Minimizing Regret when Dissolving a Partnership

Stergios Athanassoglou * Steven J. Brams † Jay Sethuraman‡

November 2008

Abstract

We study the problem of dissolving an equal-entitlement partnership when the objective is to minimize maximum regret. We initially focus on the family of linear-pricing mechanisms and derive regret-optimizing strategies. We also demonstrate that there exist linear-pricing mechanisms satisfying ex-post efficiency. Next, we analyze a binary-search mechanism which is ex-post individually rational. We discuss connections with the standard Bayesian-Nash framework for both linear and binary-search mechanisms. On a more general level, we show that if entitlements are unequal, ex-post efficiency and ex-post individual rationality impose significant restrictions on permissible mechanisms. In particular, they rule out both linear and binary-search mechanisms.

1 Introduction

1.1 Motivation and Related Work

Many partnership agreements include buy-sell clauses that stipulate that a partnership may be dissolved if one partner (the proposer) offers to buy out the other partner (the responder) at some proposed price for each share. While the responder may accept this offer, she may also turn this offer around and buy out the proposer at the same price [3, 7, 15].

We propose a class of symmetric procedures, whereby the partners make simultaneous offers; the partner who makes the higher offer becomes the buyer and the other partner the seller, where

*The Earth Institute, Columbia University, New York, NY; email: sa2164@columbia.edu

†Department of Politics, New York University, New York, NY; email: steven.brams@nyu.edu. Brams thanks James F. Ring for introducing him to standard buy-sell agreements. For an implementation of the nonstandard buy-sell mechanism analyzed in this paper, see http://www.FairOutcomes.com and the procedure called "Fair Buy-Sell."

‡IEOR Department, Columbia University, New York, NY; email: jay@ieor.columbia.edu

the price is some intermediate value. When two partners have equal entitlements (i.e., 50 percent each), this procedure satisfies several desirable properties. For example, under specified conditions for determining the price, an optimal strategy of the partners is to be truthful in order to minimize their maximum regret, a goal likely to be appealing to risk-averse partners. It also relieves the proposer in standard buy-sell from having to make herself indifferent between being the buyer or being the seller, like the cutter in "I cut, you choose." When the partners are truthful, it is efficient in awarding the partnership to the partner that values it more.

To set the stage for our analysis, consider a group of agents that jointly owns an indivisible good. Each agent is entitled to a fraction of the good, and the fractions sum to 1. In addition, each agent attaches some value to obtaining sole ownership of the good. Our objective is to design procedures for allocating the good to one agent and compensating the other agents for not obtaining it. While this general class of problems—commonly referred to as partnership-dissolution problems—has been extensively studied in the literature [4, 12, 5], almost all of the existing work (with the notable exception of Linhart [8] and Linhart and Radner [9]) assumes that agent valuations are independently drawn from distributions that are common knowledge.

The distinguishing feature of this paper is its focus on minimizing *maximum regret*, which is defined as the worst-case difference between the actual profit achieved by an agent and her optimal profit, given complete information. An important advantage of this approach is its substantial weakening of the common-prior assumption.

Specifically, it suffices to assume that agent valuations are drawn from an interval whose endpoints are common knowledge. Because our model focuses on the worst case, it is insensitive to the particular manner in which agent valuations are drawn. In contrast, the traditional model is sensitive to the *distribution* of valuations, because it seeks to optimize *expected* profit.

The partnership-dissolution problem has a rich history in the economic-theory literature. Typically, it is modeled as a bargaining game in which players maximize expected profit in a Bayesian-Nash framework. The key properties explored are ex-post efficiency, individual rationality, and incentive compatibility. A brief description of these concepts is warranted: *Ex-post efficiency* is satisfied if and only if the agent with the highest valuation receives the good; *interim individual rationality* is satisfied if the mechanism affords positive expected profit to all agents at the interim stage (i.e., after each agent learns her valuation), regardless of the agents' valuations. Finally, a mechanism is said to be *incentive-compatible* if it induces agents to be truthful in their equilibrium bidding for the good.

Chatterjee and Samuelson [4] consider an important special case of the partnership-dissolution problem in which there are two agents, one of whom (the *seller*) owns the entire good. They assume that agents' valuations for the good are independently distributed random variables and analyze linear-pricing mechanisms, wherein the price of the good is a convex combination of the players' two bids. In this context, they derive a necessary and sufficient condition for equilibrium bidding strategies. They give an explicit solution for the case of symmetric uniform [0,1] valuations and a split-the-difference price—that is, when the price is set to be the mean of the two bids, provided the seller bids less than the buyer. They show that the seller has an incentive to overstate her true valuation when it is below 3/4, whereas the buyer has an incentive to understate her true valuation when it is above 1/4; they also show that under certain conditions, a mutually beneficial trade will not occur in equilibrium, rendering the mechanism inefficient.

Myerson and Satterthwaite [12] generalize this two agent buyer-seller framework. In contrast to [4], they do not restrict their analysis to linear-pricing mechanisms. They provide a characterization of all incentive-compatible and interim individually rational mechanisms, showing these two properties to be incompatible with ex-post efficiency. Their result delineates the limitations inherent in dissolving a partnership in a satisfactory manner.

An important generalization of the Myerson-Satterthwaite model, due to Cramton, Gibbons and Klemperer [5], assumes n agents who each own a share r_i of the good, where $\sum_{i=1}^{n} r_i = 1$.[1] They depart from the standard model by allowing for a redistribution of the partnership ownership shares, thereby not limiting their attention to its dissolution (wherein one agent is assumed to take sole possession of the good). In this context, they characterize the set of all incentive-compatible and individually rational mechanisms. Furthermore, they provide a simple necessary and sufficient condition for such mechanisms to be ex-post efficient. Essentially, such a dissolution is possible if and only if initial endowments are sufficiently close to the equal-endowment vector; efficient dissolution is never possible for extreme cases of ownership asymmetry, such as in the buyer-seller framework.

McAfee [10] examines simple mechanisms for dissolving equal-share two-agent partnerships with an arbitrary degree of risk aversion and derives equilibrium strategies for the mechanisms considered. His model allows for an "outside option," which can be exercised only if both parties agree to it. One of these mechanisms (winner's bid auction), whereby the agent with the highest bid pays the loser one half of her bid, is ex-post efficient.

[1] To recover the Myerson and Satterthwaite model [12], set $n = 2$, $r_1 = 1$, and $r_2 = 0$.

1.2 Our Focus

A key element in the partnership-dissolution problem is that an agent does not know exactly the bids of the other agents. If an agent (somehow) obtains this information, her strategy is very simple to calculate. We evaluate the efficacy of any given strategy by comparing it to the agent's optimal strategy in hindsight (presuming she has complete information about the bids of the other agents). The smaller this difference, the better her strategy, because she is closer to her optimum.

Of course, this difference is sensitive to the particular values (and hence bids) of the other agents. We therefore focus on the worst difference, where "worst" is defined with respect to the valuations of all the other agents. The difference between an agent's actual profit and her optimal profit is her *regret*.

We focus on finding a strategy whose *maximum regret* is minimized. This approach has its roots in traditional decision theory [14] and is standard in the analysis of online algorithms in computer science [2] (although there the focus is on relative regret as opposed to absolute regret). Perhaps most relevant to our work is a paper by Linhart and Radner [9], in which the authors study a bilateral bargaining game under the assumption that agents seek to minimize their maximum regret. Similar to [4] and [12], they introduce a buyer-seller framework and examine a sealed bid, split-the-difference mechanism in which trade occurs if and only if the buyer outbids the seller. They assume that bargaining occurs over both price and quantity and derive regret-optimizing strategies. Linhart discusses the minimax-regret objective in bargaining games further in [8]. Important recent applications of minimax regret appear in robust newsvendor models [13] and robust monopoly pricing [1, 6].

