal is simply to win.) There are very powerful axiomatic arguments for maximizing expected utility (for an overview, see [16]), and nothing in this paper conflicts with maximizing expected utility.

Once the utility function is correctly defined, we can in principle solve such strategic settings using game-theoretic solution concepts, for example, minimax strategies and Nash equilibria. However, games such as the competitions mentioned above are very complex, and at least using current techniques, it is intractable to solve for the game-theoretically optimal strategy—although many creative approaches have been proposed to compute a strategy that is close to an optimal strategy, both for the Trading Agent Competition [21, 22, 20] and computer poker [18, 5, 11, 12]). We will make no attempt at solving them in this paper. Instead, we will study a much simpler game that nevertheless illustrates many of the key phenomena in these competitive settings.

The most basic version of the game can be described as follows. Two agents, Alice and Bob, each have a budget of chips for gambling. They each (simultaneously) place a single bet in (say) a casino. (We will assume that the outcomes of the bets are independent.) Whoever ends up with more chips is named the winner, and chips are worthless afterwards. What bets should Alice and Bob place?

To answer this question, we need to know what bets the casino is willing to accept. Let us assume that, driven by competition, the casino is willing to accept any *fair* bet.¹ That is, an agent can buy any *lottery* (probability distribution) over nonnegative real numbers whose expectation is equal to the agent’s budget.

Incidentally, if an agent were able to place a *sequence* of bets, where the choice of later bets is allowed to depend on the outcomes of the agent’s own earlier bets (but not on the outcomes of the other agent’s bets), this would make no difference to the game, for the following reason. Any plan (strategy) for betting will result in a (single) probability distribution over nonnegative numbers with expectation equal to the agent’s budget, and thus the agent can simply choose this lottery as a single bet.

In this paper, we study the equilibria of (the n -agent version of) this game, as well as variants in which agents must end up with at least a certain number of chips to win; in which agents have to first buy chips; and in which budgets are private information. There is good reason to believe that the equilibrium distributions of these games bear some resemblance to the equilibrium distributions over earnings in the agent competitions mentioned above. For example, in the stock trading competition mentioned above, (say) in the last day of trading, the agent can choose a portfolio that will result in a particular distribution over earnings at the end of the day. The expected value of this portfolio at the end of the day will be roughly the same as its value at the beginning;² however, the space of possible distributions is very large, especially if it is possible to hold derivatives such as call and put options.³ Again, the goal in the competition is simply to come out ahead of the other agents. Because the equilibrium distributions in these competitions are likely to be similar to those in the abstract game(s) in this paper, one can use our results in the following way: when creating an agent for one of these competitions, choose strategies that produce approximately the optimal distribution for the game(s) studied in this paper. Indeed, the equilibrium of our game would suggest to

¹Real-world casinos typically have payback rates of at least 90%.

²Unlike in casinos, in the stock market riskier distributions tend to have a slightly *greater* expected value.

³One issue here that is not modeled in this paper is that the values of the agents’ portfolios can be correlated, for example because they hold the same stock, or because the values of different stocks are correlated (as they typically are). However, it is at least possible to create portfolios that are roughly independent, for example by investing in small companies for which most of the risk is due to company-specific factors (diversifiable risk).

hold quite risky portfolios in the stock trading competition—which makes intuitive sense, as the goal is to come out ahead of the others.

While we have motivated our results from a multiagent systems perspective, they are also relevant to the study of several standard settings in economics. For example, previous research in economics has considered the strategic choice of lotteries as a means to characterize incentives for risk-taking in R&D environments. Here, a choice of technology leads to a distribution over the final quality (or improvement in quality) of the product, which determines which firm will dominate the market [1, 4, 6]. Patent races constitute another application, where again the choice of technology leads to a distribution over the level of innovation, and the patent is awarded to the agent with the greatest innovation; however, here, there is typically also a minimum level of innovation that needs to be reached in order to obtain the patent [8, 9]. (Later in the paper, we will study the variant of our game where agents must obtain at least a certain value to win.) Other applications include political campaigns and arms races.

In a working paper, Dulleck *et al.* [10] (independently) propose what is effectively the same game as the basic setting that we initially study in this paper, in a different context. They study all-pay auctions in which each bidder is budget constrained, has no opportunity cost for their budget, and has access to a fair insurance market. (An all-pay auction is an auction in which each agent must pay its bid, even if it did not win. For an overview on all-pay auctions, see [3]. “Access to a fair insurance market” means that agents can place any fair bet.) Dulleck *et al.* are motivated in part by a result by Laffont and Robert [15], who study the optimal (revenue maximizing) auction when bidders face (common knowledge) financial constraints. Laffont and Robert show that the optimal auction in this case takes the form of an all-pay auction. Because of the equivalence of the games, all of our results also apply to this particular type of all-pay auction. It must be admitted that this is not a very common model of an all-pay auction (especially because bidders do not care about how much money they have left in the end), and our results do not seem to have direct applications to more common all-pay auction models. Dulleck *et al.* consider different questions from the ones in this paper, and consequently their results are complementary to ours. They give an equilibrium for the case of two agents whose budgets are not necessarily equal (our Example 2) and prove that this equilibrium is unique. They also show that with n agents, an equilibrium exists. In addition, they extend their results to allow for multiple prizes—a setting that we will not study in this paper.

The remainder of our paper is organized as follows. In Section 2, we present the basic game and solve three examples. In Section 3, we show that when agents have equal budgets, there is a unique symmetric equilibrium (which we provide explicitly). We exhibit some properties of this equilibrium, and we also show that under certain restrictions on the lotteries, the symmetric equilibrium is the unique equilibrium of the equal-budget game. In Section 4, we extend our symmetric equilibrium characterization to the case where agents must surpass a minimum necessary outcome in order to win. In Section 5, we study an extension of the basic game in which agents must first select their budgets (which come at a cost). In Section 6, we study an incomplete-information variant in which agents do not know the other agents’ budgets.

