Five Essays in Economic Theory

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
Nora Szech
aus Bremen

2010

Dekan:	Prof. Dr. Christian Hillgruber
Erstreferent:	Prof. Dr. Benny Moldovanu
Zweitreferent:	JProf. Dr. Anja Schöttner

Tag der mündlichen Prüfung: 14.09.2010

Diese Dissertation ist auf dem Hochschulschriftenserver der ULB Bonn (http://hss.ulb.uni-bonn.de/diss_online) elektronisch publiziert.

Acknowledgments

Many people supported me in writing this thesis, and I am grateful to all of them.

First I want to thank my supervisor Benny Moldovanu for all his support and inspiration. I learned very much from his comments and critique, not only on my own work.

Anja Schöttner, the second referee of my thesis, influenced with her topics course especially the first two chapters. Her advice helped me very much.

Another big thank goes to Patrick Schmitz, who advised me on almost all of my projects and gave me a lot of valuable comments and critique.

From Avner Shaked I learned a lot about game theory and its place in the world. Rani Spiegler whose work inspired Chapter 2 supported me very much. Philipp Weinschenk is the coauthor of the paper presented in Chapter 3. I enjoyed very much to work with him.

I would also like to thank my office mates for the discussions we had and for the good time together: Wei Ding, Ferdinand Gider, Jasmin Gider, Thomas Rieck and Christian Westheide.

Many discussions and advice from others helped me to improve my work as well. I especially want to thank Konstanze Albrecht, Marc Armstrong, Olivier Bos, Gary Charness, Vincent Crawford, Christoph Engel, Emma von Essen, Armin Falk, Thomas Gall, Alex Gershkov, Susanne Goldlücke, Paul Heidhues, Martin Hellwig, Fabian Herweg, Jos Jansen, Philippe Jéhiel, Wilbert Kallenberg, Frederik Kaster, Kai Konrad, Sebastian Kranz, Lutz Mattner, Alexander Morell, Johannes Münster, Petra Nieken, Juliane Parys, Dan Silverman and Michael Wegener. Of course I also want to thank the many others that supported me and gave helpful comments on my projects.

The Bonn Graduate School of Economics is a great place to learn and to develop and discuss ideas. I would like to thank Urs Schweizer for his enduring effort in the school and all the others behind. I also want to thank Uwe Sunde for encouraging me to take up graduate studies, as well as Armin Falk, Ingo Lieb, Benny Moldovanu, Frank Riedel and Klaus Sandmann for their advice and encouragement.

I am deeply indebted to my family and to my friends. This thesis is dedicated to them, especially to my parents.

Contents

List of Figures

Introduction

This thesis consists of five chapters which are linked in several ways: All chapters contribute to auction theory. Yet their applications vary and cover topics such as contests, industrial organization, behavioral game theory, and health economics. From a theoretical point of view, the first three chapters form an entity as they deal with the structure of mixed strategy equilibria in asymmetric auction-type games. Chapters 4 and 5 are concerned with costly information release in standard auctions. Another link can be seen between Chapters 1, 2, and 4: They cover, respectively, all-pay, first-price, and second-price auctions in which bidders participate with asymmetric probabilities. Nevertheless the theoretical overlap between the three chapters is small - and the questions analyzed and the applications considered differ strongly. Hence this thesis can be seen as an example of how broadly applicable auction theory is.

The first chapter analyzes asymmetric all-pay auctions with two types. All-pay auctions are frequently used models of R&D, contests, lobbying, rent-seeking, and many other situations of competition. So far, the complete information case is well understood, and the same is true for the symmetric incomplete information case. Yet the asymmetric incomplete information case still poses difficulties and is generally quite complex to study. Chapter 1 provides a thorough analysis of a simple yet truly asymmetric all-pay auction model. For the two bidder case, the asymmetries can lie in the bidders' respective high and low types as well as in the type-probabilities. With n bidders, we study asymmetries in probabilities of participation. For both settings we characterize the explicit form of the unique equilibrium.

We believe these results are not only a theoretical contribution, but are also useful for various applications: So far, applications where asymmetries should be involved typically assume complete information on the bidders' side as this assumption makes those models tractable. Yet it seems rather unrealistic for most real-world situations that the bidders are that well informed about each other. Likewise, for many applications, it seems too restrictive to assume that the bidders know each other so little that they hold a symmetric common prior about each others' strengths. Hence only a very partial picture can be obtained without models that allow for a coexistence of asymmetry and informational uncertainty.