Organization of the Paper. The structure of the paper is as follows. In Section 2 we formally define the model and our concept of regret. In Section 3 we focus on the family of linear mechanisms and derive regret-optimizing strategies for the two-agent model. In Section 4 we discuss a special binary search mechanism and prove that it induces a truthful regret-optimizing equilibrium. If agents' valuations are drawn from independent uniform [0,1] distributions, then the binary search mechanism induces a truthful Bayesian-Nash equilibrium as well. In Section 5 we show that, when entitlements are unequal, the linear and binary-search mechanisms do not maintain their attractive properties. In such entitlement environments, ex-post efficiency and ex-post individual rationality impose significant restrictions on permissible mechanisms.

2 Model Description

2.1 Two Agents

Our model has two agents, denoted 1 and 2. Each agent i owns a 50-percent share of the partnership. In addition, each values the partnership at v_i, which is private information but is known to be in the interval $[0, 1]$. We emphasize that no additional assumptions are made about the valuations. In particular, we do not assume that they come from distributions that are common knowledge, which is a restrictive assumption made in the standard approach.

We analyze *direct* mechanisms in which the two agents simultaneously submit sealed bids b_1 and b_2 for the partnership. A mechanism p is a function which takes as input the bids of the agents and determines who gets the partnership (or good), and at what price. We restrict ourselves to mechanisms in which the good is always awarded to the agent submitting the higher bid, at a price $p(b_1, b_2)$, which is also the amount of money transferred to the low bidder (thus, the mechanism is *budget-balanced*). Even within this class, we focus on mechanisms that satisfy two additional properties:

(a) **Convexity:** The price $p(b_1, b_2)$ is at most the high bid and at least the low bid; and

(b) **Anonymity:** The price $p(b_1, b_2)$ does not depend on the identity of the bidders, i.e., $p(b, b') = p(b', b)$.

We assume that the agents have linear utilities, so that an agent with valuation v and owning a share r of the partnership has a utility of rv. Suppose $b_1 > b_2$ so that agent 1 is awarded the partnership. When $r = 1/2$, the utilities of agents 1 and 2 after the dissolution are given by $v_1 - p(b_1, b_2)/2$ and $p(b_1, b_2)/2$, respectively, whereas their initial utilities are, respectively, $v_1/2$ and $v_2/2$. The profit of agent 1 is therefore

$$v_1 - p(b_1, b_2)/2 - v_1/2 = (v_1 - p(b_1, b_2))/2.$$

Similarly, agent 2's profit is given by

$$p(b_1, b_2)/2 - v_2/2 = (p(b_1, b_2) - v_2)/2.$$

2.2 Regret

The distinguishing feature of our approach is the performance measure used to evaluate a bidding strategy. This measure is standard in computer science [2], and is increasingly used in economics [1] and operations research [13].

Define the *regret* of an agent to be the difference between her *optimal* profit and her *actual* profit, where the *optimal* profit is calculated by assuming that the agent knows the other agent's bid. That is, an agent's optimal profit is the best that she could have done in hindsight. This is especially easy to calculate in the case of two agents: Suppose agent i's valuation is v, and the other agent's bid is $\hat{b}$. Then it is optimal for agent i to bid *slightly above* $\hat{b}$ to obtain the good if $\hat{b} < v$; and it is optimal for agent i to bid *slightly below* $\hat{b}$ to sell her share of the partnership if $\hat{b} > v$. Note that in the first case, agent i gets the good at the lowest possible price, and in the second case, agent i sells the good at the highest possible price. In both cases, i's optimal bid is the bid of the other agent, slightly perturbed upwards or downwards.

The measure that we use to evaluate a bidding strategy is *worst-case* regret, where the worst case is over all possible bids of the other agent. In other words, we hypothesize that each agent acts as if to *minimize* her *maximum regret*. The information in Table 1 is useful in finding a bidding strategy that minimizes the maximum regret.

Focusing on player 1, we suppose she has a valuation of v_1 and bids b_1, while we denote her opponent's bid by $\hat{b}$. The columns Actual and Optimal refer to an agent's actual and optimal profits, respectively. An agent's regret is taken to equal the difference between her optimal and actual profits. We then enumerate four cases that give rise to different actual and optimal profits, and therefore regret. For example, the third row of the table describes the situation in which agent 1 wins the good with a bid of b_1, and the optimal strategy (in hindsight) is for agent 1 to lose the bidding with a bid that is slightly less than her opponent's. In each case we now take $\hat{b}$ so

Cases	Actual	Optimal	Regret
$b_1 > \hat{b}, \hat{b} \leq v_1$	$(v_1 - p(b_1, \hat{b}))/2$	$(v_1 - p(\hat{b}, \hat{b}))/2$	$(p(b_1, \hat{b}) - \hat{b})/2$
$b_1 \leq \hat{b}, \hat{b} > v_1$	$(p(b_1, \hat{b}) - v_1)/2$	$(p(\hat{b}, \hat{b}) - v_1)/2$	$(\hat{b} - p(b_1, \hat{b}))/2$
$b_1 > \hat{b}, \hat{b} > v_1$	$(v_1 - p(b_1, \hat{b}))/2$	$(p(\hat{b}, \hat{b}) - v_1)/2$	$\hat{b}/2 + p(b_1, \hat{b})/2 - v_1$
$b_1 \leq \hat{b}, \hat{b} \leq v_1$	$(p(b_1, \hat{b}) - v_1)/2$	$(v_1 - p(\hat{b}, \hat{b}))/2$	$v_1 - \hat{b}/2 - p(b_1, \hat{b})/2$

Table 1: Actual and Optimal Profits: Case Analysis

as to maximize regret—this is the worst possible regret of agent 1, assuming she knows nothing at all about agent 2's valuation or bid (except for the range of agent 2's valuation). We arrive at the following expression for maximum regret, as a function of agent 1's valuation and bid. An equivalent expression holds for agent 2.

$$R_1(b_1) = \max \begin{cases} \max\{p(b_1,\hat{b}) - \hat{b}\}/2, & \hat{b} \le \min\{b_1, v_1\} \\ \max\{\hat{b} - p(b_1,\hat{b})\}/2, & \hat{b} \ge \max\{b_1, v_1\} \\ b_1 - v_1, & b_1 > v_1 \\ v_1 - b_1, & b_1 \le v_1 \end{cases}$$

Thus, given a mechanism p, both agents wish to compute bidding strategies, b_1 and b_2, which, in equilibrium, minimize the functions $R_1(b_1)$ and $R_2(b_2)$, respectively.

2.3 Properties of Mechanisms

Armed with the particular way in which bidding strategies are evaluated, we now state properties that we would like a mechanism to satisfy. These properties have analogs in the standard Bayesian-Nash setting that we will also investigate.

Efficiency. A mechanism p is said to be *ex-post efficient* if there are equilibrium (regret-optimizing) strategies which always award the good to the agent with the highest *valuation*. In other words, if $b_1(\cdot)$ and $b_2(\cdot)$ are equilibrium (regret-optimizing) strategies induced by the mechanism p, ex-post efficiency implies:

$$v_1 \le v_2 \Leftrightarrow b_1(v_1) \le b_2(v_2) \;\; \forall v_1, v_2.$$

Individual Rationality. A mechanism p is said to be *ex-post individually rational* if it guarantees a non-negative payoff under any realization of the valuations. Specifically, if $b_1(\cdot)$ and $b_2(\cdot)$ are the equilibrium regret-optimizing strategies induced by the mechanism p, ex-post individual rationality implies for $i = 1, 2$ and $j \neq i$:

$$\begin{aligned} v_i - p(b_i(v_i), b_j(v_j)) &\ge 0, \;\; \forall \{(v_i, v_j) : b_i(v_i) \ge b_j(v_j)\}, \\ p(b_i(v_i), b_j(v_j)) - v_i &\ge 0, \;\; \forall \{(v_i, v_j) : b_i(v_i) \le b_j(v_j)\}. \end{aligned}$$

Truthfulness. A mechanism p is said to be *truthful* if it admits a regret-optimizing equilibrium in which agents bid truthfully. Truthful bidding for an agent i is taken to mean $b_i(v_i) = v_i$.
The following proposition shows that ex-post individual rationality and truthfulness are equivalent.