2. THE BASIC GAME

Let there be n agents, and let agent $i \in \{1, \dots, n\}$ be endowed with budget b_i , which is common knowledge. (In Section 6, we extend the model to allow private budgets.) The basic game consists of two periods. In the first period, each agent (simultane-

ously) selects any fair lottery over nonnegative real numbers.⁴ We describe a lottery by its cumulative distribution function (CDF) $F(x) : \mathbb{R}^{\geq 0} \rightarrow [0, 1]$. That is, for any x , $F(x)$ is the probability that the realized lottery outcome is less than or equal to x . Agent i 's lottery F_i is *fair* if its expectation is equal to b_i , that is, $\int_0^\infty x dF_i(x) = b_i$. Thus, a *pure* strategy for an agent in this game is any fair lottery over nonnegative numbers. Any *mixed* strategy (consisting of a distribution over lotteries—a *compound lottery* in the [2] framework) can be reduced to a pure strategy by considering its *reduced lottery*, the (simple) lottery that generates the same ultimate distribution over outcomes. Hence, we do not need to consider mixed strategies. (To eliminate any chance of confusion, because each distribution over outcomes is a pure strategy, there is no requirement that agents are indifferent among the outcomes in their supports—in fact, naturally, they will prefer the higher outcomes.)

In the second period, each lottery's outcome is randomly selected according to its corresponding probability distribution. The agent whose outcome is the highest wins. For now, we assume that agents only care about winning. Thus, without loss of generality, we assume that an agent gets utility 1 for winning and 0 for not winning, so that the game is zero-sum. (In Section 5, we extend the model to allow costly budgets.) Ties are broken (uniformly) at random. This gives rise to the following *ex ante* expected utility for agent i :⁵ $U_i(F_i, F_{-i}) = \int_0^\infty \prod_{j \neq i} F_j(x) dF_i(x)$. We will be interested in the Nash equilibria $\vec{F}^* = (F_1^*, F_2^*, \dots, F_n^*)$ of the simultaneous move game.

Example 1. Consider the game between two agents, 1 and 2, with identical budgets b . Agent 1's expected utility from playing F_1 given that agent 2 selects F_2 is $\int_0^\infty F_2(x) dF_1(x)$. Suppose that F_2 is uniform over $[0, 2b]$, so that $F_2(x) = x/2b$ for $x \in [0, 2b]$ and $F_2(x) = 1$ for $x > 2b$. Then, there is no reason for agent 1 to select a lottery that places positive probability on outcomes strictly larger than $2b$. This is because any probability placed above $2b$ can be shifted down to $2b$ without lowering agent 1's probability of winning. Then, to make the lottery fair again, mass elsewhere can be shifted up, which can only improve agent i 's expected utility. It follows that agent 1's problem is to select a distribution F_1 so as to maximize $\frac{1}{2b} \int_0^{2b} x dF_1(x)$ subject to the fairness condition (henceforth *budget constraint*) $\int_0^{2b} x dF_1(x) = b$. We note that the integral in the objective must equal b for any F_1 that satisfies the budget constraint. Hence, any such F_1 constitutes a best-response to agent 2's strategy. Thus, it is an equilibrium for each agent to select the uniform lottery $U[0, 2b]$. Moreover, because this is a two-agent zero-sum game, lottery $U[0, 2b]$ is also a minimax strategy; it guarantees the agent an expected utility of at least $1/2$. This is in contrast to the trivial strategy of just holding on to one's budget b , which can lead to an arbitrarily low expected utility: for any $\epsilon \in (0, 1)$, the opponent can put probability ϵ on 0 and probability $1 - \epsilon$ on $b/(1 - \epsilon)$, so that the opponent wins with probability $1 - \epsilon$.

Example 2. Now, consider two agents with different budgets, b_1

⁴If negative lottery outcomes are allowed, then an agent can place an infinitesimal mass on an extremely negative outcome, and distribute the rest of its mass on large positive outcomes. As a result, no equilibrium would exist.

⁵Technically, the expression is only well-defined if the distributions are continuous, that is, they have no mass points. In a slight abuse of notation, we use the same expression for distributions with mass points (as is common in the literature). It should be noted that (for example) in the two-agent case, if agent 2 has a mass point at x , so that $F_2(x) > \lim_{\epsilon \rightarrow 0} F_2(x - \epsilon)$, then the probability for 1 of winning given that it obtains outcome x is not $F_2(x)$, but rather $\lim_{\epsilon \rightarrow 0} F_2(x - \epsilon) + (F_2(x) - \lim_{\epsilon \rightarrow 0} F_2(x - \epsilon))/2$. This is only relevant if agent 1 also has a mass point at x .

and b_2 , and without loss of generality suppose that $b_1 < b_2$. Suppose that agent 2's strategy F_2 is the uniform lottery $U[0, 2b_2]$. First, we note that similarly to Example 1, there is no reason for agent 1 to select a lottery that places probability on outcomes strictly larger than $2b_2$. Thus, agent 1's problem is to select F_1 to maximize $\int_0^{2b_2} \frac{x}{2b_2} dF_1(x)$ subject to $\int_0^{2b_2} x dF_1(x) = b_1$. As before, any F_1 that satisfies the constraint constitutes a best-response for agent 1. Now, consider the following compound lottery F_1 :

1. Choose the lottery that with probability b_1/b_2 generates outcome b_2 , and with probability $1 - b_1/b_2$ generates outcome 0.
2. If outcome b_2 was generated, then subsequently choose the lottery $U[0, 2b_2]$.

Formally, $F_1(x) = 1 - b_1/b_2 + (b_1/b_2)(x/2b_2)$ over $[0, 2b_2]$. That is, agent 1's lottery has a probability mass at 0. (p is a *mass point* of a cumulative distribution function F if $\lim_{\epsilon \rightarrow 0} F(p + \epsilon) - F(p - \epsilon) > 0$.) Lottery F_1 satisfies the constraint, and is thus a best response to F_2 . Now, consider agent 2's problem given that agent 1 uses F_1 . With probability $1 - b_1/b_2$, agent 1 gets 0 (and given this, agent 2 wins with probability 1, as long as agent 2 does not have a mass point at 0), and with probability b_1/b_2 , agent 2 faces the lottery $U[0, 2b_2]$. Since we have already determined that $U[0, 2b_2]$ is a best response against $U[0, 2b_2]$, it follows that $U[0, 2b_2]$ is a best response against F_1 . Thus, we have found an equilibrium. Again, because this is a two-agent zero-sum game, the agents' strategies are also minimax strategies. Figure 1 shows the equilibrium strategies graphically.