To give a flavor of the applications we have in mind, we thoroughly analyze two settings to answer questions which were so far difficult to study: First, we consider a game of information release, where a bidder can release information about his valuation to the other bidder. We find that whenever the release of strictly intermediate amounts of information is possible, a bidder wants to provide additional information to his competitor. This result stands in contrast to the results found in Kovenock, Morath and Münster (2009), where the bidders can only reveal their valuations completely or not at all. In their setting, revelation is not attractive for the bidders. Our analysis shows that this result indeed hinges on the assumption of all-or-no information.

The second application we analyze in more detail is competition between contests: There are two all-pay contests differing in the prizes for the winning bidder. Bidders decide which contest to attend. We find that the fractions of bidders active in the respective contests do not depend on the prize size if there are many bidders. Furthermore, for any fixed number of bidders, adding a second contest leads to a loss in revenues for the seller compared to holding a single contest if the prize of the second contest is small.

We then move on to a setting with three contests - a global contest where all bidders can participate and two local ones. The bidders are heterogeneous: Each bidder can only participate in "his" local contest. Every bidder decides between taking part in the global contest or in his local contest. A concrete example for such a setting would be a general interest journal and two field journals: Only part of the bidders can submit to each field journal, but all bidders can submit to the general interest journal. We characterize the limit distribution of bidders across contests when there are many bidders. This distribution only hinges on the relation between the group sizes, not on the prize differences.

As a third application, we briefly discuss the following two-stage model: In the first stage contestants costlessly choose their probabilities of having low effort costs in the all-pay contest which takes place in the second stage. Contestants opt for surprisingly low probabilities of having low costs in order to avoid strong competition in the contest stage. All these results can be obtained because the theoretical part of the chapter provides a tractable model of incomplete information asymmetric all-pay auctions.

Chapter 2 deals with a market model in which the consumers are unfamiliar with the market and rely on word of mouth to get an impression about the qualities of the different firms. The firms know the market well and know the distributions from which the qualities of their competitors are drawn. They engage in price competition.

Assuming that consumers pick the cheapest firm among all firms about which they heard a positive story, we find that stronger competition may lower welfare. In the equilibrium where higher quality firms set higher prices, welfare even goes to zero in the number of firms. Welfare may first increase in the number of firms, but typically starts to decrease substantially as soon as a small threshold-number of firms is exceeded.

We furthermore characterize all equilibria of the game. As a by-product, we prove equilibrium uniqueness of the special case when firms hold complete information about each others' qualities. This setting was analyzed by Ireland (1993) and McAfee (1994) in the context of advertising. Both pointed out the question of equilibrium uniqueness as an open problem.[1]

[1]For such advertising models, the natural definition of welfare is very different. Notably, in the

The inspiration to interpret this game in terms of boundedly rational consumers and rational firms comes from Spiegler (2006a) who mostly focuses on the symmetric complete information case. In his model, competition can be detrimental to consumer surplus, but only up to a critical number of firms. With more firms than the critical number, competition leads to higher consumer surplus. Hence our effect that strong competition leads to very negative results concerning consumer surplus and welfare is novel: It arises due to the incomplete knowledge of the firms about the quality realizations of each other.

That people rely on anecdotes has been shown by a broad literature.[2] Even if statistical evidence is provided, untrained persons typically prefer to trust in anecdotal evidence, like personal stories they hear from their family, friends or colleagues. In the medical literature, it is even discussed whether statistical information about the success of different therapies should be enriched by personal stories to make patients pay more attention to the statistical results.[3] One reason why anecdotal evidence is that compelling is seen in its dichotomy:[4] A personal story would typically result in something like "this therapy was bad" or "that pill was good". Statistical evidence is much more differentiated and hence much more difficult to grasp. Our model captures the dichotomic nature of anecdotes.

Applied to the US health market, our model gives an explanation for the recent surprising result that the Veterans Health Administration (VHA) outperforms the other insurance systems prevalent in the US with regard to a large variety of factors:[5] Unlike most other systems in the US, the VHA is non-competitive. VHA patients typically stay inside the institution for the rest of their lives. Other patients, like those of Health Maintenance Organizations (HMOs), usually switch their medical plans on a regular basis. Hence competition plays an important role for these insurers. Comparing health plans with regard to quality is tedious: A patient would

interpretation of advertising, competition does not have any detrimental welfare effects.

[2] Kahneman and Tversky (1973) is seminal.

[3] See Enkin and Jadad (1998).