Proposition 1 *A mechanism is ex-post individually rational if and only if it is truthful.*

Proof. That truthfulness in conjunction with convexity implies ex-post individual rationality is obvious. To show the converse, assume that the mechanism p is ex-post individually rational and that agent 1 observes a valuation of v_1. Then for $b_1(\cdot)$ and $b_2(\cdot)$ to be equilibrium strategies induced by p, we must have:

$$\begin{aligned} p(b_1(v_1), b_2(v_2)) &\leq v_1, \ \forall v_2 : b_2(v_2) \leq b_1(v_1), \\ p(b_1(v_1), b_2(v_2)) &\geq v_1, \ \forall v_2 : b_2(v_2) \geq b_1(v_1). \end{aligned}$$

Hence, $p(b_1(v_1), b_1(v_1)) = v_1$. On the other hand, convexity implies $p(b_1(v_1), b_1(v_1)) = b_1(v_1)$. Putting the two equalities together, we conclude that $b_1(v_1) = v_1$. ■

Because truthfulness implies ex-post efficiency, we have the following:

Proposition 2 *If a mechanism p is ex-post individually rational, it is ex-post efficient.*

3 Linear Mechanisms

In this section we focus on a natural class of mechanisms in which the price is set to be a convex combination of the bids. If the agents bid b_1 and b_2, respectively, the partnership is sold to the highest bidder at the price $p = \lambda \min\{b_1, b_2\} + (1 - \lambda) \max\{b_1, b_2\}$ for $\lambda \in [0, 1]$. (Note that p is λ times the low bid plus $(1 - \lambda)$ times the high bid.) If $\lambda = 1/2$, the two agents split the difference, yielding the canonical mechanism in this class.

3.1 Linear Mechanisms and Regret

The main result in this section is the derivation of regret-optimizing strategies for two agents. That this analysis extends to the case of more than two agents is straightforward.[2] We compare different linear mechanisms in terms of their min-max regret and show that there is no dominance relation between any two mechanisms within this class.

Theorem 1 *Fix a linear mechanism λ. The regret-optimizing bidding strategy for agents $i = 1, 2$ is given by:*

$$b_i(v_i) = \begin{cases} \frac{2}{2+\lambda} v_i + \frac{\lambda}{3} & 0 \leq v_i \leq \frac{\lambda(2+\lambda)}{3} \\ \lambda & \frac{\lambda(2+\lambda)}{3} \leq v_i \leq \frac{(4-\lambda)\lambda}{3} \\ \frac{2}{3-\lambda} v_i + \frac{\lambda(1-\lambda)}{3(3-\lambda)} & \frac{(4-\lambda)\lambda}{3} \leq v_i \leq 1 \end{cases} \quad (1)$$

[2] Proof available upon request.

Proof. Let $b_1(\cdot)$ be the bidding strategy of agent 1. We search for a symmetric equilibrium. Upon learning her valuation to be v_1, she places a bid of $b_1(v_1)$. Suppose also that agent 2's bids are in the interval $[c, d]$ for $0 \leq c \leq d \leq 1$. As we showed in Section 2, the function that agent 1 wishes to minimize is the following.

$$R_1(b_1) = \max \begin{cases} \lambda(d - b_1)/2, & d \geq \max\{v_1, b_1\} \\ (1-\lambda)(b_1 - c)/2, & c \leq \min\{v_1, b_1\} \\ b_1 - v_1, & b_1 > v_1 \\ v_1 - b_1, & b_1 \leq v_1 \end{cases}$$

The optimal bidding strategy—a $b(\cdot)$ that minimizes the maximum regret—can now be derived by examining three cases separately.

First, consider the case $v_1 < c$. In this case, the regret is the larger of $\lambda(d-b_1)/2$ and $|v_1 - b_1|$. For $b_1 \leq v_1$, both of these expressions decrease with an increase in b_1, so the optimal b_1 is at least v_1 and is given by the point of intersection of $\lambda(d - b_1)/2$ and $b_1 - v_1$, which results in

$$b_1 \ = \ \frac{v_1}{1+\lambda/2} \ + \ \frac{\lambda d/2}{1+\lambda/2} \ .$$

Next, consider the case $v_1 > d$. In this case, the regret is the larger of $(1-\lambda)(b_1 - c)/2$ and $|v_1 - b_1|$. For $b_1 \geq v_1$, both of these expressions increase with an increase in b_1, so the optimal b_1 is at most v_1 and is given by the point of intersection of $(1-\lambda)(b_1 - c_2)/2$ and $v_1 - b_1$, which results in

$$b_1 \ = \ \frac{v_1}{1+(1-\lambda)/2} \ + \ \frac{(1-\lambda)c/2}{1+(1-\lambda)/2} \ .$$

These two cases already fix the values of c and d. If agent 1 observes a value of $v_1 \leq c$. she bids *at least* her valuation, so $c \geq 0$; similarly, if she observes $v_1 \geq d$, she bids at most her valuation, so $d \leq 1$. Furthermore, as the bidding range is assumed to be contained in $[c, d]$, we may assume that agent 1 with a value of 0 bids exactly c , and that agent 1 with a value of 1 bids exactly d. These observations lead to the equations

$$c = \frac{\lambda d/2}{1+\lambda/2} \ \ d = \frac{1+(1-\lambda)c/2}{1+(1-\lambda)/2},$$

which, when solved, yield

$$c = \frac{\lambda}{3} \ \ d = \frac{2+\lambda}{3}$$

We now turn to the case in which $v_1 \in [c, d]$. We consider two subcases, depending on whether or not $b_1 \leq v_1$. In each of these cases, the optimal strategy is determined by either the intersection of $\lambda(d-b_1)/2$ and $(1-\lambda)(b_1-c)/2$, or by the intersection of one of these two terms with v_1-b_1. The first two intersect at $\lambda d+(1-\lambda)c$, which becomes λ when we substitute for the values of c and d. (The other intersections have already been calculated in our analysis of the first two cases.) Specifically, for $b_1 \leq v_1$, the optimal strategy is

$$b(v) = \begin{cases} \lambda, & 0 \leq v_1 \leq u \\ \frac{2}{3-\lambda}v + \frac{\lambda(1-\lambda)}{3(3-\lambda)}, & u \leq v_1 \leq 1 \end{cases}$$

where u is defined as the point of intersection of the two functions, i.e.,

$$\lambda = \frac{2u}{3-\lambda} + \frac{\lambda(1-\lambda)}{3(3-\lambda)} \Rightarrow u = \frac{(4-\lambda)\lambda}{3}.$$

Following the same reasoning, for $b_1 \geq v_1$, the optimal strategy is

$$b_1(v_1) = \begin{cases} \frac{2}{2+\lambda}v_1 + \frac{\lambda}{3} & 0 \leq v_1 \leq \frac{\lambda(2+\lambda)}{3}, \\ \lambda, & l \leq v_1 \leq 1 \end{cases}$$

where l is defined as the point of intersection of the two functions, i.e.,

$$\lambda = \frac{2l}{2+\lambda} + \frac{\lambda}{3} \Rightarrow l = \frac{\lambda(2+\lambda)}{3}.$$

Simple algebra verifies $0 \leq l \leq u \leq 1$. Putting these cases together, we find that the optimal bidding strategy is:

$$b_1(v_1) = \begin{cases} \frac{2}{2+\lambda}v_1 + \frac{\lambda}{3} & 0 \leq v_1 \leq \frac{\lambda(2+\lambda)}{3} \\ \lambda & \frac{\lambda(2+\lambda)}{3} \leq v_1 \leq \frac{(4-\lambda)\lambda}{3} \\ \frac{2}{3-\lambda}v_1 + \frac{\lambda(1-\lambda)}{3(3-\lambda)} & \frac{(4-\lambda)\lambda}{3} \leq v_1 \leq 1 \end{cases}$$

An equivalent expression holds for agent 2. ■

We now briefly comment on the efficiency properties of linear mechanisms.

Proposition 3 *The only ex-post efficient linear mechanisms are the ones corresponding to* $\lambda = 0$ *and* $\lambda = 1$.