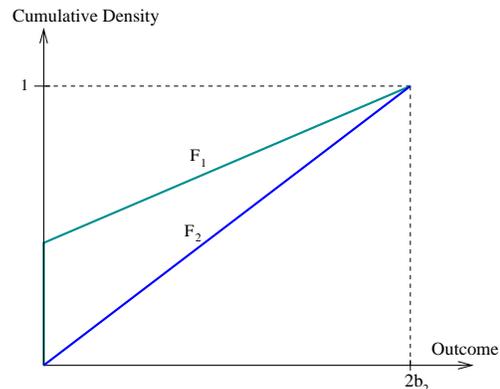


Figure 1: Equilibrium strategies in Example 2

Since agent 1 has a chance of winning only if it won its initial gamble, after which it has the same budget as agent 2, its probability of winning is $b_1/2b_2$. We note that agent 2's equilibrium strategy does not depend on b_1 (as long as $b_1 \leq b_2$). In contrast, agent 1's equilibrium strategy does depend on b_2 , because it places an initial, all-or-nothing gamble to "even the odds" and reach b_2 . [10] also study Examples 1 and 2, and show that the equilibrium described here is the unique equilibrium in each case.

Example 3. Now, suppose there are three agents with identical budgets b , and consider the lottery F such that $F(x) = (3b)^{-\frac{1}{2}} x^{\frac{1}{2}}$ over $[0, 3b]$. Given that agents 2 and 3 employ strategy F , there is no reason for agent 1 to allocate mass to outcomes larger than $3b$. Thus, agent 1's problem is to select F_1 to maximize $\int_0^{3b} F^2(x) dF_1(x) = \frac{1}{3b} \int_0^{3b} x dF_1(x)$ subject to $\int_0^{3b} x dF_1(x) = b$. As in Example 1, any lottery that satisfies the constraint is a best response. In particular, playing F is a best response for agent 1. Hence, (F, F, F) is a symmetric equilibrium. In Section 3.2 we will illustrate how symmetric equilibrium strategies change as the number of agents increases.

3. CHARACTERIZING EQUILIBRIA OF THE EQUAL-BUDGET GAME

In this section, we will study the case where all n agents have the same budget $b > 0$. We refer to this setting as the *equal-budget game*. We will show that this game has a unique symmetric equilibrium. We also show that under certain conditions on the set of strategies, there are no other equilibria.

3.1 Properties of best responses

In this subsection, we prove that any best response in our setting (even in games with unequal budgets) must have certain properties. These properties will be useful in the remainder of this section, where we analyze the equilibria of the equal-budget game.

Consider agent i . Let $F_{-i}(x)$ be the probability that all agents other than i obtain an outcome below x : $F_{-i}(x) = \prod_{j \neq i} F_j(x)$. The first three lemmas show that if i is best-responding, then F_{-i} must be linear in the support of F_i . (If this is not the case, then i is better off changing its distribution, as we will show.) For given $x_1 < x_2 < x_3$, Lemma 1 considers what happens if agent i shifts probability from (around) x_2 to x_1 and x_3 , in an expectation-preserving way. If agent i is best-responding, this cannot leave them better off, and this imposes some constraints on F_{-i} .

LEMMA 1. Consider $x_1, x_2, x_3 \in \mathbb{R}^{\geq 0}$ such that $x_1 \leq x_2 \leq x_3$. Suppose that F_{-i} is continuous at x_2 , and let F_i be a best response for i to F_{-i} . If x_2 is in the support⁶ of F_i , then the following inequality holds:

$$(x_2 - x_1)F_{-i}(x_3) + (x_3 - x_2)F_{-i}(x_1) \leq (x_3 - x_1)F_{-i}(x_2)$$

Due to space constraint, we omit all the proofs; a full version is available upon request. Nevertheless, to get some intuition for why Lemma 1 is true, suppose that F_i has mass points at x_1, x_2, x_3 . Suppose we modify F_i by shifting ϵ mass from x_2 to x_1 and x_3 . To preserve the expected value of the distribution, it must be that the mass shifted to x_1 is $\epsilon(x_3 - x_2)/(x_3 - x_1)$, and the mass shifted to x_3 is $\epsilon(x_2 - x_1)/(x_3 - x_1)$. Since we assumed F_i is a best response, this modification cannot have increased the probability that i wins. Hence, it must be that $F_{-i}(x_2)\epsilon \geq F_{-i}(x_1)\epsilon(x_3 - x_2)/(x_3 - x_1) + F_{-i}(x_3)\epsilon(x_2 - x_1)/(x_3 - x_1)$, which is equivalent to the expression in the Lemma. (The formal proof addresses the general case where F_i does not necessarily have mass points.)

Whereas Lemma 1 considers shifting probability mass from outcome x_2 to x_1 and x_3 , Lemma 2 considers the opposite. Intuitively, if outcomes x_1 and x_3 are in the support of F_i , then agent i should not find it profitable to redistribute mass from (around) x_1 and x_3 to x_2 in an expectation-preserving way.

LEMMA 2. Consider $x_1, x_2, x_3 \in \mathbb{R}^{\geq 0}$ such that $x_1 \leq x_2 \leq x_3$. Suppose that F_{-i} is continuous at x_1 and x_3 , and let F_i be a best response for i to F_{-i} . If x_1 and x_3 are in the support of F_i , then the following inequality holds:

$$(x_2 - x_1)F_{-i}(x_3) + (x_3 - x_2)F_{-i}(x_1) \geq (x_3 - x_1)F_{-i}(x_2)$$

Lemma 3 follows immediately from Lemmas 1 and 2, establishing that F_{-i} must be linear in the support of F_i if i is best-responding.

LEMMA 3. Consider $x_1, x_2, x_3 \in \mathbb{R}^{\geq 0}$ such that $x_1 \leq x_2 \leq x_3$. Suppose that F_{-i} is continuous at these outcomes and let F_i be a best response for i to F_{-i} . If x_1, x_2 , and x_3 are in the support of F_i , then the following equality holds:

$$(x_2 - x_1)F_{-i}(x_3) + (x_3 - x_2)F_{-i}(x_1) = (x_3 - x_1)F_{-i}(x_2)$$

Finally, we prove that the support of any best-response strategy has an upper bound (unless the agent can win with probability 1).

⁶In our use of the word “support”, the support is a closed set, that is, we include all the limit points in the support.

LEMMA 4. Given F_{-i} , suppose that there is no strategy for i such that i wins with probability 1. Then the support of any best response strategy F_i for i has an upper bound.

