Choosing Fair Lotteries to Defeat the Competition

Wagman, Liad and Conitzer, Vincent

Duke University, Duke University, Economics Department, Duke University, Computer Science Department

14 August 2008

Choosing Fair Lotteries to Defeat the Competition

Liad Wagman* Vincent Conitzer†
Duke University Duke Unversity

August 2008‡

Abstract

We study the following game: each agent i chooses a lottery over nonnegative numbers whose expectation is equal to his budget b_i. The agent with the highest realized outcome wins (and agents only care about winning). This game is motivated by various real-world settings where agents each choose a gamble and the primary goal is to come out ahead. Such settings include patent races, stock market competitions, and R&D tournaments. 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.

JEL classifications: C70, C72, D81, L20

Keywords: Strategic gambling, Nash equilibrium, fair lotteries

1 Introduction

The most basic version of the game that we study 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 assume that the outcomes of the bets

*Department of Economics, 213 Soc Sci, Box 90097, Durham, NC 27708, USA. Tel: (919) 660-8153, Fax: (919) 684-8974, E-mail: liad.wagman@duke.edu

†Departments of Computer Science and Economics, Levine Science Research Center, Box 90129, Durham, NC 27708, USA. Tel: (919) 660-6503, Fax: (919) 660-6519, E-mail: conitzer@cs.duke.edu

‡Early versions of this work were presented at the Victor Rothschild Memorial Symposium on Economic Aspects of Communication and Information (2007), the Conference on Autonomous Agents and Multi Agent Systems (2008), the World Congress of the Game Theory Society (2008), and the Duke Microeconomics lunch seminar; we thank the participants and reviewers for many useful comments. Wagman is supported by the Program for Advanced Research in the Social Sciences Fellowship; Conitzer is supported by an Alfred P. Sloan Fellowship and by the NSF under award number IIS-0812113.

are independent. Whoever ends up with more chips is named the winner, and chips are worthless afterwards—the only goal is to win. 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 over nonnegative real numbers whose expectation is equal to the agent's budget.[2]

As an example, suppose Alice and Bob each have a budget of 10 chips. If Alice were to choose the degenerate lottery that always results in 10 chips, Bob can win most of the time by choosing the lottery that gives 11 chips with probability $10/11$, and 0 chips with probability $1/11$. In this case, Bob wins with probability $10/11$. A better response for Alice, in turn, would be to choose the lottery that gives 12 chips with probability $9/11$, and 1 chip with probability $2/11$. Alice would then win with probability $9/11 + 2/11 \cdot 1/11$. As we will see, the unique equilibrium of this game is for both Alice and Bob to choose the uniform lottery over $[0, 20]$.

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 in order to win; in which agents have to first buy chips; and in which budgets are private information.

In spite of their simplicity, games such as the above can model real-world scenarios. Previous research 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. Examples include Anderson and Cabral [2007]; Bagwell and Staiger [1990]; Bhattacharya and Mookherjee [1986]; Cabral [1994, 2002, 2003]; Judd [2003]; Klette and de Meza [1986] and Vickers [1985]. All of these earlier papers study a constrained environment in the sense that the set of possible lotteries is limited. Also, most of the previous work studies decisions that take place over time. In particular, Cabral [1994, 2002, 2003] consider an environment with two agents and two possible lotteries, a safe lottery (no variance) and a risky one (positive variance). In each period of a repeated game, agents select between those two lotteries.

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

[2] 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 finite 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.

Cabral shows that increasing dominance, a situation in which the leader advances more and more rapidly in comparison to a laggard, can be the result of the laggard choosing a riskier strategy. Judd [2003] extends this environment to continuous time. Anderson and Cabral [2007] analyze the more general choice of lottery *variance* in a continuous time setting that follows an Ito process. Both papers focus on the dynamics and welfare implications drawn from a continuous-time game with two agents. In contrast, our work focuses on the strategic choices made by agents in a static environment, where an agent's strategy choice set is larger.

Bhattacharya and Mookherjee [1986] and Klette and de Meza [1986] consider patent race models where agents select their variance. Their models consider winner-takes-all settings with two agents, where the winning agent's utility is a function of the lottery outcome and varies across agents. They show that in equilibrium, firms may take too much risk from a social-welfare point of view due to competition. In contrast, we find that in spite of competition, firms may take too little risk when compared to a risk-neutral social planner.

An important difference between our work and all of the above work is that we allow agents to select *any* fair lottery. In addition, our work abstracts from specific environments such as patent or R&D races, leading to a simpler model. We do illustrate throughout the paper how our model can apply in those settings.

Baye and Hoppe [2003] analyze relationships among rent-seeking, innovation, and patent-race games. They establish the strategic equivalence of these types of games when the Tullock [1980] logit form is employed as the probability of success given agents' efforts. Skaperdas [1996] axiomatizes this probability of success function. Rosen [1991] examines R&D contests and shows that in equilibrium, a large firm invests more than a smaller firm but, by choosing safer R&D projects, makes fewer major innovations. In our Example 2, we show that unlike Baye and Hoppe [2003], the probability of success in equilibrium does not have the logit form. In addition, in contrast to Rosen [1991], we find that while the firm that invests more does choose a safer R&D project, it still has a higher chance of making a major innovation than the firm that invests less.

Gilbert and Shapiro [1990], van Dijk [1996], Denicol [1996], and Denicol [2000] study the optimal selection of patent breadth (among other properties) given that firms compete in quality improvement. Gallini [1992] analyzes the optimal selection of patent breadth and length. In her setting, she shows that broad, short-lived patents are optimal. In Section 4, we study the case where agents must obtain at least a threshold outcome in order to win. This threshold can be interpreted as an existing patent's breadth, or the quality of an existing product in the market (which must be exceeded for a new product to be of value). A more general discussion of incentive properties of mechanisms for intellectual property is given by Gallini and Scotchmer [2001].

There are certainly aspects of R&D competition and patent races that our model does not capture (and many of these aspects are explored in other literature). A benefit of our model is that it is simple and can be embedded in multiple frameworks, as we show throughout the paper. Incorporating aspects that are not common to all of these applications into the model is likely to make it less generally applicable. For example, we do not study repeated interaction, because how this should be done presumably depends on the specific application. (In R&D, phenomena such as increasing dominance and persistence of monopoly are of interest [Cabral, 2002]; whereas in patent races, the value of an innovation over time is affected by patent regulation, raising the question of how to regulate to encourage innovation [Denicol , 1996].) Specializing the model to particular applications is an important direction for future research. Additionally, it may be possible to add features to our model that do not significantly restrict its applicability. We will discuss future research directions in more detail at the end of this paper.

In a working paper, Dulleck *et al.* [2006] (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 his budget, and has access to a fair insurance market. (An all-pay auction is an auction in which each agent must pay his bid, even if he did not win. For an overview on all-pay auctions, see Baye *et al.* [1996]. “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 [1996], 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 (which is reminiscent of the Colonel Blotto game as in Roberson [2006])—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 (simultaneously) selects any fair lottery over nonnegative real numbers.[3] We describe a lottery by its cumulative distribution function (CDF) $F(x) : \mathbb{R}^{\geq 0} \to [0,1]$. 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 Anscombe and Aumann [1963] 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.

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

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:[4] $U_i(F_i, F_{-i}) = \int_0^\infty \prod_{j \neq i} F_j(x) dF_i(x)$. We will be interested in the Nash equilibria $\vec{F}^* = (F_1^*, F_2^*, ..., F_n^*)$ 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 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 $\int_0^{2b_2} x dF_1(x) = b_1$.

[4] Technically, 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 $F_2(x) > \lim_{\epsilon \to 0} F_2(x - \epsilon)$, then the probability for 1 of winning given that he obtains outcome x is not $F_2(x)$, but rather $\lim_{\epsilon \to 0} F_2(x - \epsilon) + (F_2(x) - \lim_{\epsilon \to 0} F_2(x - \epsilon))/2$. This is only relevant if agent 1 also has a mass point at x.