Proof. Examining the bidding strategy in (1), we see that the "flat" portion in the middle is eliminated only when $\lambda(2+\lambda) = (4-\lambda)\lambda$, which is true only when $\lambda = 0$ or 1. In each of these cases, we note that the bidding strategy in (1), which is identical for both agents, is strictly increasing. ■

Comparing Linear Mechanisms. A natural question to ask is if, among all linear mechanisms, there exists one that fares the best in terms of efficiency. We say that a mechanism λ_1 *dominates* λ_2 if its interim regret is weakly lower for every $v \in [0,1]$, and strictly lower for some $v \in [0,1]$. Conversely, we say that mechanism λ_2 is *dominated* by λ_1. The following proposition shows that there is no meaningful way to rank mechanisms.

Proposition 4 *Fix two linear mechanisms λ_1 and λ_2, where $\lambda_i \in [0,1]$. There is no dominance relation between the two mechanisms.*

Proof. It is easy to see that for $\lambda \in [0,1]$, min max regret is given by the following expression:

$$R(\lambda, v) = \begin{cases} \frac{\lambda}{3} - \frac{\lambda}{2+\lambda}v & 0 \leq v \leq \frac{\lambda(2+\lambda)}{3} \\ \frac{(1-\lambda)\lambda}{3} & \frac{\lambda(2+\lambda)}{3} \leq v \leq \frac{(4-\lambda)\lambda}{3} \\ \frac{1-\lambda}{3-\lambda}v - \frac{\lambda(1-\lambda)}{3(3-\lambda)} & \frac{(4-\lambda)\lambda}{3} \leq v \leq 1 \end{cases}$$

From the above it follows that there is no mechanism that simultaneously minimizes maximum regret for every $v \in [0,1]$. The argument goes as follows. Fix two mechanisms λ_1 and λ_2. Initially assume $0 < \lambda_2 < \lambda_1 < 1$. Then for $v \in [0, \lambda_2(2+\lambda_2)/3]$, the mechanism λ_1 has a higher regret than λ_2. Conversely, for $v \in [\lambda_1(4-\lambda_1)/3, 1]$, the mechanism λ_2 has higher regret than λ_1.

Now assume that $0 < \lambda_1 < 1$ and that $\lambda_2 = 0$, whose regret is given by $1/3v$. Then for v small enough we have $v/3 < \lambda_1/3 - \lambda_1/(2+\lambda_1)v$, whereas for v large enough, we have $v/3 > (1-\lambda_1)v/(3-\lambda_1) - \lambda_1(1-\lambda_1)/(3(3-\lambda_1))$. So again there is no dominance relation.

Finally we take $0 < \lambda_1 < 1$ and $\lambda_2 = 1$, whose regret is given by $1/3 - v/3$. For v large enough we have $1/3 - 1/3v < (1-\lambda_1)(3-\lambda_1)v - \lambda_1(1-\lambda_1)/(3(3-\lambda_1))$, whereas for v small enough we have $1/3 - v/3 > \lambda_1/3 - \lambda_1 v/(2+\lambda_1)$. Again, there is no dominance relation.[3] ■

Finally, it is not difficult to show that the mechanism $\lambda = 1/2$ minimizes the worst-case regret *ex-ante*. This is because

$$\max_{v \in [0,1]} R(\lambda, v) = \max \left\{ \frac{\lambda}{3}, \frac{(1-\lambda)}{3} \right\},$$

which is minimized at $\lambda = 1/2$. By comparison, the mechanisms $\lambda = 0$ and $\lambda = 1$ fare the worst under this measure.

The above facts are graphically depicted in Figure 1. The figure shows that while $\lambda = 1/2$ is overall optimal (i.e., achieves the lowest maximum regret over all possible v's), other λ's may be better for v's outside a "middle" range between 5/12 and 7/12. To minimize maximum regret, therefore, agents with middling valuations are well-advised to choose a split-the-difference price mechanism.

[3] Note that the mechanisms $\lambda = 0, 1$ do not dominate each other either.

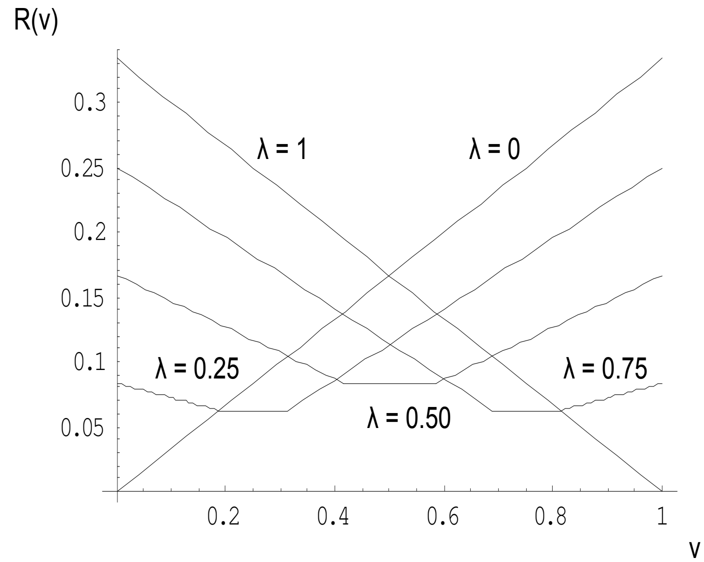

Figure 1: Regret of Linear Mechanisms with Different λ's

Linear Mechanisms under Bayesian-Nash We close this section by briefly commenting on linear mechanisms in the traditional Bayesian-Nash profit-maximizing framework. We note the following theorem due to Cramton et al. [5].

Theorem 2 *Suppose there are two agents whose valuations are drawn from a distribution* $F(\cdot)$ *in* $[0, 1]$ *that is common knowledge. Then the bidding strategy*

$$h(v) \quad = \quad v - \frac{\int_{F^{-1}(\lambda)}^{v} (F(x) - \lambda)^2 dx}{(F(v) - \lambda)^2}$$

is a Nash equilibrium.

In fact, Cramton et al. prove a more general version of this theorem, allowing for an arbitrary number of agents.

It is not difficult to show that when players use the equilibrium strategy given by Theorem 2, their regret will be given by

$$R(\lambda, v) = \frac{1}{2} \max \left\{ h(1) - h(v), h(v) - h(0) \right\},$$

which is never less and sometimes more than that given by the regret-minimizing strategy (Equation 1) for all v's and λ's. Presumably, risk-averse players would prefer the latter strategy over the Bayesian-Nash equilibrium strategy.

4 Binary Search Mechanism

We now turn to a different mechanism that resembles the familiar binary search. This mechanism is attractive, because it satisfies many of the desirable properties outlined earlier.

Suppose that the agents bid b_1 and b_2, and it is common knowledge that their valuations are in $[0, 1]$. The binary search mechanism proceeds as follows: If b_1 and b_2 are on opposite sides of $1/2$, then the procedure terminates, and the good goes to the high bidder with the price set at $1/2$. If b_1 and b_2 are both equal to or less than $1/2$, the procedure is applied to the interval of interest, $[0, 1/2]$. We now check if b_1 and b_2 are on opposite sides of $1/4$, in which case the price of the good is set at $1/4$; otherwise, the interval of interest is halved yet again, and the procedure is repeated recursively. Similarly, if b_1 and b_2 are both greater than $1/2$, the procedure is applied to the interval $[1/2, 1]$. This continues until the two agents are on different sides of the relevant candidate price.[4]

For example, if $b_1 = 1/8$ and $b_2 = 3/7$, the procedure terminates after two steps at which point the good is sold to agent 2 at a price of $1/4$. In theory, this process can take an unbounded number of iterations to terminate. Clearly, the mechanism is anonymous, and price is always between the minimum and maximum bids.

We show that this simple mechanism has some attractive properties. First, we prove that it induces a truthful regret-optimizing equilibrium. Second, in the case of two-agents and i.i.d. uniform [0,1] valuations, we show that truthful bidding is a Nash equilibrium.