The intuition behind Lemma 4 is the following. Shifting probability mass that is placed on sufficiently large outcomes downwards slightly will not decrease the probability of winning significantly. Doing so will allow the agent to shift mass on lower outcomes upwards, where this is more fruitful.

3.2 Symmetric equilibria with equal budgets

In the remainder of this section, we restrict attention to the equal-budget game. First, in this subsection, we characterize the symmetric equilibria of this game. The results we obtained in Subsection 3.1 assume that F_{-i} is continuous (at certain points). The following lemma and corollary establish that in a symmetric equilibrium, this assumption is trivially satisfied.

LEMMA 5. Consider the equal-budget case. Suppose that the strategy profile in which all agents play lottery F constitutes a (symmetric) equilibrium. Then F has no mass points.

Intuitively, if F had a mass point, then an agent would find it beneficial to deviate by shifting this mass up infinitesimally (to avoid a tie) and shifting mass down elsewhere. Since F is a cumulative distribution function with no mass points, F is continuous. F_{-i} is the product of continuous functions, and is thus continuous as well. We thus have the following corollary:

COROLLARY 1. In the equal-budget game, suppose that the strategy profile in which all agents play F constitutes a symmetric equilibrium. Then F is continuous. Furthermore, F_{-i} is continuous for all i .

We now show 0 is in the support of any symmetric-equilibrium strategy.

LEMMA 6. Consider the equal-budget game. Suppose that the strategy profile in which all agents play F constitutes a symmetric equilibrium, and that the greatest lower bound of the support of F is l . Then $l = 0$.

To give some intuition, consider the following. If all agents playing F constitutes a symmetric equilibrium and $l > 0$, then an agent’s expected utility given that it obtained an outcome in a close neighborhood of l is near 0. Hence, it is beneficial to reallocate mass in a neighborhood of l to 0 and to some higher outcomes, contrary to the equilibrium assumption. We are now ready to derive the main result of this section.

THEOREM 1. The equal-budget game has a unique symmetric equilibrium. It is for all agents to select the following lottery:

$$F(x) = (nb)^{-\frac{1}{n-1}} x^{\frac{1}{n-1}} \quad (1)$$

over support $[0, nb]$.

If all agents use the lottery described in (1), then for every agent i , F_{-i} is the uniform distribution over $[0, nb]$. Hence, any lottery over outcomes in $[0, nb]$ is a best response. Figure 2 shows how the symmetric equilibrium strategy changes with the number of agents.

A random variable that is of particular interest is the *maximum* outcome. This variable is especially interesting when we interpret the game as a model for competitive R&D, where lotteries correspond to technologies that can be used and outcomes correspond to qualities of products. In this setting, the maximum outcome corresponds to the quality of the best product—the one that will dominate the market. The cumulative distribution of the maximum outcome in equilibrium is $(F(x))^n$, and its expectation is:

$$E[x_{max}] = \int_0^{nb} x d(F(x))^n = \frac{n^2 b}{2n-1} > \frac{nb}{2}$$

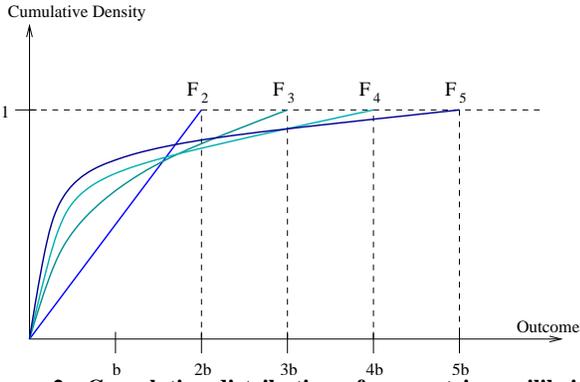


Figure 2: Cumulative distribution of symmetric equilibrium strategy for different values of n , given equal budgets $b = 5$.

This expectation is quite high, in the following sense. Suppose that we did not impose any strategic constraints on F_i . Then, $E[x_{max}] \leq E[\sum_i x_i] = \sum_i E[x_i] = nb$. That is, the expected value of the maximum outcome in equilibrium is within a factor 2 of the highest expectation that can be obtained without any equilibrium constraint. (Incidentally, without the equilibrium constraint one can in fact come arbitrarily close to achieving nb , as follows. Let F_i be the distribution that places $1 - \epsilon$ mass on 0, and ϵ mass on b/ϵ . The probability that at least one agent will receive b/ϵ is $1 - (1 - \epsilon)^n$, hence the expected quality of the product is $(b/\epsilon)(1 - (1 - \epsilon)^n)$, which as $\epsilon \rightarrow 0$ converges to nb .) Moreover, even if one can shift budgets among agents (in addition to prescribing their strategies), it still holds that $E[x_{max}] \leq nb$. By contrast, if each agent uses the degenerate strategy that places all the probability mass on b , we would have $E[x_{max}] = b$.

3.3 Uniqueness of the symmetric equilibrium

Is the symmetric equilibrium unique, or do asymmetric equilibria exist? In this subsection, we show that under mild restrictions on the strategy space, the former is the case. (We currently do not know whether these restrictions are necessary for this to be true.) Specifically, we consider the following restrictions: **(A1)** Supports have no gaps, **(A2)** F_i has no mass points for all $i \in \{1, \dots, n\}$. The next lemma shows that if (A1) holds, then all agents have 0 in their support.

LEMMA 7. *Suppose that $\vec{F}^* = (F_1^*, F_2^*, \dots, F_n^*)$ is an equilibrium strategy profile of the equal-budget game and that (A1) is satisfied. Then 0 is in the support of F_i^* for all $i \in \{1, 2, \dots, n\}$.*

We are now ready to present the main result of this subsection.

THEOREM 2. *Given (A1) and (A2), the unique equilibrium of the equal-budget game is the symmetric equilibrium described in Theorem 1.*

4. EXTENSION: MINIMUM OUTCOME REQUIREMENT

In this section, we add one feature to the equal-budget game from the previous section: in order to win, agents must end up with an outcome that is at least as high as some threshold. In other words, the winning agent must obtain the highest outcome among all agents, as well as reach or exceed some minimum outcome. If no agent reaches this threshold, then no agent receives anything. (We note that the game is no longer zero-sum.) Let us denote this threshold by r , where $r > 0$. For example, in a stock trading competition, there may be a specification that if a contestant does not outperform a risk-free asset, then the contestant cannot win. Similarly, in a trading agent competition, there may be a specification that if

no agent has positive profit (which does sometimes happen: for example, in the early rounds of the 2003 Supply Chain Management TAC, as described in [14]), then nobody wins. Also, in the patent race application from economics, a minimum level of innovation must be reached or surpassed for a patent to be granted.