[4] Compare Fagerlin et al. (2005).

[5] Compare Brooks (2008) and Longman (2010).

have to study coverage of myriads of different health problems in detail. In contrast, prices are easy to grasp and serve very likely as an important decision-tool in practice. In such a setting, if consumers take coverage of a small sample of conditions as representative for an insurance plan's quality or if they rely on word of mouth when choosing their health plans, our model predicts that patients often end up with low quality health plans.

Chapter 3 is joint work with Philipp Weinschenk. We extend the classical Bertrand Game by assuming that firms compete over a heterogeneous base of consumers: Part of the consumers get a rebate at one of the firms, while the remaining group receives a rebate at the other firm. Rebates and group-sizes can be asymmetric. We provide an explicit characterization of the mixed pricing strategies firms employ in equilibrium. The supports of these strategies are made up by two parts: one "aggressive" subinterval where a firm plays prices that could be attractive for the home-base consumers of the other firm as well and one "defensive" part where the firm plays high prices that are only potentially attractive for its own home-base of customers who receive the rebate at the firm. Furthermore, we show that firms have a general incentive to make the consumer base heterogeneous: With rebates, firms earn positive payoffs, in contrast to the standard Bertrand result that firms earn no payoffs in equilibrium. Viewed as an auction, this game is a two-bidder complete information first-price auction where the bidders can control their bids only up to some binary, asymmetric noise.

Chapters 4 and 5 deal with the same question in two different settings: A seller wants to sell an object to bidders who initially only know the distribution from which their valuation is independently drawn. At costs that increase with the amount of transmitted information, the seller can make the bidders learn their valuations better. Yet by providing information, the seller does not learn how the information affects the valuation realizations - he only knows that releasing more information makes bidders know their valuations better.

We study the question whether the seller should concentrate his informational efforts on few bidders or whether he should divide his efforts among many bidders to maximize his revenues. The main difference between the chapters lies in how the process of information transmission is modeled.

In Chapter 4, we assume that the seller controls the probabilities with which the bidders learn their valuations in a second-price auction. Bidders who do not learn their valuation realizations either bid their ex-ante estimate or do not take part in the auction. The first case is thought to capture the situation of information release, whereas the second case applies to the situation of advertising an auction.

Our first main result is that if costs are quasi-concave in the information transmission probabilities, under mild conditions the revenue or welfare-maximizing seller concentrates a total amount of informational efforts on as few bidders as possible. Hence he maximally unlevels the playing field.

One mild but crucial condition for the optimality of asymmetric allocations is that $f/(1-F)^2$ is increasing. F denotes the distribution from which the bidders' valuations v_i are drawn and f is its density. The condition ensures that the sequence of expected second order statistics $E[v_{2:n}]$ is strictly concave in n. It reminds of the increasing failure rate condition from reliability theory.[6]

Furthermore, we show that the revenue-maximizing seller invests more into advertising than socially optimal if the distribution of the bidders' valuations has increasing failure rate. Under a decreasing failure rate, the opposite is observed. Both, the revenue- and the welfare-optimal levels of advertising, are larger the more dispersed the distribution of valuations is in terms of the excess wealth order.

Chapter 5 deals with the same question as Chapter 4 in a different setting. The bidders' valuations are given by sums of independent identically distributed random variables, so-called information packages. The seller decides how many information packages to reveal to each bidder at a cost which depends on the total number of information packages released. We mostly focus on the case of two bidders in the

[6]See, e.g., Barlow and Proschan (1981).

situation where the seller can sell his information and can thus extract all surplus through individual entry fees. Essentially, we show that giving the same number of packages to both bidders is dominated by any other choice of dividing the same number of packages. The main results of the chapter hinge on the following little-known, elementary but non-trivial inequality:[7] For two independent, identically distributed random variables X and Y it holds that

$$E[|X+Y|] \geq E[|X-Y|].$$

The inequality becomes an equality iff the distribution of X and Y is symmetric around its mean.

[7]See Jagers, Kallenberg and Kroese (1995).

Chapter 1

Asymmetric All-Pay Auctions with Two Types

We characterize the equilibria of asymmetric all-pay auctions with incomplete information. The bidders' types are independently drawn from different two-point probability distributions: Types as well as probabilities differ among bidders. The realized types are private information.

We characterize the unique equilibrium of the two bidder case. For the n bidder case, we consider the restriction in which bidders participate in the auction with different probabilities. Our explicit formula for the unique equilibrium payoffs transfers to auctions with multiple prizes and to asymmetric cost functions which fulfill a local symmetry condition.