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 $[0, 2b]$. That is, agent 1's lottery has a probability mass at 0. (p is a *mass point* of a cumulative distribution function F if $\lim_{\epsilon\to 0} F(p+\epsilon) - F(p-\epsilon) > 0$.) 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.

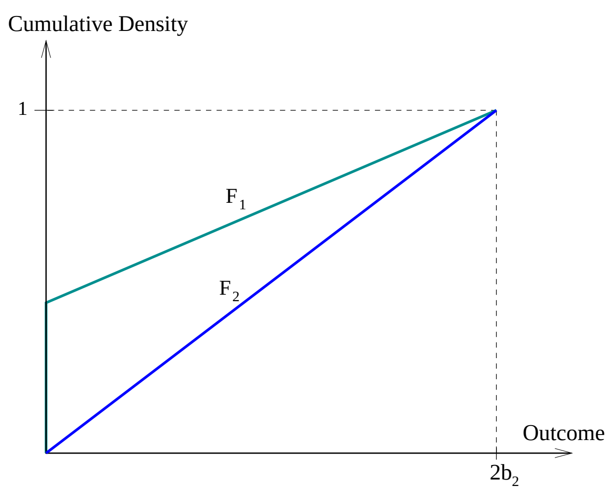

Figure 1: Equilibrium strategies in Example 2

Since agent 1 has a chance of winning only if he won his initial gamble, after which he has the same budget as agent 2, his 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 $b_1 \le b_2$). 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. Dulleck *et al.* [2006] 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 $\int_0^{3b} F^2(x)dF_1(x) = \frac{1}{3b}\int_0^{3b} xdF_1(x)$ subject to $\int_0^{3b} xdF_1(x) = b$. 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 his 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 support[5] of F_i, then the following inequality holds:*

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

[5] In our use of the word "support", the support is a closed set, that is, we include all the limit points in the support.

The proofs of Lemmas 1 and 2 are in the appendix. 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).

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.*

Proof: Consider a best response F_i. Because agent i does not win with probability 1, there must exist some x in the support of F_i, some $\epsilon > 0$, and some δ, such that $F_{-i}(x + \delta) - F_{-i}(x) > \epsilon$ (and F_{-i} does not have a mass point at $x + \delta$). Now suppose that F_i

has no upper bound. Then, there must exist some y in the support of F_i such that $F_{-i}(y-\delta) > 1-\epsilon/4$. For sufficiently small m, there exists some $m' > m/2$ such that we can change F_i an expectation-preserving way, as follows:

- Move mass m from around y to $y-\delta$,
- Move mass m' from around x to $x+\delta$.

For sufficiently small m, this results in an increase in the probability of winning for i of at least $m(F_{-i}(y-\delta)-F_{-i}(y))+(m/2)(F_{-i}(x+\delta)-F_{-i}(x)) > -m(\epsilon/4)+(m/2)\epsilon = m\epsilon/4 > 0$, which contradicts the original F_i being a best response. ■

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.*

Proof: Suppose on the contrary that F places some positive mass m on outcome k. Then there is a positive probability of a tie at k. Consider agent i. Agent i's budget constraint implies that i has some mass on outcomes equal to or larger than b. Let $\epsilon > 0$ satisfy

$$m\epsilon < \int_b^\infty x\,dF(x)$$

Agent i can shift the mass at k up to outcome $k+\epsilon$. This will create an upward pressure of $m\epsilon$ on i's budget constraint. In order to mitigate this pressure, mass can be shifted from outcomes equal to or larger than b down to 0. As ϵ approaches 0, the mass that needs to be shifted down becomes infinitesimally small, so that the cost of shifting down the mass becomes infinitesimally small as well. However, due to a positive probability of a tie at k, agent i's gain from redistributing as prescribed is bounded away from 0. Hence, agent i possesses a profitable deviation, which is contrary to the equilibrium assumption. ■

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 $l = 0$.*

Proof: Consider agent i. Since F constitutes a symmetric equilibrium, Corollary 1 tells us that both F and F_{-i} are continuous. Suppose on the contrary that $l > 0$. Continuity of F_{-i} implies that for any $\epsilon > 0$, there exists a $\delta > 0$ such that for $|x - l| < \delta$, $|F_{-i}(x) - F_{-i}(l)| = F_{-i}(x) < \epsilon$ (where we make use of the fact that $F_{-i}(l) = 0$). Let h denote the least upper bound of the support, which exists by Lemma 4. Note that $h > l$ and $F_{-i}(h) = 1$ hold by continuity. We set $\epsilon = l/h$. Consider an upper neighborhood of l, $[l, l + \psi]$, where $0 < \psi < \delta$. Denote the probability mass spread over $[l, l + \psi]$ by ϵ_l, so that

$$\int_{l}^{l+\psi} dF(x) = \epsilon_l \tag{1}$$

Note that $\epsilon_l > 0$ by continuity of F and the fact that l is in the support. Also, we have that

$$\int_{l}^{l+\psi} F_{-i}(x) dF(x) < F_{-i}(l + \psi)\epsilon_l \tag{2}$$

and

$$F_{-i}(l + \psi) < \epsilon \tag{3}$$

where (3) holds since $\psi < \delta$. Define ϵ_h by

$$\int_{l}^{l+\psi} x dF(x) dx = \epsilon_h h \tag{4}$$

In words, ϵ_h is the probability mass that would need to be placed on outcome h when mass is removed from $[l, l+\psi]$, so as not to change the expected outcome of the lottery. Note that

$$\epsilon_l(l+\psi) > \epsilon_h h > \epsilon_l l \tag{5}$$

holds by definition of ϵ_h. Thus, $\epsilon_h > \epsilon_l(l/h)$. Lastly, define ϵ_0 by

$$\epsilon_0 = \epsilon_l - \epsilon_h \tag{6}$$

We plan on reallocating mass from $[l, l+\psi]$ to outcomes 0 and h. Specifically, we will shift mass ϵ_0 to outcome 0 and ϵ_h to outcome h. Conditions (4) and (6) ensure that the magnitude of the mass and the budget constraint will be preserved. By reallocating this mass, agent i's expected utility changes by

$$\begin{aligned}
&\epsilon_h F_{-i}(h) - \int_l^{l+\psi} F_{-i}(x)dF(x) \\
&> \epsilon_h - \epsilon_l F_{-i}(l+\psi) \qquad (7)\\
&> \epsilon_h - \epsilon_l \epsilon \\
&= \epsilon_h - \epsilon_l \frac{l}{h} \\
&> 0
\end{aligned}$$

The first two inequalities follow from (1)-(3). The equality follows from the definition of ϵ, and the last inequality follows from (5). Hence, agent i possesses a profitable deviation, which is in contradiction to the equilibrium assumption. Thus, $l = 0$. ■

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 he 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:*

$$F(x) = (nb)^{-\frac{1}{n-1}} x^{\frac{1}{n-1}} \tag{8}$$

over support $[0, nb]$.

Proof: First, note that lottery F is a viable strategy:

$$\int_0^{nb} (nb)^{-\frac{1}{n-1}} x^{\frac{1}{n-1}} dx = b$$

Given that all agents other than i employ strategy F, agent i will not allocate mass to outcomes larger than nb. Thus, agent i's problem is to select F_i to maximize

$$\int_0^{nb} \prod_{j \neq i} F_j(x) dF_i(x) = \frac{1}{nb} \int_0^{nb} x dF_i(x) \tag{9}$$

subject to

$$\int_0^{nb} x dF_i(x) = b \tag{10}$$

Note that because of the constraint, the integral in (9) must equal b for any F_i that satisfies (10). Hence, playing F is a best-response to F_{-i} for agent i, and so all agents playing F constitutes a symmetric equilibrium.