4.1 Binary Search and Regret

We start with the following result:

Theorem 3 *The binary search mechanism induces a truthful regret-optimizing equilibrium. It is therefore ex-post individually rational.*

Proof. The worst-case regret is written as:

$$R(b) = \max \begin{cases} 1/2 \max(\hat{b} - p(b, \hat{b})), & \hat{b} \geq \max\{v, b\} \\ 1/2 \max(p(b, \hat{b}) - \hat{b}), & \hat{b} \leq \min\{v, b\} \\ b - v, & b > v \\ v - b, & b \leq v \end{cases}$$

[4]If an agent's bid is equal to the candidate price, then the mechanism treats the bid as if it were greater than the price. For instance, $p(b_1, 1/2) = 1/2$ if $b_1 < 1/2$.

Assume that one agent bids truthfully, i.e., $b(v) = v$. Then truthful bidding by the other agent will result in a worst-case regret of 1/4. The reasoning is as follows. First, truthfulness implies that the expressions on the third and fourth lines of the right-hand side of $R(b)$ are zero. Now assume that a player observes $v < 1/2$ and bids $b(v) = v < 1/2$. The worst regret is realized when her opponent has a valuation of 1 and bids $\hat{b} = 1$, leading to $1/2 \max(\hat{b} - p(b, \hat{b}))$ that gives a regret of $1/2(1 - p(1, b(v))) = 1/2(1 - 1/2) = 1/4$. Similarly, when $b(v) = v \geq 1/2$, the worst regret is realized when her opponent has a valuation of 0 and bids $\hat{b} = 0$, leading to $1/2 \max(p(b, \hat{b}) - \hat{b})$ that gives regret of $1/2(1/2 - 0) = 1/4$. Thus, we have established that truthfulness yields a maximum regret of 1/4.

Now assume that a player deviates and bids $b(v) \neq v$. If $|b(v) - v| > 1/4$, then she does strictly worse than if she were truthful. If $|b(v) - v| \leq 1/4$, then the expressions on the first and second lines of the right-hand side of $R(b)$ will again have a maximum regret of 1/4, because the other agent will be truthful and will bid up in the whole interval $[0, 1]$. In fact, any strategy of the sort $|b(v) - v| \leq 1/4$ will be a best response to truthful bidding. ■

Efficient Equilibrium. While truthful bidding is an equilibrium, it is not an efficient one. That is, it does not achieve the lowest maximum regret among all equilibria. Suppose that an agent does not bid truthfully and rather bids in an interval $[c, d]$, where $0 \leq c \leq 1/2 \leq d \leq 1$ and $c = 1 - d$. Then the function to be minimized is the following:

$$R(b) = \max \begin{cases} \frac{1}{2}(d - 1/2), & v \leq d, b \leq 1/2 \\ \frac{1}{2}(1/2 - c), & v \geq c, b \geq 1/2 \\ b - v, & b > v \\ v - b, & b \leq v \end{cases}$$

A similar analysis to that in the previous section shows that $c = 1/6$ and $d = 5/6$, from which we obtain the symmetric equilibrium strategy $b(v) = 2/3v + 1/6$, yielding a worst-case regret of 1/6. This equilibrium is clearly better than the truthful one. But the gain in efficiency comes at a price: The strategies are no longer ex-post individually rational.

4.2 Uniform Distributions and Nash Equilibria

In this section we focus on two agents with uniform [0,1] valuations and prove that bidding one's valuation is a Nash equilibrium. We show that this result is not true in the case of general distributions.

Theorem 4 *Suppose that the agents' valuations are iid uniform [0,1] random variables. The binary search mechanism induces a truthful Nash equilibrium. It is therefore ex-post individually rational.*

Proof. Suppose that agents do not submit sealed bids but rather are allowed simultaneously to declare if they are above or below the candidate price. That is, in round 1 they declare if they are above or below 1/2. If there is a second round, they declare if they are above or below 1/4 or 3/4, and so on. In this dynamic game, we will show that truthfulness is a subgame-perfect Nash equilibrium.

Suppose that agent 2 bids truthfully and that agent 1's valuation is $v_1 = 1/2 - \epsilon$, where $\epsilon \geq 0$.[5] Since agent 2 is truthful, her bid will be uniformly distributed in [0,1].

Assuming agent 1 is truthful, let her payoff, should the game end in the k^{th} round, be indexed by $\pi_1(k)$. Denoting the candidate price at stage k by $p(k)$, we may write

$$\pi_1(k) = \frac{1}{2}\Big|v_1 - p(k)\Big|$$

Clearly, if the game ends in period 1, then $\pi_1(1) = \epsilon/2$. If $k > 1$, then $\pi_1(k)$ is bounded from below by the quantity $\frac{1}{2}\left[\frac{1}{2^k} - \epsilon\right]$. We can establish this with the following simple argument:

$$\pi_1(k) = \frac{1}{2}\Big|v_1 - p(k)\Big| = \frac{1}{2}\Big|\frac{1}{2} - \epsilon - p(k)\Big| \geq \frac{1}{2}\Big(\frac{1}{2} - \epsilon - p(k)\Big) \geq \frac{1}{2}\Big(\frac{1}{2^k} - \epsilon\Big).$$

Since agent 2 is truthful, the probability that the game ends at stage k is simply the probability that the agent's valuation lies in the relevant interval of length $1/2^k$. It is thus equal to $1/2^k$. So, player 1's expected payoff π_1 will satisfy:

$$\pi_1 = \frac{1}{2}\sum_{i=1}^{\infty}\frac{1}{2^k}\pi_1(k) \geq \frac{1}{2}\left[\frac{\epsilon}{2} + \sum_{k=2}^{\infty}\frac{1}{2^k}\left(\frac{1}{2^k} - \epsilon\right)\right] = \frac{1}{24}$$

with equality if and only if $\epsilon = 0$.

Now let us return to agent 1's actual strategy. Without loss of generality, assume that she is not truthful in the first round and declares a valuation that is greater than 1/2. After her initial misrepresentation, suppose that she is truthful up to time $k_1 - 1$, where $k_1 \geq 2$. Should the game end in round k, where $2 \leq k \leq k_1 - 1$, her payoff will be $\tilde{\pi}_1(k) = 1/2^k + \epsilon$. Now suppose that at time k_1 she is again untruthful. Then should the game end at that time, her payoff will be $\tilde{\pi}_1(k_1) = -1/2^{k_1} - \epsilon$. If the game does not end then, assume that she is again truthful from time

[5] An analogous argument works for $\epsilon \leq 0$.

k_1+1 until k_2-1. In all those rounds her payoff will be $1/2^k + 1/2^{k_1} + \epsilon$. At time k_2 it will be $-1/2^{k_2} - 1/2^{k_1} - \epsilon$. We can repeat this reasoning for all subsequent successive time periods in which agent 1 is untruthful.