We wish to solve for the symmetric equilibrium of this modified equal-budget game. We will make use of the following observations. First, it is never in agents' interest to select lotteries that place mass on outcomes in $(0, r)$. This is because outcomes in this interval can never lead to winning, so an agent would always be better off reallocating mass from this interval to 0 and to outcomes larger than r . Second, Lemmas 3, 4, and 6 still hold in this context. Moreover, Lemma 3 can be extended to hold at 0 even when F_{-i} is discontinuous there, because outcomes close to 0 can never lead to winning when $r > 0$. (We call this the "extended" Lemma 3.) Third, Lemma 5 also holds, but only over outcomes that are at or above r . Agents may have a mass point at 0.

4.1 The two-agent equal-budget game with a minimum necessary outcome

Let us begin by solving for the symmetric equilibrium of the two-agent equal-budget game. By the above discussion, for some $h \geq r$, the support of the symmetric strategy will be contained in $\{0\} \cup [r, h]$. (Let h be the smallest number for which this holds.) The next lemma shows that r must be in the support.

LEMMA 8. *Consider the equal-budget game with a minimum necessary outcome of r . Suppose that the strategy profile in which all agents play F constitutes a symmetric equilibrium. Let S denote the support of F , and let l be the greatest lower bound of $S - \{0\}$. Then $l = r$.*

Intuitively, the reason for this result is as follows. Suppose $l > r$. Then, outcomes in a close neighborhood of l have a significant chance of leading an agent to winning only if all other agents obtain outcome 0. Because of this, outcome r provides almost the same probability of winning as these outcomes. Thus, shifting mass from a neighborhood of l to r does not have a large impact on an agent's probability of winning, while it allows the agent to shift some mass to higher outcomes. For sufficiently small neighborhoods of l , doing so increases the agent's probability of winning. Therefore, r must be the greatest lower bound of $S - \{0\}$.

Lemmas 3, 5, and 8 imply that any symmetric equilibrium strategy has the form $F(x) = a + cx$ over $[r, h]$, where a and c are positive constants. Furthermore, this strategy may place a mass $m > 0$ at 0 (so that $F(r) = m$). The following claim establishes that for $x \in [r, h]$, $F(x)$ must lie on a line originating from the origin.

CLAIM 1. *In the two-agent equal-budget game with a minimum necessary outcome of r , there is some c so that for $x \in [r, h]$, $F(x) = cx$. (That is, $a = 0$.)*

Since $F(r) = m$, it holds that $m = cr$. In addition, since $F(h) = 1$, we have that $h = c^{-1}$. Finally, the budget constraint requires $\int_r^{c^{-1}} x dF(x) = b$. Substituting for F in the constraint and rearranging, we obtain $c(b, r) = \frac{\sqrt{b^2 + r^2} - b}{r^2}$. Thus, the unique candidate symmetric equilibrium strategy is for each agent to select the lottery specified by

$$F(x) = \begin{cases} \frac{\sqrt{b^2 + r^2} - b}{r} & \text{if } 0 \leq x < r \\ \frac{\sqrt{b^2 + r^2} - b}{r^2} x & \text{if } r \leq x \leq \frac{r^2}{\sqrt{b^2 + r^2} - b} \\ 1 & \text{if } x > \frac{r^2}{\sqrt{b^2 + r^2} - b} \end{cases} \quad (2)$$

It remains to verify that (2) indeed constitutes an equilibrium strategy. To check this, suppose agent 1 employs strategy F . Given this, agent 2 would not find it optimal to place mass on outcomes higher than $c(b, r)^{-1}$. Thus, agent 2's problem is to choose lottery F_2 to maximize $\int_r^{c(b, r)^{-1}} F(x) dF_2(x) = c(b, r) \int_r^{c(b, r)^{-1}} x dF_2(x)$ subject to $\int_0^{c(b, r)^{-1}} x dF_2(x) = b$. For any F_2 that satisfies the constraint and places no mass on $(0, r)$, $\int_r^{c(b, r)^{-1}} x dF_2(x)$ equals b , so the objective becomes $c(b, r) \cdot b$. Hence, any such F_2 is a best response, including F . Figure 3 shows how the symmetric equilibrium strategy varies as r increases.

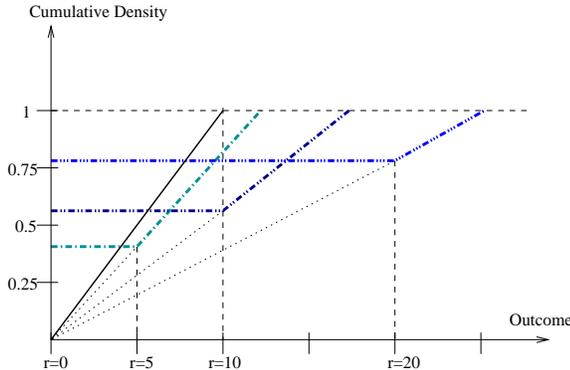


Figure 3: Cumulative distribution of symmetric equilibrium strategies for different values of r , given equal budgets $b = 5$.

We can observe the following facts about the equilibrium strategies from (2) and Figure 3. First, as r approaches 0, $c^{-1}(b, r)$ approaches $2b$, so that we converge to the equilibrium of Example 1. Second, $c(b, r)$ is decreasing in r , so that, as r grows larger, the cumulative distribution of the lottery chosen over outcomes larger than r becomes flatter. Meanwhile, the mass m at 0 approaches 1. Thus, the equilibrium strategy becomes ever riskier as r increases.

4.2 The n -agent equal-budget game with a minimum necessary outcome

We now extend the equilibrium result to n agents.