To show that this is the only symmetric equilibrium, we proceed as follows. Consider lottery G. Using Lemma 4, let h be the least upper bound of G (since we assume supports to be closed, h is in the support), and suppose that G constitutes a symmetric equilibrium. Note that by definition, $G(h) = 1$. By Lemmas 5 and 6, 0 is in the support of G, and $G(0) = 0$. Consider agent i. By Lemma 3 and Corollary 1, we know that for x_1, x_2, and x_3 in the support of G, such that $x_1 \leq x_2 \leq x_3$, we have

$$\begin{aligned}(x_2 - x_1)G_{-i}(x_3) + (x_3 - x_2)G_{-i}(x_1) \\ = (x_3 - x_1)G_{-i}(x_2)\end{aligned} \tag{11}$$

Let $x_3 = h$ and $x_1 = 0$. Substituting in (11), we obtain that for any x_2 in the support of G

$$G_{-i}(x_2) = \frac{x_2}{h}$$

By symmetry, we also have

$$G(x) = (\frac{x_2}{h})^{\frac{1}{n-1}} \tag{12}$$

To show that G has no gaps, suppose the contrary. Then, there exist l' and h', $0 < l' < h' < h$, such that l' and h' are in the support of G but the interval (l', h') is not. Since (l', h') is not in the support, and by continuity of G, $G(l') = G(h')$. However, since $l' < h'$, this contradicts (12). Hence, G has no gaps. Since G has no gaps and G must satisfy the budget constraint, we have that

$$\int_0^{h} x dG(x) = b \tag{13}$$

From equalities (12) and (13) we can derive $h = nb$. Substituting for h in (12), we obtain that $F = G$. ∎

In the appendix, we provide an alternative method to derive Theorem 1 using results from the common-value all-pay auction literature and some of the lemmas here. If all agents use the lottery described in (8), 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.

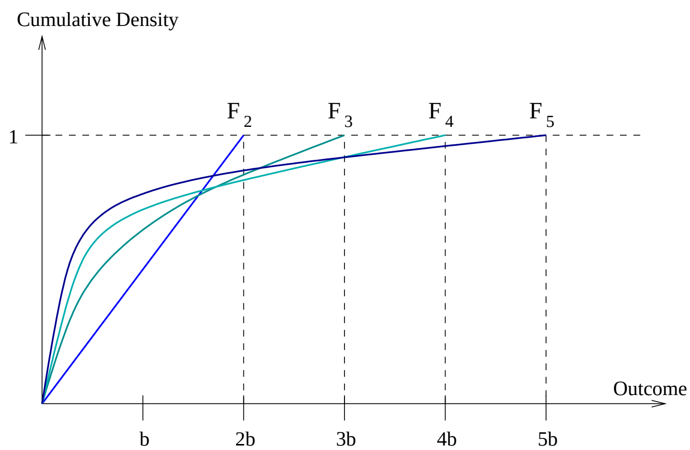

Figure 2: Cumulative distribution of symmetric equilibrium strategy for different values of n, given equal budgets $b = 5$.

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:

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

This expectation is quite high, in the following sense. Suppose that we did not impose any strategic constraints on F_i. Then, $E[x_{max}] \le E[\sum_i x_i] = \sum_i E[x_i] = nb$. 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 $\epsilon \to 0$ 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 $i \in \{1,\ldots,n\}$. The next lemma shows that if (A1) holds, then all agents have 0 in their support.

Lemma 7. *Suppose that $\vec{F}^* = (F_1^*, F_2^*, \ldots, F_n^*)$ 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 $i \in \{1,2,\ldots,n\}$.*

Proof: First, (A1) implies that all supports must have the same greatest lower bound (henceforth *GLB*). To see this, note that if agent i has a higher GLB than j, then agent j is guaranteed to lose the game given an outcome in the interval between his GLB and agent i's GLB. Hence, j would prefer to shift some of this mass down to 0, and the remainder to outcomes that give him a chance of winning, resulting in a strategy with a gap. Thus, all agents' supports must have the same GLB. If this GLB were greater than 0, then any agent would prefer to shift mass from a neighborhood of that GLB down to 0 in order to reallocate other mass to higher outcomes (the formal argument here is similar to that made in Lemma 6). ■

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.*

Proof: Suppose, for the sake of contradiction, that an equilibrium that is not symmetric exists. In this equilibrium, consider any two agents with different strategies and denote their chosen lotteries by F and G. Denote the distribution of the maximum outcome of all *other* agents by H. Let h_F and h_G denote the least upper bounds of the supports of

F and G, respectively. (By Lemma 4, equilibrium strategies must always have an upper bound.) Assume without loss of generality that $h_F \le h_G$. Because of (A1) and Lemma 7, we know that every agent i's support has the form $[0, h_i]$. Also, F_i is continuous because F_i is a nondecreasing function and (A2) rules out mass points. Since F_{-i} is the product of continuous functions, F_{-i} is continuous as well. Finally, (A2) implies that $F_{-i}(0) = 0$. Hence, we can apply Lemma 3 to obtain

$$(x-0)G(h_F)H(h_F) + (h_F - x)G(0)H(0) = (h_F - 0)G(x)H(x)$$

Using the fact that $G(0)H(0) = 0$, we obtain:

$$G(x)H(x) = c_1 x$$

for some positive c_1. Similarly,

$$F(x)H(x) = c_2 x$$

for some positive c_2. Combining these conditions, we obtain that for x in $[0, h_F]$,

$$F(x) = \frac{c_2}{c_1} G(x) \tag{14}$$

Now suppose that $h_F < h_G$. Because supports have no gaps by (A1), it must be that $G(h_F) < 1$. Hence, in order for $F(h_F) = 1$ to hold, we need

$$\frac{c_2}{c_1} > 1$$

It follows that G first-order stochastically dominates F on $[0, h_G]$. This entails that G has a higher expectation, which contradicts our premise that all agents have equal budgets. Therefore, $h_G = h_F$. It follows that all agents' lotteries must have identical supports $[0, h]$. However, by (14),

$$F(h) = \frac{c_2}{c_1} G(h)$$

Since $F(h) = G(h) = 1$, it must be that $c_1 = c_2$. This means that F equals G, contrary to the initial assumption that they were unequal. It follows that any equilibrium must be symmetric. But Theorem 1 tells us that there is only a single symmetric equilibrium. ■

4 Adding a 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 a specification that if a contestant does not outperform a risk-free asset, then the contestant cannot win. Under the R&D interpretation, r represents the existing product quality in the market (a "reserve" quality), a quality that research departments must improve upon to generate any business value. In an innovation tournament or in a patent race, r represents the breadth of the current patent on some product. To be able to register a new patent, innovators must reach a level of innovation that surpasses the breadth of the current patent. (For technical simplicity, we assume that an innovation of quality exactly r can be registered.)

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 $l = r$.*

Proof: Let l denote the greatest lower bound of the support of F excluding 0. Then $l \geq r$. Consider agent i. Lemma 5 and the fact that F constitutes a symmetric equilibrium tell us that both F and F_{-i} are continuous over outcomes greater or equal to r. Suppose on the contrary that $l > r$. Continuity of F_{-i} over outcomes greater or equal to r implies that for any $\epsilon > 0$, there exists a $\delta > 0$ such that for $|x-l| < \delta$, $|F_{-i}(x) - F_{-i}(l)| < \epsilon$. Furthermore, $F_{-i}(r) = F_{-i}(l) = m^{n-1}$, where m is the mass F places at 0. Let h denote the least upper bound of the support, which exists by Lemma 4. Note that $h > l$ and $F_{-i}(h) = 1$ hold by continuity. We set

$$\epsilon = \frac{l-r}{h-r}(1-m^{n-1})$$

and

$$\delta = min\{\delta(\epsilon), h-l\}$$

Consider an upper neighborhood of l, $[l, l+\psi]$, where $0 < \psi < \delta$. Denote the probability mass spread over $[l, l+\psi]$ by ϵ_l, so that

$$\epsilon_l = \int_l^{l+\psi} dF(x) \tag{15}$$

$\epsilon_l > 0$ by continuity of F in this region and the fact that l is in the support. Also, we have that

$$\int_l^{l+\psi} F_{-i}(x)dF(x) < F_{-i}(l+\psi)\epsilon_l \tag{16}$$

and

$$F_{-i}(l+\psi) < \epsilon + m^{n-1} \tag{17}$$

where (17) holds since $\psi < \delta$. Since $r < l < l+\psi < h$, there exist strictly positive ϵ_r and ϵ_h such that