It is evident that for any strategy that agent 1 adopts, all relevant information is captured by these successive intervals of truthful and untruthful behavior. Writing the expected profit we obtain

$$\begin{aligned}\tilde{\pi}_1 &= \sum_{k=1}^{\infty} \frac{1}{2^k}\tilde{\pi}_1(k) = \frac{1}{2}\Big[-\frac{\epsilon}{2} + \sum_{k=2}^{k_1-1} \frac{1}{2^k}\big(\frac{1}{2^k}+\epsilon\big) - \frac{1}{2^{k_1}}\big(\frac{1}{2^{k_1}}+\epsilon\big) + \\ &+ \sum_{k=k_1+1}^{k_2-1} \frac{1}{2^k}\big(\frac{1}{2^k}+\frac{1}{2^{k_1}}+\epsilon\big) - \frac{1}{2^{k_2}}\big(\frac{1}{2^{k_2}}+\frac{1}{2^{k_1}}+\epsilon\big) + \ldots\Big].\end{aligned}$$

Focus on the terms $\{\epsilon, 1/2^{k_1}, 1/2^{k_2}, \ldots\}$. These terms are subtracted from the agent's payoff should the game end in rounds $\{\{1, k_1, k_2, \ldots\}, \{k_1, k_2, ..\}, \{k_2, \ldots\}, \ldots\}$, respectively. Furthermore, these losses will never be fully recovered by truthful behavior in previous and later rounds. For example, focusing on ϵ we have

$$-\frac{\epsilon}{2} + \sum_{k=2}^{k_1-1}\big(\frac{\epsilon}{2^k}\big) - \frac{\epsilon}{2^{k_1}} + \sum_{k=k_1+1}^{k_2-1}\big(\frac{\epsilon}{2^k}\big) - \frac{\epsilon}{2^{k_2}} + \ldots \leq 0,$$

with equality if and only if $k_1 = \infty$. Applying this logic to all terms $1/2^k$ such that k corresponds to a round in which agent 1 is untruthful, we may write:

$$\tilde{\pi}_1 \leq \frac{1}{2}\Big[\frac{\epsilon}{2} + \sum_{k=2}^{\infty}\frac{1}{2^k}\big(\frac{1}{2^k} - \epsilon\big)\Big].$$

Thus we can conclude that

$$\tilde{\pi}_1 \leq \pi_1$$

with equality if and only if $\epsilon = 0$ and $k_1 = \infty$. In fact, we also have

$$\tilde{\pi}_1 \leq \hat{\pi}_1$$

where $\hat{\pi}_1$ is the agent's payoff when she is untruthful only in round 1 and is truthful thereafter. This observation establishes that the equilibrium is subgame-perfect. ■

General Distributions. The binary search mechanism as described is designed specifically with the uniform distribution in mind. For general distributions, the most obvious way to extend it would be to fix prices in terms of the distribution quantiles. Thus, if the two bids are on opposite sides of the 50-percent quantile, then the price is set to that quantile. The process iterates on either the 25-percent or 75-percent quantile, and so on.

When valuations are generally distributed, it is easy to construct examples in which truthful reporting is not a Nash equilibrium. Let $\alpha_1, \alpha_2, \alpha_3$ be the 50, 25, and 75-percent quantiles, respectively. Assume that $r_1 = r_2$, agent 2 is truthful, and $v_1 < \alpha_1$. A truthful response will yield the payoff

$$\pi_1 = \frac{1}{2}\left[\frac{1}{2}(\alpha_1 - v_1) + \frac{1}{4}|\alpha_2 - v_1| + ...\right].$$

Now assume that, instead, agent 1 bids exactly α_1. This will yield the payoff

$$\tilde{\pi}_1 = \frac{1}{2}\left[\frac{1}{2}(v_1 - \alpha_1) + \frac{1}{4}(\alpha_3 - v_1) + ...\right].$$

Now we can see that if v_1 is close enough to α_1 and α_2, and if α_3 is far enough from α_1, then

$$\frac{1}{2}(\alpha_1 - v_1) + \frac{1}{4}|\alpha_2 - v_1| < \frac{1}{2}(v_1 - \alpha_1) + \frac{1}{4}(\alpha_3 - v_1).$$

We observe that if the distribution's right tail is sufficiently heavy, bidding α_1 dominates bidding v_1.

5 Unequal Entitlements and General Mechanisms

As we saw in sections 3 and 4, when both agents have a 50-percent entitlement to the partnership and wish to minimize maximum regret, linear mechanisms with $\lambda \in \{0, 1\}$ achieve ex-post efficiency. The binary-search mechanism performs even better: It is ex-post individually rational. Furthermore, as we established in Theorem 4, it maintains this appealing property even when agents have uniform [0,1] i.i.d. valuations and wish to maximize expected profit.

In this section we show that the attractive properties of the linear mechanisms and the binary-search mechanism do not extend when entitlements are unequal. In addition, we show that, on a general level, ex-post efficiency and ex-post individual rationality place significant restrictions on the kinds of mechanisms that may be used.

5.1 Linear Mechanisms

Proposition 5 *When $r_1 \neq r_2$, there is no ex-post efficient linear mechanism.*

Proof. When $r_1 \neq r_2$, using a similar technique as in the proof of Theorem 1,[6] the regret-optimizing bidding strategy for agent i is given by

$$b_i(v_i) = \begin{cases} \frac{v_i}{1+\lambda r_i} + \frac{\lambda r_i}{1+r_i} & 0 \leq v_i \leq l_i \\ \frac{\lambda r_i \frac{1+\lambda r_i}{1+r_i} + (1-\lambda) r_j \frac{\lambda r_j}{1+r_j}}{\lambda r_i + (1-\lambda) r_j} & l_i \leq v_i \leq u_i \\ \frac{v_i}{1+(1-\lambda) r_j} + \frac{(1-\lambda) r_j \frac{\lambda r_j}{1+r_j}}{1+(1-\lambda) r_j} & u_i \leq v_i \leq 1 \end{cases}$$

where l_i and u_i are such that $b_i(v_i)$ is continuous.

Let us first assume that $\lambda > 0$. This immediately implies that $l_1, l_2 > 0$. Now, pick $v_1 = v_2 = v \leq \min\{l_1, l_2\}$. By efficiency and continuity we must have

$$v_1 = v_2 \Rightarrow b_1(v_1) = b_2(v_2).$$

Recalling that $v_1 = v_2 = v$, $\lambda > 0$, and $r_2 = 1 - r_1$, we may write:

$$\frac{v}{1+\lambda r_1} + \frac{\lambda r_1}{1+r_1} = \frac{v}{1+\lambda(1-r_1)} + \frac{\lambda(1-r_1)}{1+(1-r_1)} \Rightarrow r_1 = \frac{1}{2}.$$

But this contradicts our hypothesis that $r_1 \neq r_2$.

Now let us examine the case where $\lambda = 0$. Here we have $u_1 = u_2 = 0$. Assume we have $v_1 = v_2 \geq 0 = u_1$. Applying the same argument as before and substituting $\lambda = 0$, we obtain:

$$\frac{v}{1+(1-r_1)} = \frac{v}{1+r_1} \Rightarrow r_1 = \frac{1}{2}.$$

Once again, we arrive at a contradiction. ■

5.2 Binary-Search Mechanism

Our next result shows that the truthfulness of the binary search mechanism depends critically on the equal-entitlement profile. This limitation is present regardless of whether agents wish to minimize maximum regret or maximize expected profit.

Proposition 6 *Assume that agents wish to minimize maximum regret. The binary search mechanism cannot be truthful for any unequal-entitlement profile.*

[6] Details available upon request.

Proof. Without loss of generality, assume that $r_1 > r_2$. The worst-case regret for player 1 is

$$R_1(b_1) = \max \begin{cases} r_1 \max(\hat{b} - p(b_1, \hat{b})), & \hat{b} \geq \max\{v_1, b_1\} \\ r_2 \max(p(b_1, \hat{b}) - \hat{b}), & \hat{b} \leq \min\{v_1, b_1\} \\ b_1 - v_1, & b_1 > v_1 \\ v_1 - b_1, & b_1 \leq v_1 \end{cases}$$

Assume that agent 2 bids truthfully, i.e., $b_2(v_2) = v_2$. Then truthful bidding by agent 1 will result in a worst-case regret that is given by the following function:

$$R_1(b_1) = \max \begin{cases} \max\{\frac{r_1}{2}, r_2 p(v_1, 0)\} = \frac{r_1}{2}, & v_1 < \frac{1}{2} \\ \max\{r_1(1 - p(v_1, 1)), \frac{r_2}{2}\}, & v_1 \geq \frac{1}{2} \end{cases}$$

Now assume that agent 1, instead of being truthful, bids according to the strategy:

$$b_1(v_1) = \max \left\{ v_1 + \frac{r_2}{2}, 1 \right\}.$$

Then it is easy to see that the maximum regret associated with this strategy is

$$\tilde{R}_1(b_1) = \max \begin{cases} \frac{r_1}{2}, & v_1 < \frac{1}{2} - \frac{r_2}{2} \\ \max\{r_1(1 - p(v_1 + \frac{r_2}{2}, 1)), \frac{r_2}{2}\}, & \frac{1}{2} - \frac{r_2}{2} \leq v_1 \leq 1 - \frac{r_2}{2} \\ \frac{r_2}{2}, & v_1 \geq 1 - \frac{r_2}{2} \end{cases}$$