We apply our results to models of information disclosure in contests, endogenous choice of type-probabilities, and competing contests. With information disclosure, bidders always want to disclose some information.

1.1 Introduction

We analyze asymmetric all-pay auctions where each bidder has one of two possible types, characterized by either low or high costs of exerting effort. Bidders know their own type, but they only know the probabilities of types for the other bidders. The asymmetry between the bidders may lie in the respective effort costs of the two types as well as in the type-probabilities. We are going to characterize the equilibria of this auction.

All-pay auctions are frequently used models of contests and related situations of competition.[1] Complete information all-pay auctions are by now well-understood.[2] The same is true for symmetric all-pay auctions with incomplete information.[3] One case which has posed relative difficulties so far is the asymmetric case with incomplete information. In this case, bidders are aware of some differences in valuations or costs between each other but do not possess complete information. For the case of continuously distributed valuations, Amann and Leininger (1996) analyze the two bidder case. They show existence and uniqueness of the equilibrium and characterize some properties of the equilibrium, e.g., one participant bidding zero with positive probability. Parreiras and Rubinchik (2010) generalize the setting of Amann and Leininger to n bidders with asymmetric risk-attitudes. They characterize classes of examples where the equilibrium has properties that have been found in experimental studies, such as non-decreasing densities of bids and complete drop-out of some bidders. Konrad and Kovenock (2010) also consider two-bidder all-pay auctions where marginal costs are drawn from different distribution functions. Yet in their model all private information is revealed before the auction takes place. Thus the bidding behavior is covered by the complete information case.

We complement these works by considering asymmetric incomplete information all-pay auctions in which each bidder has one of two possible types. The assumption of two types allows us to analyze the incomplete information case while retaining some of the tractability and some techniques from the complete information setting. We provide a complete analysis of the two bidder case in which the bidders' types are drawn from different two-point distributions. Additionally we characterize the unique equilibrium of the n bidder case where bidders participate in the auction with asymmetric probabilities. We show that payoff-uniqueness transfers to auctions with

[1]Konrad (2009) provides an overview with many references.

[2]See for instance Hillman and Samet (1987), Hillman and Riley (1989), Baye, Kovenock and de Vries (1996), Clark and Riis (1998), and Siegel (2009, 2010).

[3]See for example Moldovanu and Sela (2006).

multiple prizes and to auctions with asymmetric cost functions which fulfill a local symmetry condition. This local condition follows the spirit of Siegel (2009): Adapting Siegel's terminology, we require that bidders are identical in "reach" given they are active. Yet note that in our setting, equilibrium considerations are necessary to determine the reach of each bidder as it depends on the other bidders' participation probabilities in a non-trivial way.

Our results for the two-bidder case also complement those of Maskin and Riley (1985) who characterize equilibria of first- and second-price winner-pay auctions for the two-bidder case with asymmetric two-point distributions. The n bidder auction we analyze in Section 1.4 can be interpreted as a contest with stochastic participation. There has been some recent research about contests with stochastic numbers of participants.[4] Yet to our knowledge our study is the first to provide a discussion of an asymmetric model of this type.

The two-types approach is a bit restrictive, but it has some clear advantages: The equilibria we find are completely explicit. This gives us a simple tool for studying comparative statics of the equilibrium with regard to the parameters of the distribution. The general two-point distributions we consider in the two-bidder case are rich enough to capture many different types of asymmetries. This is also reflected in the rich structure of the equilibrium.

In the final part of the chapter we consider another way to exploit our results which goes beyond studying the effects of asymmetries: We discuss situations where our explicit results admit an easy analysis of asymmetric all-pay auctions with incomplete information nested in richer settings. Concretely, we consider three models where asymmetries occur endogenously as the result of a preliminary stage of the game: We first study contests where bidders have the opportunity to release *some* information about their type. We find that in our two-types setting it is always profitable to release information. This stands in stark contrast to recent negative results about information-sharing in contests by Kovenock, Morath and Münster

[4] See Münster (2006), Myerson and Wärneryd (2006), and Lim and Matros (2009).