$$\epsilon_r r + \epsilon_h h = \int_l^{l+\psi} x dF(x)dx$$

and

$$\epsilon_r + \epsilon_h = \int_l^{l+\psi} dF(x)dx$$

In words, there exists a mean- and mass-preserving spread from $[l, l+\psi]$ to outcomes r and h. By definition,

$$\epsilon_l(l+\psi) > \epsilon_h h + \epsilon_r r > \epsilon_l l$$

Substituting for $\epsilon_r = \epsilon_l - \epsilon_h$, we obtain

$$\epsilon_h > \frac{l-r}{h-r}\epsilon_l \tag{18}$$

By reallocating mass from $[l, l+\psi]$ to outcomes r and h as described above, agent i's expected utility changes by

$$\epsilon_h F_{-i}(h) + \epsilon_r F_{-i}(r) - \int_l^{l+\psi} F_{-i}(x)dF(x)$$

$$
\begin{aligned}
&> \epsilon_h + \epsilon_r m^{n-1} - \epsilon_l F_{-i}(l+\psi) \\
&> \epsilon_h + \epsilon_r m^{n-1} - \epsilon_l(\epsilon + m^{n-1}) \\
&= \epsilon_h(1-m^{n-1}) - \epsilon_l \epsilon \\
&= (1-m^{n-1})(\epsilon_h - \epsilon_l \frac{l-r}{h-r}) \\
&> 0
\end{aligned}
$$

The first two inequalities follow from continuity of F_{-i} and (16). The next two equalities follow from the definitions of ϵ_r, ϵ_h, and ϵ. The last inequality follows directly from (18). Hence, agent i possesses a profitable deviation, which is in contradiction to the equilibrium assumption. Thus, $l = r$. ■

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, $a = 0$.)*

Proof: Suppose on the contrary that $a > 0$. Let $x_1 = r$, $x_2 = x \in (r,h)$, and $x_3 = h$. Applying the result of of Lemma 3 and substituting for F gives

$$x(1-m) + hm = ha + hcx + r - ra - rcx \tag{19}$$

Now set $x_1 = 0$, keeping $x_2 = x \in (r,h)$, $x_3 = h$. Applying the result of the extended Lemma 3 and substituting for F gives

$$x(1-m) + hm = ha + hcx \tag{20}$$

Combining equations (19) and (20), we obtain $c = 0$, implying that the symmetric equilibrium strategy places a mass of 1 at 0. This is in contradiction to agents having positive budgets. ∎

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

$$F(x) = \begin{cases} \frac{\sqrt{b^2+r^2}-b}{r} & \text{if} \quad 0 \le x < r \\ \frac{\sqrt{b^2+r^2}-b}{r^2} x & \text{if} \quad r \le x \le \frac{r^2}{\sqrt{b^2+r^2}-b} \\ 1 & \text{if} \quad x > \frac{r^2}{\sqrt{b^2+r^2}-b} \end{cases} \tag{21}$$

It remains to verify that (21) 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.

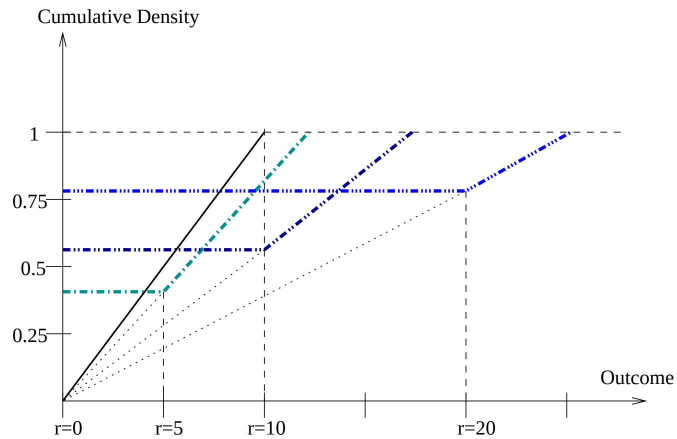

Figure 3: Cumulative distribution of symmetric equilibrium strategies for different values of r, given equal budgets $b = 5$.

We can observe the following facts about the equilibrium strategies from (21) 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*

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

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

Proof: As in the two-agent game, the symmetric equilibrium strategy F will have support in $\{0\} \cup [r,h]$, where $h > r$ is some least upper-bound, which exists by Lemma 6. The support is contained in this set because outcomes in $(0,r)$ can never lead to winning, and an agent is better off redistributing mass over this interval to 0 and outcomes greater than r. Now, for a given i, Lemmas 3, 8, and 5 imply that that over $[r,h]$, $F_{-i} = F^{n-1}$ takes the form $F^{n-1}(x) = a + cx$, where a and c are positive constants. Let $m \geq 0$ denote the mass F places at 0. Then by Lemma 5, $F^{n-1}(0) = F_{-i}(r) = m^{n-1}$. The following claim establishes that $F^{n-1}(x)$ must lie on a line originating from the origin.

Claim 2. *In the two-agent equal-budget game with minimum necessary outcome, the symmetric equilibrium strategy F, with $F^{n-1}(x) = a + cx$ over $[r,h]$, has intercept 0. Hence, $a = 0$.*

The proof follows a similar argument to the one made in the proof of Claim 1. It follows that $F(x) = (cx)^{\frac{1}{n-1}}$ over $[r,h]$. From $F(h) = 1$ we obtain $h = c^{-1}$. Also, from $F^{n-1}(r) = m^{n-1}$ we obtain

$$m = (cr)^{\frac{1}{n-1}} \tag{22}$$

Finally, the budget constraint requires

$$\int_r^{c^{-1}} x dF(x) = b \tag{23}$$

Substituting for F in (23) we obtain

$$\frac{1}{n}(c^{-1} - c^{\frac{1}{n-1}} r^{\frac{n}{n-1}}) = b \tag{24}$$

Equality (24) implicitly and uniquely defines[6] $c(b,r)$, whereas $m(b,r) = (c(b,r)r)^{\frac{1}{n-1}}$ from (22). The candidate symmetric equilibrium strategy is for each agent to select the lottery specified by

$$F(x) = \begin{cases} m(b,r) & \text{if} \quad x < r \\ \\ (c(b,r)x)^{\frac{1}{n-1}} & \text{if} \quad x \in [r, (c(b,r))^{-1}] \\ \\ 1 & \text{if} \quad x > (c(b,r))^{-1} \end{cases} \tag{25}$$

By construction, the specification in (25) provides the unique candidate symmetric equilibrium strategy. It remains to verify that (25) indeed constitutes an equilibrium strategy. To check this, suppose all agents other than i employ strategy F. Given this, agent i would not find it optimal to place mass on outcomes higher than $c(b,r)^{-1}$. Then, agent i's problem is to select lottery F_i to maximize

$$\int_r^{c(b,r)^{-1}} F_{-i}(x) dF_i(x) = c(b,r) \int_r^{c(b,r)^{-1}} x dF_i(x)$$

subject to

$$\int_r^{c(b,r)^{-1}} x dF_i(x) = b$$

Playing F is a best-response to F_{-i} for agent i and thus F constitutes a unique symmetric equilibrium strategy. ∎

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 $n \to \infty$.

[6] $c(b,r)$ exists and is unique because the left-hand side of (24) is continuously decreasing in c, positive when c is small, and negative when c is large.

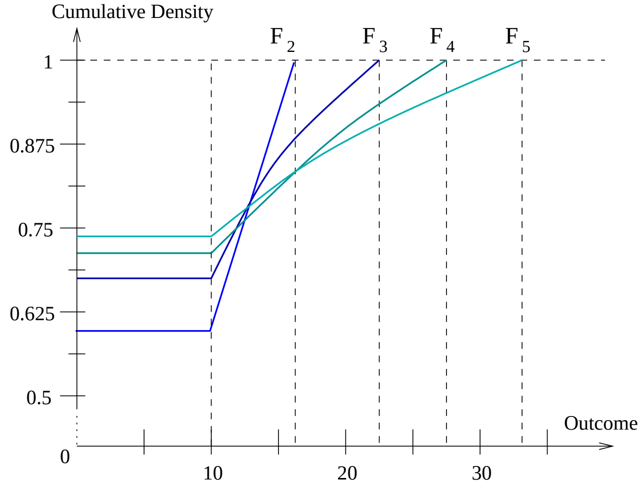

Figure 4: Cumulative distribution of symmetric equilibrium strategies for different values of n, given equal budgets $b = 5$ and $r = 10$.