THEOREM 3. *In the n -agent equal-budget game with a minimum necessary outcome of r , the unique symmetric equilibrium strategy is for each agent to play F described by*

$$F(x) = \begin{cases} m(b, r) & \text{if } x < r \\ (c(b, r)x)^{\frac{1}{n-1}} & \text{if } x \in [r, (c(b, r))^{-1}] \\ 1 & \text{if } x > (c(b, r))^{-1} \end{cases}$$

where $m(b, r) = (c(b, r)r)^{\frac{1}{n-1}}$ and $c(b, r)$ is implicitly and uniquely defined by $\frac{1}{n}(c^{-1} - c^{\frac{1}{n-1}} r^{\frac{n}{n-1}}) = b$.

As in the two-agent game, it can be verified that $c(b, r)$ is increasing in r . Also, as r approaches 0, $c(b, r)$ approaches $1/nb$, so that F becomes the unique symmetric equilibrium strategy described in Theorem 1. Figure 4 shows how the symmetric equilibrium strategy changes as n increases.

Figure 4 resembles Figure 2 (where there is no minimum outcome requirement). One additional effect that the minimum outcome requirement introduces is that as n gets larger, the mass that the equilibrium strategy places on 0 increases—in fact, this mass converges to 1 as $n \rightarrow \infty$.

5. EXTENSION: COSTLY BUDGETS

In this section, we study a variant in which agents can choose their budgets at the beginning of the game, and each budget comes

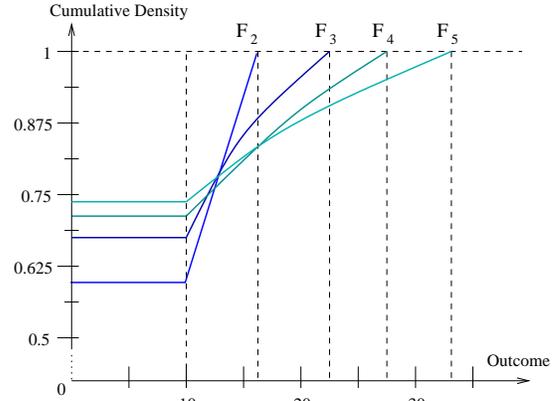


Figure 4: Cumulative distribution of symmetric equilibrium strategies for different values of n , given equal budgets $b = 5$ and $r = 10$.

at a cost. After the budgets have been chosen, the game proceeds as before. Thus, in the first period, agents choose their budgets b_i ; in the second period, they choose their lotteries F_i (whose expectation must equal b_i); and in the third period, outcomes are drawn from the lotteries and the winner is determined. An agent's utility is $-b_i$ if it does not win, and $D - b_i$ if it does win, where D is a constant. Agents try to maximize expected utility. This variant is especially natural in many of the applications in economics, where agents must make some initial investment. We only consider the 2-agent case, and we also do not consider the possibility of a minimum necessary outcome.

To solve this game, we apply backward induction. Suppose agent i has chosen budget b_i in the first period. To solve the subgame starting at the second period, we make use of the equilibrium derived in Example 2 (which, by the work of Dulleck *et al.* [10], is unique). Assume without loss of generality that $b_1 \leq b_2$. (Even though the game is symmetric at the beginning, the agents may choose different budgets in the first period.) From Example 2, we know that it is an equilibrium for agent 1 to select lottery $F_1(x) = 1 - b_1/b_2 + (b_1/b_2)(x/2b_2)$ and for agent 2 to select lottery $F_2(x) = x/2b_2$, both with supports $[0, 2b_2]$. (In fact, these are minimax strategies.) Given this, we can analyze the first period. Since the game is symmetric between agents at this point, it will suffice to focus on agent 1. Given that agent 2 has decided on budget $b_2 > 0$, agent 1's expected utility as a function of b_1 is given by

$$E[u_1(b_1, b_2)] = \begin{cases} \frac{b_1}{2b_2}D - b_1 & \text{if } b_1 \leq b_2 \\ (1 - \frac{b_2}{2b_1})D - b_1 & \text{if } b_1 > b_2 \end{cases}$$

When $b_1 \leq b_2$, agent 1's expected utility is linear in b_1 . Hence, it will choose to set $b_1 \geq b_2$ whenever $D > 2b_2$. Furthermore, by differentiating the expected utility function when $b_1 > b_2$, it can be shown that $b_1 = \sqrt{b_2 D/2}$ maximizes expected utility, given that $D > 2b_2$. (We note that in this case, indeed, $b_1 = \sqrt{b_2 D/2} > b_2$.) Moreover, it will choose to set $b_1 = 0$ whenever $D < 2b_2$, because in this case, any other budget will give it a negative expected utility. Finally, when $D = 2b_2$, any $b_1 \in [0, D/2]$ is optimal. To summarize, agent 1's (set-valued) best-response function is

$$b_1(b_2) = \begin{cases} \{0\} & \text{if } b_2 > \frac{D}{2} \\ [0, \frac{D}{2}] & \text{if } b_2 = \frac{D}{2} \\ \{\sqrt{\frac{b_2 D}{2}}\} & \text{if } 0 < b_2 < \frac{D}{2} \end{cases}$$

We note that if $b_2 = 0$, agent 1 would want to choose an infinitesimally small budget in order to win, so the best response is not well-defined in this case. Figure 5 shows the agents' best-response

curves. (To eliminate any chance of confusion, we note that the

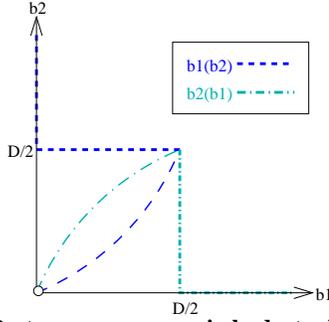


Figure 5: Best-response curves in budget selection stage

variables on the axes of this graph are budgets, not probabilities; this graph is not intended to show mixed-strategy equilibria.) The best-response curves intersect at $(D/2, D/2)$. The unique subgame perfect pure-strategy equilibrium of this game is thus for both agents to choose a budget of $D/2$ in the first period, and select the uniform lottery over $[0, D]$ in the second. Each agent's expected utility is 0 in equilibrium. This is reminiscent of the equilibrium of a common-value sealed-bid all-pay auction, where both agents choose their bids uniformly at random from $[0, D]$ (where D is the common value), leading to an expected utility of 0 for each agent. We emphasize that while the equilibria are similar, the games are quite different.