It is immediate that

$$\tilde{R}_1(b_1(v_1)) \leq R_1(v_1) \quad \forall v \in [0, 1].$$

If $r_1 < 1$, then we also have that $\tilde{R}_1(b_1(v_1)) < R_1(v_1)$ for at least one $v_1 \in [0, 1]$, so the dominance is strict. If $r_1 = 1$, then the maximum regret of a policy b_1 is given by

$$\tilde{R}_1(b_1) = \max \begin{cases} \frac{1}{2}, & b_1 < \frac{1}{2} \\ 1 - p(b_1, 1), & b_1 \geq \frac{1}{2} \\ |b_1 - v_1| \end{cases}$$

In this case, it is easy to see that a truthful bidding strategy is dominated by a host of others, for example by b_1 such that

$$b_1(v_1) = \max \begin{cases} v_1 + \frac{1}{4}, & v_1 \leq \frac{1}{4} \\ \frac{1}{2}, & \frac{1}{4} < v_1 \leq \frac{1}{2} \\ v_1, & \frac{1}{2} < v_1 \leq 1 \end{cases}$$

■

Proposition 7 *Assume that agents wish to maximize expected profit and that their valuations are iid uniform [0,1] random variables. The binary search mechanism cannot be truthful for any profile of unequal entitlements.*

Proof. Without loss of generality, assume $r_1 < r_2$. Now suppose that agent 2 is truthful and agent 1 observes a valuation of $v_1 = 1/2$. Recalling that the binary search mechanism treats a bid of 1/2 as if it were greater than 1/2, agent 1's expected profit under a truthful bidding strategy will be

$$\pi_1 = \frac{1}{2}0 + \frac{1}{4}\left(r_1\frac{1}{4}\right) + \frac{1}{8}\left(r_1\frac{1}{8}\right) + ... = \frac{r_1}{12}.$$

Now suppose agent 1 deviates and instead bids just slightly below 1/2. Her expected profit will be arbitrarily close to

$$\tilde{\pi}_1 = \frac{1}{2}0 + \frac{1}{4}\left(r_2\frac{1}{4}\right) + \frac{1}{8}\left(r_2\frac{1}{8}\right) + ... = \frac{r_2}{12}.$$

Since $r_2 > r_1$, we have $\tilde{\pi}_1 > \pi_1$. Thus, the mechanism cannot be truthful. ■

We end this section by noting that although the binary-search mechanism satisfies certain desirable properties, it is applicable only under restrictive conditions. By contrast, the class of linear mechanisms can be used in more general situations, but it it not ex-post individually rational for the special cases that we have examined.

5.3 General Mechanisms

We now consider the general mechanism-design version of the problem. Specifically, we ask if it is possible to design any mechanism satisfying all the desirable properties outlined earlier.

Ex-post Efficiency and Regret. As we saw in sections 3 and 4, when entitlement profiles are equal and agents wish to minimize maximum regret, both linear mechanisms (with $\lambda \in \{0, 1\}$) and the binary-search mechanism are ex-post efficient. At the same time, both fail to achieve ex-post efficiency when entitlements are unequal. The following theorem establishes that a large class of mechanisms cannot hope to ever achieve ex-post efficiency.

Theorem 5 *Suppose that the agents' objective is to minimize maximum regret. Consider the class of mechanisms p such that* $\frac{\partial p}{b_1}, \frac{\partial p}{b_2} > 0$ *for all* b_1, b_2*. Any such mechanism cannot be ex-post efficient for unequal-entitlement profiles.*

Proof. Consider two agents with $r_1 \neq r_2$, and fix a mechanism p. Once again we write the worst-case regret for agent 1:

$$R_1(b_1) = \max \begin{cases} r_2 \max\{p(b_1, \hat{b}) - \hat{b}\}, & \hat{b} \leq \min\{b_1, v_1\} \\ r_1 \max\{\hat{b} - p(b_1, \hat{b})\}, & \hat{b} \geq \max\{b_1, v_1\} \\ b_1 - v_1, & b_1 > v_1 \\ v_1 - b_1, & b_1 \leq v_1 \end{cases}$$

Let b_1, b_2 be the regret-optimizing strategies for players 1 and 2, respectively, and denote $b_1(0) = c_1, b_1(1) = d_1$ and $b_2(0) = c_2, b_2(1) = d_2$.

Ex-post efficiency along with the continuity of b_1, b_2 dictates that both players will have the same bidding strategy, i.e., $b_1(v) = b_2(v)$ for all $v \in [0, 1]$. Our argument proceeds in two steps. First, ex-post efficiency implies that the agent with the higher valuation will have to outbid her opponent, i.e.,

$$v_1 \leq v_2 \Leftrightarrow b_1(v_1) \leq b_2(v_2), \quad \forall (v_1, v_2)$$

Now if we assume that the bidding functions are continuous and we let $v_1, v_2 \to v$, the previous inequality becomes an equality:

$$b_1(v) = b_2(v), \quad \forall v$$

Applying this condition to $v = 0$ and $v = 1$, we obtain $c_1 = b_1(0) = b_2(0) = c_2 = c$ and $d_1 = b_1(1) = b_2(1) = d_2 = d$.

Assume that agent 1 has a valuation of $v_1 = 0$. Then we know that she must bid $b(0) = c$. For c to be the optimal bid, it must be at the intersection of $b_1 - v_1$ and $r_1 \max_{\hat{b} \geq c}\{\hat{b} - p(c, \hat{b})\}$. Since $\frac{\partial p}{b_1} > 0$, c will be the unique minimizer, as $\max_{b \geq \hat{c}}\{\hat{b} - p(\hat{c}, \hat{b})\} > \max_{b \geq c}\{\hat{b} - p(c, \hat{b})\}$ for all $\hat{c} < c$. Substituting $b_1 = c$ and $v_1 = 0$, c will have to satisfy

$$c = r_1 \max_{\hat{b} \geq c}\{\hat{b} - p(c, \hat{b})\}.$$

Recalling that agent 2 will also have to bid c upon seeing a valuation of $v_2 = 0$, an equivalent argument establishes that c will also need to satisfy

$$c = r_2 \max_{\hat{b} \geq c}\{\hat{b} - p(\hat{b}, c)\}.$$

Thus we obtain

$$r_1 \max_{\hat{b} \geq c}\{\hat{b} - p(c, \hat{b})\} = r_2 \max_{\hat{b} \geq c}\{\hat{b} - p(\hat{b}, c)\}.$$

Anonymity implies that $p(c, \hat{b}) = p(\hat{b}, c)$. Therefore, unless $\max_{\hat{b}>c}\{\hat{b} - p(c, \hat{b})\} = 0$, the above implies that $r_1 = r_2$, a contradiction. So we must have

$$\max_{\hat{b}\geq c}\{\hat{b} - p(c, \hat{b})\} = 0 \Rightarrow p(c, d) = d.$$

Applying the same reasoning to $v = 1$, we obtain

$$\max_{\hat{b}\leq d}\{p(d, \hat{b}) - \hat{b}\} = 0 \Rightarrow p(d, c) = c.$$

Because ex-post efficiency implies that $c < d$, we have reached a contradiction. ■

This proposition establishes that a large class of mechanisms can never hope to achieve ex-post efficiency. For instance, many anonymous mechanisms of the form $p(b_1, b_2) = f(b_1, b_2)b_1 + (1 - f(b_1, b_2))b_2$, where f is continuously differentiable and $0 < f(b_1, b_2) < 1$ for all b_1, b_2, will fail to be ex-post efficient. This class obviously includes linear mechanisms for which $0 < \lambda < 1$.