5 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 at a cost. After the budgets have been chosen, the game proceeds as before. This variant is especially natural in many real-world applications, where agents must make some initial investment. For instance, a game can model an R&D competition between two risk-neutral firms: to improve their product, each firm can choose to pursue various technologies, each of which bears a different cost. The chosen technology stochastically determines the final product quality, and the firm with the highest realized product quality wins the entire market. Specifically, the game proceeds as follows. 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 he does not win, and $D - b_i$ if he does win, where D is a constant (e.g. , the benefit from winning the market). Agents maximize their expected utilities. We only consider the 2-agent case, and we 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.* [2006], 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

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

When $b_1 \le b_2$, agent 1's expected utility is linear in b_1. Hence, he will choose to set $b_1 \ge b_2$ 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, he will choose to set $b_1 = 0$ whenever $D < 2b_2$, because in this case, any other budget will give him 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

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

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

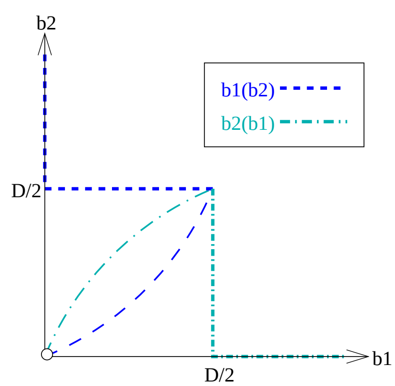

Figure 5: Best-response curves in budget selection stage

the 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 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 $j \in \{1, ..., n\}$, 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

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

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

$$\int_0^\infty \Big[\int_0^\infty \dots \int_0^\infty \prod_{j \neq i} G_{b_j}^j(x) dW_1(b_1) \dots \\ \dots dW_{i-1}(b_{i-1}) dW_{i+1}(b_{i+1}) \dots dW_n(b_n) \Big] dF(x) \tag{26}$$

Here, the bracketed expression in (26) 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 $\vec{G}^* = (G^{*1},...,G^{*n})$, for which for every $i \in \{1,...,n\}$ the bracketed term in (26) 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

$$\int_0^{nk} G_b(x)dW(b) = (nk)^{-\frac{1}{n-1}} x^{\frac{1}{n-1}} \tag{27}$$

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 (27). 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* $\int_b^\infty x dW(x) \leq$

[7] If 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.

$\int_{J^{-1}(W(b))}^{\infty} x dJ(x)$.

Proof: Consider any b in the support of W and the probability $W(b)$ of a budget at or below it. The conditional expectation of this probability mass (*i.e.* the conditional expectation of W given that the resulting budget is at or below b) is

$$(W(b))^{-1} \int_0^b x dW(x)$$

Given a G that transforms W into J, all of the probability mass that W places on budgets at or below b must correspond (*i.e.* get mapped by G) to probability mass in J. Moreover, by the fairness constraint, that mass in J must have the same conditional expectation. But a subset of the mass of J with a total probability of $W(b)$ must have a conditional expectation of at least

$$(W(b))^{-1} \int_0^{J^{-1}(W(b))} x dJ(x)$$

(because the mass in J with the lowest conditional expectation is the mass that is placed on the smallest outcomes in the support). It follows that

$$\int_0^b x dW(x) \geq \int_0^{J^{-1}(W(b))} x dJ(x)$$

Similarly, consider the probability mass that W places on outcomes greater than b (a total mass of $1 - W(b)$). The conditional expectation of this mass is equal to

$$(1 - W(b))^{-1} \int_b^{\infty} x dW(x)$$

(Note that any mass at b should not be included in this integral.) Again, the mass must correspond to mass in J, with the same conditional expectation. But a subset of the mass of J with a total probability of $1 - W(b)$ must have a conditional expectation of at most

$$(1 - W(b))^{-1} \int_{J^{-1}(W(b))}^{\infty} x dJ(x)$$

(because the mass in J with the highest conditional expectation is the mass that is placed on the largest outcomes in the support). It follows that

$$\int_b^{\infty} x dW(x) \leq \int_{J^{-1}(W(b))}^{\infty} x dJ(x)$$

■

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 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).*

Proof: For any budget b in the support of W, define $p(b)$ by

$$\frac{k}{2}p(b) + \frac{3k}{2}(1 - p(b)) = b$$

So that

$$p(b) = \frac{3k - 2b}{2k}$$

Note that $p(b) \in [0, 1]$ because $b \in [k/2, 3k/2]$. Now, consider the following compound (fair) lottery F_b:

1. Choose the lottery that with probability $p(b)$ generates outcome $k/2$, and with probability $1 - p(b)$ generates outcome $3k/2$.
2. If outcome $k/2$ was generated, then subsequently choose the lottery $U[0, k]$. If outcome $3k/2$ was generated, then subsequently choose the lottery $U[k, 2k]$.

Suppose agent i plays F_b given budget b. Since k is the expectation of W and strategy F_b involves only fair lotteries, agent i must play $U[0, k]$ with probability $1/2$ and $U[k, 2k]$ with probability $1/2$ (so that the overall expected budget outcome equals k). Therefore, the *ex ante* distribution over agent i's outcome is $U[0, 2k]$. ■

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

We studied the following game: each agent i chooses a lottery over nonnegative numbers whose expectation is equal to his 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 a social planner could obtain if the goal were 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. 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. In the setting where there is a minimum necessary outcome, it is also possible to consider the case where, if no agent reaches the minimum outcome, a fixed agent wins. (This agent would represent the incumbent who owns the product that is the current market leader, or who holds the current patent.) We can also consider different utility functions: for example, the agent may 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. There are many more open-ended modeling questions for future research. Specifically, it would be desirable to model other important aspects of applications such as R&D and patent races. Examples include increasing dominance, barriers to entry, optimal patent regulation, and mergers of R&D departments or joint research. (In the introduction, we discussed related research that considers these phenomena in the context of other models.)

References

Axel Anderson and Luis M. B. Cabral. Go for broke or play it safe? Dynamic competition with choice of variance. *RAND Journal of Economics*, 2007. Forthcoming.

Frank J. Anscombe and Robert J. Aumann. A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205, 1963.

Kyle Bagwell and Robert W. Staiger. Risky R&D in oligopolistic product markets. In *Discussion Paper No. 872, CMSEMS, Northwestern University*, 1990.

Michael R. Baye and Heidrun C. Hoppe. The strategic equivalence of rent-seeking, innovation, and patent-race games. *Games and Economic Behavior*, 44:217–226, 2003.

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

Sudipto Bhattacharya and Dilip Mookherjee. Portfolio choice in research and development. *RAND Journal of Economics*, 17(4):594–605, 1986.

Luis M. B. Cabral. Bias in market R&D portfolios. *International Journal of Industrial Organization*, 12:533–547, 1994.

Luis M. B. Cabral. Increasing dominance with no efficiency effect. *Journal of Economic Theory*, 102:471–479, 2002.

Luis M. B. Cabral. R&D competition when firms choose variance. *Journal of Economics and Management Strategy*, 12:139–150, 2003.

Vincenzo Denicol . Patent races and optimal patent breadth and length. *The Journal of Industrial Economics*, 44(3):249–265, 1996.

Vincenzo Denicol . Two-stage patent races and patent policy. *RAND Journal of Economics*, 31(3):450–487, 2000.

Uwe Dulleck, Paul Frijters, and Konrad Podczeck. All-pay auctions with budget constraints and fair insurance. Working paper 0613, Department of Economics, Johannes Kepler University of Linz, Altenberger Strasse 69, A-4040 Linz - Aufhof, Austria, July 2006.