6. EXTENSION: PRIVATE BUDGETS

In this section, we consider an incomplete-information setting, where agents do not know the other agents' budgets. We consider the n -agent case, but do not consider the possibility of a minimum necessary outcome or costly budgets. Suppose that for every $j \in \{1, \dots, n\}$, agent j 's (nonnegative) budget is selected by Nature according to some commonly known prior, described by the CDF $W_j(b)$. Thus, this is a Bayesian game, and we will use Bayes-Nash equilibrium as our solution concept. Suppose that agent $j \neq i$ chooses lottery G_b^j when endowed with budget b , and consider agent i 's problem. Given b_i , agent i selects lottery F to maximize

$$\int_0^\infty \dots \int_0^\infty \prod_{j \neq i} G_{b_j}^j(x) dF(x) dW_1(b_1) \dots dW_{i-1}(b_{i-1}) dW_{i+1}(b_{i+1}) \dots dW_n(b_n)$$

subject to $\int_0^\infty x dF(x) = b_i$. Since agent i 's expected utility is bounded by 1, Fubini's Theorem allows us to change the order of integration in the objective function, which is hence equivalent to

$$\int_0^\infty \left[\int_0^\infty \dots \int_0^\infty \prod_{j \neq i} G_{b_j}^j(x) dW_1(b_1) \dots dW_{i-1}(b_{i-1}) dW_{i+1}(b_{i+1}) \dots dW_n(b_n) \right] dF(x) \quad (3)$$

Here, the bracketed expression in (3) gives the *ex ante* cumulative distribution over the maximum outcome of all agents other than i , evaluated at x . Hence, the bracketed term has a role that is analogous to the role of $F_{-i}(x)$ earlier in the paper: whereas before the uncertainty derived only from the other agents' strategies, now it derives both from the other agents' strategies and from Nature's choice of their budgets. In order to use our previous techniques for deriving equilibria, we would need this expression to be proportional to x . This is illustrated by the following two examples

of prior distributions and corresponding strategies that constitute symmetric equilibria:

1. Consider the two-agent game with identical prior $W = U[0, h]$ for some $h > 0$. One equilibrium is for both agents to acquire the degenerate lottery at b when endowed with a budget b . (This is because given these strategies, the distribution over the other agent's outcome is uniform over $[0, h]$, hence any strategy that uses only outcomes in $[0, h]$ is a best response.)

2. For some $b > 0$, let $b_L = \frac{1}{2}b$ and $b_H = \frac{3}{2}b$. In a two-agent game with an identical prior $P(b_i = b_L) = \frac{1}{2}$ and $P(b_i = b_H) = \frac{1}{2}$, $i \in \{1, 2\}$, the strategy that chooses $U[0, b]$ when $b_i = b_L$ and $U[b, 2b]$ when $b_i = b_H$, constitutes a symmetric equilibrium. (This is because given these strategies, the distribution over the other agent's outcome is uniform over $[0, 2b]$, hence any strategy that uses only outcomes in $[0, 2b]$ is a best response.)

More generally, a strategy profile $\vec{G}^* = (G^{*1}, \dots, G^{*n})$, for which for every $i \in \{1, \dots, n\}$ the bracketed term in (3) is proportional to x for all x that are used in i 's supports, constitutes an equilibrium. This is because, as in the complete-information case, the objective function reduces to the constraint for every agent. Hence, any strategy that satisfies the constraint is a best response, including that suggested by \vec{G}^* . For example, if the prior over all agents' budgets is W , with expectation k , then a strategy G that satisfies

$$\int_0^{nk} G_b(x) dW(b) = (nk)^{-\frac{1}{n-1}} x^{\frac{1}{n-1}} \quad (4)$$

for all $x \in [0, nk]$, constitutes a symmetric equilibrium. In order to obtain such a strategy, we need to be able to *transform* the prior distribution W into another distribution. Specifically, we need strategy G to map budgets in the support of the prior W to fair lotteries, so that the ensuing (expected) distribution over outcomes is as in (4). Let us say that prior CDF W is *transformable* into another CDF J if there exists a strategy G such that the ensuing distribution is J . The following theorem provides necessary conditions for a prior W to be transformable into a CDF J .

THEOREM 4. *Consider a CDF W and a CDF J , with supports contained in $\mathbb{R}^{\geq 0}$. Suppose that W is transformable into J . Then for any b in the support of W , the following two inequalities must hold:⁷ $\int_0^b x dW(x) \geq \int_0^{J^{-1}(W(b))} x dJ(x)$, and $\int_b^\infty x dW(x) \leq \int_{J^{-1}(W(b))}^\infty x dJ(x)$.*

Specifically, consider the case where the prior over each agent's budget is W , with expectation k . In order for there to exist a strategy G that satisfies $\int_0^{nk} G_b(x) dW(b) = (nk)^{-\frac{1}{n-1}} x^{\frac{1}{n-1}}$ for all $x \in [0, nk]$ (and hence constitutes a symmetric equilibrium), Theorem 4 tells us that for any budget b in the support of W , it is necessary that $E_W[x | 0 \leq x \leq b] \geq k(W(b))^{n-1}$ and $E_W[x | x > b] \leq k \sum_{j=0}^{n-1} (W(b))^j$. It is an open question whether these conditions are also sufficient for the strategy to be transformable in the desired way. However, the following theorem does provide a (more limited) sufficient condition:

THEOREM 5. *Consider a 2-agent private-budget game in which both agents' budgets are distributed according to a commonly known*

⁷If J has mass points, then $J^{-1}(W(b))$ is not necessarily defined. In this case, $\int_0^{J^{-1}(W(b))} x dJ(x)$ should be interpreted to integrate x only over the lowest $W(b)$ mass of J . Letting y be the point such that $J(y) > W(b)$ and $J(y - \epsilon) < W(b)$ for all $\epsilon > 0$, a more precise expression would be $\int_0^y x dJ(x) - (J(y) - W(b))y$. The interpretation of $\int_{J^{-1}(W(b))}^\infty x dJ(x)$ is similar.

CDF W with expectation k . If the support of W is a subset of $[k/2, 3k/2]$, then W is transformable into $U[0, 2k]$ (and hence a symmetric equilibrium exists).

Intuitively, if W 's support is a subset of $[k/2, 3k/2]$, then given any budget, an agent can choose a fair lottery over outcomes $k/2$ and $3k/2$. Since W has expectation k , choosing such lotteries results in a mass of $1/2$ at each of these outcomes. The agent can subsequently select lottery $U[0, k]$ given outcome $k/2$, and $U[k, 2k]$ given outcome $3k/2$. The resulting distribution over outcomes is $U[0, 2k]$.