5.4 Ex-post Individual Rationality and Bayesian-Nash Equilibria

In this section we analyze the implications of requiring ex-post individual rationality in the traditional Bayesian-Nash profit-maximizing setting. As we saw in Section 4, the binary-search mechanism is ex-post individually rational in the case of equal entitlements and uniform [0,1] utilities. This result does not extend to general distributions. If entitlements are unequal, we prove that the only mechanisms that could satisfy this property are ones in which, like the binary-search mechanism, the price is a kind of a step-function of the bids.

Theorem 6 *Suppose the agents' objective is to maximize expected profit and that $r_1 \neq r_2$. Then if a mechanism $p(b_1, b_2)$ is ex-post individually rational, it must satisfy:*

$$\frac{\partial p}{\partial b_1} = \frac{\partial p}{\partial b_2} = 0 \ \textit{almost everywhere.}$$

Proof. Fix a mechanism p. Assume that agent 1 and 2's utilities are i.i.d. with a cdf F. Suppose agent 2 bids according to $h_2(v_2)$, and let $G_1(v_1, b_1)$ denote agent 1's expected profit, given a valuation of v_1 and a bid of b_1. We may write

$$\begin{aligned} G_1(v_1, b_1) &= \int_0^{h_2^{-1}(b_1)} r_2(v_1 - p(b_1, h_2(v_2)))f(v_2)dv_2 \\ &+ \int_{h_2^{-1}(b_1)}^1 r_1(p(b_1, h_2(v_2)) - v_1)f(v_2)dv_2. \end{aligned}$$

Differentiating the above with respect to b_1 and setting the derivative equal to 0, we obtain

$$\int_0^{h_2^{-1}(b_1)} -r_2(\frac{\partial}{\partial b_1}p(b_1,h_2(v_2))f(v_2)dv_2 \;+\; r_2(v_1-b_1)\frac{d[h_2^{-1}(b_1)]}{db_1}f(h_2^{-1}(b_1)) +$$

$$\int_{h_2^{-1}(b_1)}^1 r_1\frac{\partial}{\partial b_1}p(b_1,h_2(v_2))f(v_2)dv_2 \;-\; r_1(b_1-v_1)\frac{d[h_2^{-1}(b_1)]}{db_1}f(h_2^{-1}(b_1)) \;=\; 0.$$

Therefore in order for $h_1(v_1)=v_1$ and $h_2(v_2)=v_2$ to be an equilibrium, the following first-order conditions need to be satisfied:

$$\int_0^{v_1} -r_2\frac{\partial}{\partial b_1}p(v_1,v_2)f(v_2)dv_2 + \int_{v_1}^1 r_1\frac{\partial}{\partial b_1}p(v_1,v_2)f(v_2)dv_2 = 0,\;\; \forall v_1$$

$$\int_0^{v_2} -r_1\frac{\partial}{\partial b_2}p(v_1,v_2)f(v_1)dv_1 + \int_{v_2}^1 r_2\frac{\partial}{\partial b_2}p(v_1,v_2)f(v_1)dv_2 = 0,\;\; \forall v_2.$$

Now let $v_1=v_2=v$. By anonymity we have that $\frac{\partial}{\partial b_1}p(v_1,v_2)=\frac{\partial}{\partial b_2}p(v_2,v_1)$. Adding the two previous equalities gives us

$$\int_0^v \frac{\partial}{\partial b_1}p(v,u)f(u)du = \int_v^1 \frac{\partial}{\partial b_2}p(u,v)f(u)du,\;\; \forall v.$$

On the other hand, subtracting the same equalities, and assuming $r_1 \neq r_2$, we obtain

$$\int_0^v \frac{\partial}{\partial b_1}p(v,u)f(u)du = -\int_v^1 \frac{\partial}{\partial b_2}p(u,v)f(u)du,\;\; \forall v.$$

Combining the two equalities, we obtain

$$\int_0^v \frac{\partial}{\partial b_1}p(v,u)f(u)du = \int_v^1 \frac{\partial}{\partial b_2}p(u,v)f(u)du = 0,\;\; \forall v.$$

Since $f(v)>0$ and $\frac{\partial p}{b_1},\frac{\partial p}{b_2}\geq 0$, we conclude that

$$\frac{\partial p}{\partial b_1} = \frac{\partial p}{\partial b_2} = 0 \;\text{ almost everywhere.}$$

■

6 Conclusion

Standard buy-sell agreements are flawed by forcing the proposer to make herself indifferent to buying out, or being bought out, by the responder, who will generally not be indifferent and can choose her preferred option [11]. By contrast, we proposed a class of symmetric mechanisms that, by treating each partner equally, is fair to both.

We studied partnership dissolution models under the novel assumption that agents act as maximum-regret minimizers. We analyzed linear and binary search mechanisms and derived expressions for the regret-optimizing equilibrium strategies that they induce. When agents have equal entitlements to the partnership, we showed that linear mechanisms can achieve ex-post efficiency; the binary search mechanism satisfies the even stronger property of ex-post individual rationality.

Switching to a Bayesian-Nash framework, linear mechanisms remain ex-post efficient. The binary-search mechanism is ex-post individually rational for the case of uniformly distributed utilities, but this is not the case for generally distributed utilities.

When entitlements are unequal, linear mechanisms, as well as the binary-search procedure, do not satisfy these desirable properties. In fact, when entitlements are unequal and agents seek to minimize regret, there is no mechanism with a strictly increasing price function that is ex-post efficient. Likewise, when agents wish to maximize profit in such an asymmetric environment, the only mechanisms that can satisfy ex-post individual rationality are ones in which the price has a step-function structure, akin to the binary search mechanism.

The failure of several mechanisms to satisfy certain desirable properties is counterbalanced by the positive results we found in one important case—when the two partners have equal entitlements (i.e., 50 percent each). Under specified conditions, there is an ex-post efficient mechanism linear mechanism, and the binary mechanism is truth-inducing and, therefore, ex-post individually rational. These mechanisms deserve to be experimented with, if not actually used, in this important setting.

References

[1] D. Bergemann and K. Schlag (2005), "Robust Monopoly Pricing," Yale University, Cowles Foundation Discussion Paper, 1527.

[2] A. Borodin and R. El-Yaniv (1998), *Online Computation and Competitive Analysis,* Cambridge University Press.

[3] R.R. Brooks and K.E. Spier (2004), "Trigger Happy or Gun Shy? Dissolving Common-Value Partnerships with Texas Shootouts," School of Law, Northwestern University, Law and Economics Paper 19.

[4] K.Chatterjee and W. Samuelson (1983), "Bargaining under Incomplete Information," *Operations Research*, 31, 835–851.

[5] P. Cramton, R. Gibbons, and P. Klemperer (1987), "Dissolving a Partnership Efficiently," *Econometrica*, 55, 615–632.

[6] S. Eren and C. Maglaras (2006), "Pricing without Market Information," Graduate School of Business, Columbia University, working paper.

[7] M. Y. Gerchak and J. D. Fuller (1992), "Optimal-Value Declaration in 'Buy-Sell' Situations," *Management Science*, 38, 48–56.

[8] P. B. Linhart (2001), "Bargaining Solutions with Non-Standard Objectives," *Review of Economic Design*, 6, 225–239.

[9] P. B. Linhart and R. Radner (1989), "Minimax-Regret Strategies for Bargaining over Several Variables," *Journal of Economic Theory*, 48, 152–178.

[10] R. P. McAfee (1991), "Amicable Divorce: Dissolving a Partnership with Simple Mechanisms," *Journal of Economic Theory*, 56, 266–293.

[11] J. Morgan (n.d.), "Dissolving a Partnership (Un)Fairly," Hass School of Business and Department of Economics, University of California, Berkeley, working paper.

[12] R. Myerson and M. Satterthwaite (1983), "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265–281.

[13] G. Perakis and G. Roels (2008), "Regret in the Newsvendor Model with Partial Information," *Operations Research*, 56, 188–203.

[14] L.J. Savage (1954), "The Foundations of Statistics", John Wiley and Sons, New York.

[15] J. L. Turner (2006), "Dissolving (In)effective Partnerships," Department of Economics, University of Georgia, working paper.