Nancy T. Gallini and Suzanne Scotchmer. Intellectual property: When is it the best incentive mechanism? *Innovation Policy and the Economy*, 2:51–78, 2001.

Nancy T. Gallini. Patent policy and costly imitation. *RAND Journal of Economics*, 23(1):52–63, 1992.

Richard Gilbert and Carl Shapiro. Optimal patent length and breadth. *RAND Journal of Economics*, 21(1):106–112, 1990.

Kenneth L. Judd. Closed-loop equilibrium in a multi-stage innovation race. *Economic Theory*, 21:673–695, 2003.

Tor J. Klette and David de Meza. Is the market biased against risky R&D? *RAND Journal of Economics*, 17:133–139, 1986.

Jean-Jacques Laffont and Jacques Robert. Optimal auction with financially constrained buyers. *Economic Letters*, 52:181–186, 1996.

Brian Roberson. The Colonel Blotto game. *Economic Theory*, 29(1):1–24, 2006.

Richard J. Rosen. Research and development with asymmetric firm sizes. *RAND Journal of Economics*, 22:411–429, 1991.

Stergios Skaperdas. Contest success functions. *Economic Theory*, 7:283–290, 1996.

Gordon Tullock. Efficient rent-seeking. In James M. Buchanan, Robert D. Tollison, and Gordon Tullock, editors, *Toward a Theory of the Rent-Seeking Society*, pages 97–112. College Station: Texas A&M University Press, 1980.

Theon van Dijk. Patent height and competition in product improvements. *The Journal of Industrial Economics*, 44(2):151–167, 1996.

John Vickers. *Patent Races and Market Structure*. PhD thesis, Nuffield College, Oxford University, 1985.

APPENDIX

PROOF OF LEMMA 1:

If $x_1 = x_2$ or $x_2 = x_3$, the lemma is trivial, so suppose without loss of generality that $x_1 < x_2 < x_3$. The proof proceeds by contradiction. Suppose on the contrary that

$$(x_2 - x_1)F_{-i}(x_3) + (x_3 - x_2)F_{-i}(x_1) > (x_3 - x_1)F_{-i}(x_2) \tag{A-1}$$

Then

$$\frac{(x_2 - x_1)}{(x_3 - x_1)}F_{-i}(x_3) + \frac{(x_3 - x_2)}{(x_3 - x_1)}F_{-i}(x_1) > F_{-i}(x_2) \tag{A-2}$$

For any $\epsilon_2 > 0$, we have that

$$\epsilon_2\frac{(x_2 - x_1)}{(x_3 - x_1)}F_{-i}(x_3) + \epsilon_2\frac{(x_3 - x_2)}{(x_3 - x_1)}F_{-i}(x_1) > \epsilon_2 F_{-i}(x_2) \tag{A-3}$$

Define ϵ_1 by

$$\epsilon_1 = \epsilon_2\frac{(x_3 - x_2)}{(x_3 - x_1)}$$

Similarly, define ϵ_3 by

$$\epsilon_3 = \epsilon_2\frac{(x_2 - x_1)}{(x_3 - x_1)}$$

By definition, $\epsilon_2 = \epsilon_1 + \epsilon_3$ and $\epsilon_1 x_1 + \epsilon_3 x_3 = \epsilon_2 x_2$. Inequality (A-3) reduces to

$$\epsilon_1 F_{-i}(x_1) + \epsilon_3 F_{-i}(x_3) > \epsilon_2 F_{-i}(x_2) \tag{A-4}$$

There are now two possible scenarios:

(i) If F_i has positive mass at outcome x_2, that is, there is a positive probability that i will get exactly x_2, then the contradiction follows immediately: setting ϵ_2 to equal this mass, inequality (A-4) implies that agent i would be better off redistributing ϵ_2 to outcomes x_1 and x_3. The definitions of ϵ_1 and ϵ_3 ensure that i would be shifting mass in a way that satisfies his budget constraint.

(ii) If F_i has no mass at outcome x_2, we can still show that agent i has a profitable deviation by gathering up mass in a *neighborhood*[8] of x_2 for which inequality (A-1) holds, and redistributing this mass to outcomes x_1 and x_3 in a mean-preserving way. We now show this formally. Define θ by

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

[8]By neighborhood of x_2, we refer to a closed interval that contains x_2 in its interior.

$\theta > 0$ by inequality (A-2). Continuity of F_{-i} implies that for any $\epsilon > 0$, there exists a $\nu > 0$, such that for $|x - x_2| < \nu$, $|F_{-i}(x) - F_{-i}(x_2)| < \epsilon$. Let

$$\epsilon = \frac{\theta}{2}$$
$$\delta = \min\{\frac{1}{2}\nu, \epsilon(x_3 - x_1)\} \tag{A-5}$$

and

$$\psi = F_i(x_2 + \delta) - F_i(x_2 - \delta) \tag{A-6}$$

Since x_2 is in the support of F_i and $\delta > 0$, F_i has positive mass over $[x_2 - \delta, x_2 + \delta]$. Thus, $\psi > 0$. Define ϕ by

$$\psi(\frac{(x_3 - x_2)}{(x_3 - x_1)} + \phi)x_1 + \psi(\frac{(x_2 - x_1)}{(x_3 - x_1)} - \phi)x_3 = \int_{x_2-\delta}^{x_2+\delta} x dF_i(x) \tag{A-7}$$

ϕ is the adjustment required in the coefficients of x_1 and x_3 (which correspond to ϵ_1 and ϵ_3) in order to ensure that the budget constraint is preserved after redistributing mass from $[x_2 - \delta, x_2 + \delta]$ to x_1 and x_3 (ϕ could be negative). By definition,

$$(\frac{(x_3 - x_2)}{(x_3 - x_1)} + \phi)x_1 + (\frac{(x_2 - x_1)}{(x_3 - x_1)} - \phi)x_3 \geq x_2 - \delta \tag{A-8}$$

Furthermore,

$$\frac{x_3 - x_2}{x_3 - x_1}x_1 + \frac{x_2 - x_1}{x_3 - x_1}x_3 = x_2 \tag{A-9}$$

Combining (A-5)-(A-9) and using the definition of δ, we obtain

$$\phi \leq \frac{\delta}{x_3 - x_1} \leq \epsilon \tag{A-10}$$

Utilizing the above construction, we have that

$$\begin{aligned}
\psi(\frac{(x_3 - x_2)}{(x_3 - x_1)} + \phi)F_{-i}(x_1) + \psi(\frac{(x_2 - x_1)}{(x_3 - x_1)} - \phi)F_{-i}(x_3) & \\
= \psi\frac{(x_3 - x_2)}{(x_3 - x_1)}F_{-i}(x_1) + \psi\frac{(x_2 - x_1)}{(x_3 - x_1)}F_{-i}(x_3) & \\
-\psi\phi(F_{-i}(x_3) - F_{-i}(x_1)) & \\
\geq \psi F_{-i}(x_2 + \delta) & \\
> \int_{x_2-\delta}^{x_2+\delta} F_{-i}(x) dF_i(x) &
\end{aligned} \tag{A-11}$$

The first inequality follows from (A-10), the construction of δ and ϵ, and by continuity of F_{-i} at x_2. In addition, we make use of the facts that $F_{-i}(x_3) - F_{-i}(x_1) \leq 1$ and $\psi \leq 1$. The last inequality holds by the definition of ψ. Lastly, the budget constraint is preserved by (A-7). Thus, the inequalities in (A-11) imply that agent i is better off redistributing mass from $[x_2 - \delta, x_2 + \delta]$ to outcomes x_1 and x_3, which contradicts F_i being i's best-response to F_{-i}. The lemma follows.