7. CONCLUSIONS

In many multiagent settings, each agent's goal is to come out ahead of the other agents on some metric, such as the currency obtained by the agent. Examples include trading agent competitions, computer poker tournaments, stock trading competitions, etc. In such settings, it is not appropriate for an agent to try to maximize its expected score on the metric; rather, the agent should maximize its expected probability of winning. In principle, given this objective, the game can be solved using game-theoretic techniques. However, the games above are far too large and complex to solve exactly. To get some intuition as to what an optimal strategy in such games should look like, we introduced a simplified game that captures some of their key aspects, and solved it (and several variants) exactly. We expect that the equilibria of the large games will display some similarity to the equilibria obtained in this paper.

Specifically, the basic game that we studied is the following: each agent i chooses a lottery over nonnegative numbers whose expectation is equal to its budget b_i . The agent with the highest realized outcome wins (and agents only care about winning). We began by solving a few examples. Then, we studied the case where each agent has the same budget. We showed that there is a unique symmetric equilibrium, in which each agent chooses a lottery that randomizes over a continuum of monetary outcomes. The expectation of the highest realized outcome in this equilibrium is within a factor 2 of what could be obtained if all agents cooperated to maximize the expectation of the highest realized outcome. We also showed that under some restrictions on the lotteries, the symmetric equilibrium is the unique equilibrium of the equal-budget game.

We proceeded to study variants of the basic game. First, we extended our symmetric equilibrium characterization to the case where agents must surpass a minimum necessary outcome in order to win. Next, we studied a game in which agents first choose their budgets, which come at a cost. We found the unique pure-strategy subgame perfect equilibrium of this game, which gives the agents an expected utility of 0. Then, we introduced an incomplete-information model in which agents do not know the other agents' budgets—a common situation. We showed that our complete-information techniques can be applied to this setting if it is possible to transform the prior over budgets into the appropriate distribution over outcomes. We gave a necessary condition as well as a (more restrictive) sufficient condition for this to be possible.

Future research can take a number of specific technical directions. The most obvious directions are to extend our results to the setting of unequal budgets, as well as to investigate whether the symmetric equilibrium is the unique equilibrium of the equal-budget game (without any restrictions on the lotteries). Another important direction is to consider lottery spaces that are restricted (for example, allowing only lotteries over a discretized space), or extended with unfair lotteries. Even more generally, we can allow agents to choose lotteries that are correlated with each other. Yet another direction is to consider versions of these games in which agents may observe other agents' budgets over time. We can also

consider different utility functions: for example, the agent may also derive some utility from coming in second place. Finally, in the private-budgets setting, we left as an open question whether our necessary condition is also sufficient.

8. REFERENCES

- [1] A. Anderson and L. M. B. Cabral. Go for broke or play it safe? Dynamic competition with choice of variance. *RAND Journal of Economics*, 2007. Forthcoming.
- [2] F. J. Anscombe and R. J. Aumann. A definition of subjective probability. *Annals of Math. Statistics*, 34:199–205, 1963.
- [3] M. R. Baye, D. Kovenock, and C. G. de Vries. The all-pay auction with complete information. *Economic Theory*, 8(2):291–305, 1996.
- [4] S. Bhattacharya and D. Mookherjee. Portfolio choice in research and development. *RAND Journal of Economics*, 17(4):594–605, 1986.
- [5] D. Billings, N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and D. Szafron. Approximating game-theoretic optimal strategies for full-scale poker. *IJCAI*, 2003.
- [6] L. M. B. Cabral. Increasing dominance with no efficiency effect. *Journal of Economic Theory*, 102:471–479, 2002.
- [7] J. Collins, R. Arunachalam, N. Sadeh, J. Eriksson, N. Finne, and S. Janson. The supply chain management game for the 2007 trading agent competition. Technical Report CMU-ISRI-07-100, Carnegie Mellon University, 2006.
- [8] V. Denicolò. Patent races and optimal patent breadth and length. *J. of Industrial Economics*, 44(3):249–265, 1996.
- [9] V. Denicolò. Two-stage patent races and patent policy. *RAND Journal of Economics*, 31(3):450–487, 2000.
- [10] U. Dulleck, P. Frijters, and K. Podczeck. All-pay auctions with budget constraints and fair insurance. Working paper 0613, Department of Economics, Johannes Kepler University of Linz, Austria, 2006.
- [11] A. Gilpin and T. Sandholm. A competitive Texas Hold'em poker player via automated abstraction and real-time equilibrium computation. *AAAI*, 2006.
- [12] A. Gilpin and T. Sandholm. Better automated abstraction techniques for imperfect information games, with application to Texas Hold'em poker. *AAMAS*, 2007.
- [13] M. Kearns and L. Ortiz. The Penn-Lehman automated trading project. *IEEE Intelligent Systems*, 18(6):22–31, 2003.
- [14] C. Kiekintveld, Y. Vorobeychik, and M. P. Wellman. An analysis of the 2004 supply chain management trading agent competition. In *Workshop on Trading Agent Design and Analysis (TADA)*, 2005.
- [15] J.-J. Laffont and J. Robert. Optimal auction with financially constrained buyers. *Economic Letters*, 52:181–186, 1996.
- [16] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [17] C. McMillen and M. Veloso. Thresholded rewards: Acting optimally in timed, zero-sum games. *AAAI*, 2007.
- [18] J. Shi and M. Littman. Abstraction methods for game theoretic poker. *Computers and Games*, 333–345, 2001.
- [19] University of Alberta. American Association for Artificial Intelligence computer poker competition, 2006. <http://www.cs.ualberta.ca/~pokert/>.
- [20] Y. Vorobeychik, C. Kiekintveld, and M. Wellman. Empirical mechanism design: Methods, with application to a supply chain scenario. *ACM-EC*, 2006.
- [21] M. P. Wellman, J. Estelle, S. Singh, Y. Vorobeychik, C. Kiekintveld, and V. Soni. Strategic interactions in a supply chain game. *Computational Intelligence*, 21:1–26, 2005.
- [22] M. P. Wellman, P. R. Jordan, C. Kiekintveld, J. Miller, and D. M. Reeves. Empirical game-theoretic analysis of the TAC market games. In *Workshop on Game Theoretic and Decision Theoretic Agents (GTDT)*, 2006.