Proof of Lemma 2:

If $x_1 = x_2$ or $x_2 = x_3$, the lemma is trivial, so suppose without loss of generality that $x_1 < x_2 < x_3$. The proof proceeds by contradiction. Suppose on the contrary that

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

Then

$$\frac{(x_2 - x_1)}{(x_3 - x_1)}F_{-i}(x_3) + \frac{(x_3 - x_2)}{(x_3 - x_1)}F_{-i}(x_1) < F_{-i}(x_2) \tag{A-12}$$

For any $\epsilon_2 > 0$, we have that

$$\epsilon_2\frac{(x_2 - x_1)}{(x_3 - x_1)}F_{-i}(x_3) + \epsilon_2\frac{(x_3 - x_2)}{(x_3 - x_1)}F_{-i}(x_1) < \epsilon_2 F_{-i}(x_2) \tag{A-13}$$

Define ϵ_1 by

$$\epsilon_1 = \epsilon_2\frac{(x_3 - x_2)}{(x_3 - x_1)} \tag{A-14}$$

Similarly, define ϵ_3 by

$$\epsilon_3 = \epsilon_2\frac{(x_2 - x_1)}{(x_3 - x_1)} \tag{A-15}$$

By definition, $\epsilon_2 = \epsilon_1 + \epsilon_3$ and $\epsilon_1 x_1 + \epsilon_3 x_3 = \epsilon_2 x_2$. Inequality (A-13) reduces to

$$\epsilon_1 F_{-i}(x_1) + \epsilon_3 F_{-i}(x_3) < \epsilon_2 F_{-i}(x_2) \tag{A-16}$$

There are now four possible scenarios:

(i) If F_i has positive mass at outcomes x_1 and x_3, then the contradiction follows immediately, as agent i would be better off shifting some mass to outcome x_2 by (A-16). The construction of ϵ_1, ϵ_2, and ϵ_3 ensures that mass can be redistributed in a way that preserves agent i's budget constraint (e.g. if F_i has mass m_1 at x_1 and m_3 at x_3, then let $\epsilon_2 = min\{\frac{x_3 - x1}{x_3 - x_2}m_1, \frac{x_3 - x1}{x_2 - x1}m_3\}$. Define ϵ_1 and ϵ_3 as in (A-14) and (A-15). This ensures that $\epsilon_1 \le m_1$ and $\epsilon_3 \le m_3$). The contradiction follows.

(ii) If F_i has no mass at both outcomes x_1 and x_3, we can still show that agent i has a profitable deviation by gathering up mass in neighborhoods of x_1 and x_3 and reallocating this mass to outcome x_2. We now show this formally. Since x_1 and x_3 are in the support of F_i, F_i has mass over neighborhoods of these outcomes. Define θ by

$$\theta = F_{-i}(x_2) - \frac{(x_2 - x_1)}{(x_3 - x_1)}F_{-i}(x_3) - \frac{(x_3 - x_2)}{(x_3 - x_1)}F_{-i}(x_1)$$

$\theta > 0$ by (A-12). Continuity of F_{-i} implies that for any $\epsilon > 0$, there exist $\delta_1 > 0$ and $\delta_3 > 0$, such that for $|x - x_1| < \delta_1$ and $|y - x_3| < \delta_3$, $|F_{-i}(x) - F_{-i}(x_1)| < \epsilon$ and $|F_{-i}(y) - F_{-i}(x_3)| < \epsilon$. Let

$$\epsilon = \frac{\theta}{3}$$

and

$$\delta = min\{\frac{1}{2}\delta_1, \frac{1}{2}\delta_3, x_3 - x_2, x_2 - x_1, \epsilon(x_3 - x_1)\}$$

Define $M(F_i, x, \epsilon)$ to be the distribution of outcomes in a neighborhood of x, derived from F, that has total mass ϵ. (Technically, $M(F_i, x, \epsilon)$ would need to be normalized by a factor of $1/\epsilon$ in order to be a CDF.) By definition, the expectation of $M(F_i, x, \epsilon)$ is continuous in ϵ. Denote this expectation by $E[M(F_i, x, \epsilon)]$ (This expectation is the upward pressure placed on the budget constraint by $M(F_i, x, \epsilon)$). Define $m_1(t)$ by

$$m_1(t) = F_i(x_1 + \frac{\delta}{2^t}) - F_i(x_1 - \frac{\delta}{2^t})$$

In words, m_1 denotes mass taken in the $\delta/2^t$ neighborhood of x_1. Also, define $\psi(t)$ by

$$\psi(t) = \int_{x_1 - \frac{\delta}{2^t}}^{x_1 + \frac{\delta}{2^t}} x dF_i(x)$$

$\psi(t)$ denotes the upward pressure on the budget constraint added by probability mass distributed over outcomes in this neighborhood. Since $x_1 < x_2 < x_3$, continuity of $E[M(F_i, x, \epsilon)]$ implies that for any $t \geq 0$, there exists some mass $m_3(t)$ such that

$$\psi(t) + E[M(F_i, x_3, m_3(t))] = (m_1(t) + m_3(t))x_2 \tag{A-17}$$

We can now take t sufficiently high, so that $M(F_i, x_3, m_3(t))$ is distributed only over outcomes in $[x_3 - \delta, x_3 + \delta]$. Denote such t by T. We know T exists because $x_3 > x_2$ and because F_i has no mass at x_1, so that for sufficiently high t, $m_1(t)$ becomes arbitrarily small. From now on, we will refer to $m_j(T)$ by m_j, $j \in \{1,3\}$, to $\psi(T)$ by ψ, and to $M(F_i, x_3, m_3(T))$ by M. By construction, we have that

$$|\psi/m_1 - x_1| < \delta \tag{A-18}$$

and

$$|E[M]/m_3 - x_3| < \delta \tag{A-19}$$

In order to perturb masses m_1 and m_3 so as to fit the setting of (A-16), define ϕ by

$$(m_1 + \phi)x_1 + (m_3 - \phi)x_3 = (m_1 + m_3)x_2 \tag{A-20}$$

Note that ϕ can be negative. By (A-17)-(A-20), we have that

$$\psi - m_1\delta + \phi x_1 + E[M] - m_3\delta - \phi x_3 \leq (m_1 + m_3)x_2$$

Simplifying and rearranging, we obtain that

$$\phi \geq -\delta \frac{m_1 + m_3}{x_3 - x_1}$$

By construction of δ (which further implies that $m_1 + m_3 \leq 1$), we have that

$$\phi \geq -\epsilon \tag{A-21}$$

We are now ready to derive a contradiction. By inequality (A-16) and the definitions of δ and ϕ, we have that

$$\begin{aligned} &(m_1 + m_3)F_{-i}(x_2) \\ &> (m_1 + \phi)F_{-i}(x_1 + \delta) + (m_3 - \phi)F_{-i}(x_3 + \delta) \\ &= m_1 F_{-i}(x_1 + \delta) + m_3 F_{-i}(x_3 + \delta) - \phi(F_{-i}(x_3) - F_{-i}(x_1)) \end{aligned} \tag{A-22}$$

Furthermore, because $F_{-i}(x_3) - F_{-i}(x_1) \leq 1$, by construction of ϵ, and by (A-21), (A-22) implies that

$$(m_1 + m_3)F_{-i}(x_2) > m_1 F_{-i}(x_1 + \delta) + m_3 F_{-i}(x_3 + \delta) \tag{A-23}$$

Lastly, by definition of m_1 and m_3, it follows that

$$\begin{aligned} &m_1 F_{-i}(x_1 + \delta) + m_3 F_{-i}(x_3 + \delta) \\ &> \int_{x_1 - \frac{\delta}{2T}}^{x_1 + \frac{\delta}{2T}} F_{-i}(x) dF_i(x) + \int_{x_3 - \delta}^{x_3 + \delta} F_{-i}(y) dM(y) \end{aligned} \tag{A-24}$$

Combining inequalities (A-23) and (A-24), we obtain

$$\begin{aligned} &(m_1 + m_3)F_{-i}(x_2) > \\ &\int_{x_1 - \frac{\delta}{2T}}^{x_1 + \frac{\delta}{2T}} F_{-i}(x) dF_i(x) + \int_{x_3 - \delta}^{x_3 + \delta} F_{-i}(y) dM(y) \end{aligned} \tag{A-25}$$

Therefore, agent i would find it profitable to redistribute mass from neighborhoods of outcomes x_1 and x_3 to outcome x_2 in a mean-preserving way, which contradicts the premise that F_i constitutes a best response to F_{-i}.

(iii and iv) In these scenarios, F_i has positive mass at either outcome x_1 or x_3 (but not at both outcomes). In this case, we apply the same argument as in scenario (ii), with the exception that we gather mass only around the outcome that has no mass. The lemma follows.

An alternative method to derive Theorem 1:

In this section, we provide an alternative proof of Theorem 1 (the symmetric equilibrium strategy and its uniqueness) using results from common-value all-pay auctions along with some of the intermediate results that we proved before Theorem 1 (for a review of the common-value all-pay auction literature, see Baye *et al.* [1996]).

Consider a common-value all-pay auction whose prize is nb (this prize is chosen so that the supports of the equilibrium strategies in the two games will coincide). We will show that F is a symmetric equilibrium strategy of our game (with budget b for each agent) if and only if it is a symmetric equilibrium strategy of this common-value all-pay auction. It is known that the common-value all-pay auction has a unique symmetric equilibrium [Baye *et al.*, 1996], so Theorem 1 follows.

First, we prove the easier direction: if F is a symmetric equilibrium strategy in the common-value all-pay auction, then it is a symmetric equilibrium strategy in our game. The unique equilibrium strategy of the common-value all-pay auction is known to have expectation b, so it is a valid strategy in our game. Moreover, if there were a beneficial deviation from this strategy in our game, then it would also constitute a beneficial deviation in the common-value all-pay auction, because the player would obtain a higher probability of winning with the same expected payment; but this is contrary to the assumption that F is an equilibrium strategy of the common-value all-pay auction.

Now, we will prove the more difficult direction: if F is a symmetric equilibrium strategy of our game, then it is a symmetric equilibrium strategy of the common-value all-pay auction with prize nb. We will show this as follows. Suppose, for the sake of contradiction, that G is a beneficial deviation (in the common-value all-pay auction setting) when everyone plays F. We will derive a strategy H such that H is also a beneficial deviation, but $E(H) = E(F) = b$. (Here, $E[F]$ refers to the expectation of a random variable distributed according to F.) Hence, H is a valid strategy in our game, and it will give a higher probability of winning than F when everyone else plays F, contrary to the assumption that F is a symmetric equilibrium strategy of our game. All that remains to do is to show how to construct H.

Since $E[F] = b$, playing F in the common-value all-pay auction (when everyone else does so as well) yields an expected utility of $nb(1/n) - b = 0$. By Corollary 1, we know that F is continuous. Let $W(x)$ denote the probability that i wins given that i realizes outcome x, when all other agents use F. If G constitutes a beneficial deviation, then there must exist an outcome x in the support of F such that

$$(nb)W(x) - x > 0$$

Because $W(x) \leq 1$, we have $x < nb$.

First, suppose that $x \geq b$. Let H be the lottery that places mass b/x at x and $1 - b/x$ at 0; its expectation is b, so it is a valid strategy in our game. If an agent plays H in our game when everyone else plays F, then the agent wins with probability $bW(x)/x$.

But we know $bW(x)/x > 1/n$, so it constitutes a beneficial deviation, contrary to the assumption that F is a symmetric equilibrium strategy.

Now, suppose $x < b$. Let U denote the least upper bound of the support of F, which exists by Lemma 4. We have that $U > b$ (since the degenerate distribution at b is never a best response) and, by continuity, $W(U) = 1$. Let α and β satisfy $\alpha x + \beta U = b$ and $\alpha + \beta = 1$. Now, we let distribution H place mass α at x and β at U, so that $E(H) = b$ and thus H is a valid strategy for our game. An agent's expected utility from playing H in the common-value all-pay auction, given that all other agents play F, is given by

$$\alpha(nbW(x) - x) + \beta(nb - U) \tag{A-26}$$

We note that the term on the left in (A-26) is positive. Furthermore, when $U \leq nb$, the right term in (A-26) is non-negative, so that (A-26) is strictly positive. In this case, the expectation of H is equal to the expectation of F (which is b), so it follows that the probability of winning using H is greater than the probability of winning using F (when everyone else uses F). Since H is a valid strategy in our game, we obtain the desired contradiction. All that remains to show is that $U \leq nb$, which we prove below in Lemma 9. This completes the proof.

Lemma 9. *Let F denote a symmetric equilibrium of the equal-budget game and let U denote the least upper bound of its support. Then $U \leq nb$.*

Proof: Suppose on the contrary that $U > nb$. By Lemma 6, 0 is the greatest lower bound of the support of F. It follows that F places positive probability mass on outcomes larger than nb, and similarly, F places positive probability mass on outcomes in the neighborhood of 0. Since F is a symmetric equilibrium strategy, the probability of winning (and subsequent expected utility) from playing F is $1/n$ (when everyone else plays F). We will show that under the premise that $U > nb$, there exists a beneficial deviation strategy.

Consider a probability mass ϵ_{nb} spread over some region $[nb + \phi, nb + \phi']$, where $0 < \phi < \phi'$ and $nb + \phi' < U$. Continuity of F_{-i} implies that for any $\epsilon > 0$, there exists a $\delta > 0$ such that for $|x - 0| < \delta$, $|F_{-i}(x) - F_{-i}(0)| = F_{-i}(x) < \epsilon$. Set $\epsilon = (1/n)(1 - F_{-i}(nb + \phi'))$, so that $F_{-i}(nb + \phi') = 1 - n\epsilon$, and consider a neighborhood of 0, $[0, \psi]$, where $0 < \psi < \delta$. Denote the probability mass spread over $[0, \psi]$ by ϵ_0, so that

$$\int_0^{\psi} dF(x) = \epsilon_0$$

We note that

$$\int_0^{\psi} F_{-i}(x)dF(x) < F_{-i}(\psi)\epsilon_0$$

and since $\psi < \delta$,

$$F_{-i}(\psi) < \epsilon$$

The weighted expectation of the regions over which ϵ_0 and ϵ_{nb} are spread is given by $\int_0^{\psi} x dF(x) + \int_{nb+\phi}^{nb+\phi'} x dF(x)$. Without loss of generality, we can assume that

$$\int_0^{\psi} x dF(x) + \int_{nb+\phi}^{nb+\phi'} x dF(x) > (\epsilon_0 + \epsilon_{nb})b \tag{A-27}$$

If that is not the case, we can choose a smaller ψ (and correspondingly, ϵ_0) such that (A-27) is indeed satisfied. For such ψ, as we increase ϕ, ϵ_{nb} shrinks. We shrink ϵ_{nb} until ϵ_0 becomes sufficiently large relative to ϵ_{nb} that

$$\int_0^{\psi} x dF(x) + \int_{nb+\phi}^{nb+\phi'} x dF(x) = (\epsilon_0 + \epsilon_{nb})b \tag{A-28}$$

We can then modify F into a distribution H that has the same expectation, as follows:

- Remove mass ϵ_{nb} from the region $[nb + \phi, nb + \phi']$,
- Remove mass ϵ_0 from $[0, \psi]$,
- Place the combined mass of $\epsilon_{nb} + \epsilon_0$ on playing F again.

Since ϵ_{nb} is taken from outcomes larger than nb, in order for (A-28) to hold we must have $\epsilon_{nb} < (1/n)(\epsilon_0 + \epsilon_{nb})$. Hence, such deviation would result in an increase in the probability of winning for the deviating agent of at least

$$(1/n)(\epsilon_{nb} + \epsilon_0) - \epsilon_0 F_{-i}(\psi) - \epsilon_{nb} F_{-i}(nb + \phi') >$$

$$(1/n)(\epsilon_{nb} + \epsilon_0) - \epsilon\epsilon_0 - \frac{\epsilon_0 + \epsilon_{nb}}{n}(1 - n\epsilon) = \epsilon\epsilon_{nb} > 0$$

It follows that F is not a best response, which is contrary to assumption. ■