(2009) who focus on the case where contestants can either reveal all information or none. As a second application, we consider an all-pay auction where bidders freely choose their probability of being of strong type in a preliminary stage. A concrete example could be competing lobbyists who control chances of having good access to the decision-maker. We show that - at first sight surprisingly - bidders prefer to set not too high access probabilities. As a third application, we adapt our results about contests with random participations to a model of competing contests: Bidders simultaneously choose how much effort to exert and in which out of two contests to participate. We show that adding a second, competing contest with a smaller prize to an existing contest always reduces aggregate efforts with two or three bidders. With more than three bidders an additional contest with a small prize is detrimental while sufficiently large prizes lead to an increase in aggregate efforts. Finally, we characterize the distribution of large numbers of bidders across contests in a generalized setting: Bidders belong to one of two groups and can either submit to a group-intern contest or to a contest which is open to all bidders. (An example could be researchers deciding between submitting to a field journal or to a general-interest journal.) We find that the asymptotic distribution depends only on proportions between group sizes but not on asymmetries in prize money across contests.

Our road-map is as follows: In Section 1.2, we introduce the model and state some elementary observations. Section 1.3 characterizes the unique equilibrium of the two bidder case. Section 1.4 provides the equilibrium of the n bidder case under the restriction that bidders only differ in their type-probabilities p_i and have prohibitively high costs as weak types. In Section 1.5, we study three applications of our analysis, information disclosure in contests, endogenous choice of type-probabilities, and competing contests with different prizes. Section 1.6 concludes.

1.2 The Model

There is an object for sale in an all-pay auction with n bidders, who each have a valuation of 1 for the object. With probability $p_i \in (0, 1)$, bidder i is of strong type and has low marginal costs of exerting effort, c_i.[5] With the counter probability, he is of weak type and has high marginal costs of C_i. We assume $c_i < C_i \leq \infty$. The probability distributions are common knowledge. Each bidder knows his own type but not the type-realizations of his opponents. The bidder who exerts the highest effort e wins the object. Ties are broken arbitrarily. The symmetric two bidder case of this model has been analyzed in Konrad (2004). Münster (2009) studies the case $c_i = c$ and $C_i = \infty$.

Standard arguments show that any equilibrium of this game must be in mixed strategies and that there are no atoms except possibly in zero just as in the complete information auction. Thus each bidder's strategy can be represented by two distribution functions which are atomless (with zero as possible exception): Bidder i utilizes F_i^c if he has low marginal costs c_i and F_i^C if he has high marginal costs C_i.

For fixed i, the supports of F_i^c and F_i^C must be disjoint (except possibly for boundary points and zero): Let $P_i(e)$ denote bidder i's probability of winning via an effort of e (given the other bidders' strategies). Postulate for convenience that $P_i(e)$ is differentiable. Note that $P_i(e)$ does not depend on i's type. Thus bidder i maximizes either

$$P_i(e) - c_i e \qquad \text{or} \qquad P_i(e) - C_i e.$$

At no e both maxima can be attained (except in zero when the first summands are zero), as the respective first order conditions are $P_i'(e) = c_i$ and $P_i'(e) = C_i$.

Taking this argument one step further, we see that the strong type's payoff from exerting effort in the weak type's interval must be increasing in effort: The weak type earns constant expected payoffs on his interval. Since the strong type has lower marginal costs, increasing the effort must then be profitable for the strong type on these effort levels. Likewise, the weak type's payoff must be decreasing in effort on the strong type's interval. This implies that the strong type of a bidder must play strictly higher effort levels than the weak type in equilibrium.

[5] We assume that the fixed costs equal zero.

1.3 The Two Bidder Case

In this section we analyze the case of two bidders, $n = 2$. We also assume $C_i < \infty$ for both i.[6]

Depending on the values of the parameters c_1, C_1, p_1 and c_2, C_2 and p_2 we have to distinguish six different cases defined by the following conditions:

$$p_1 c_1 > p_2 c_2, \tag{A1}$$

$$p_1 c_1 > p_2 c_2 + (1 - p_2) C_2, \tag{A2}$$

$$p_1 c_1 + (1 - p_1) C_1 > p_2 c_2 + (1 - p_2) C_2. \tag{A3}$$

Before we come to the intuition for these conditions, let us start with some general properties of the equilibrium. Most of these follow from arguments known from the complete information case:[7]

- In equilibrium, bidders do not leave gaps in their equilibrium supports.
- No bidder sets an atom on strictly positive effort-levels.

Thus bidders will play strictly mixed strategies if they choose strictly positive effort-levels. Furthermore we know from Section 1.2:

- At a given effort level, a bidder competes either against the strong type of the other bidder, or against the weak type. At no effort level, a bidder competes against both types of the other bidder. The strong type of a bidder plays higher effort-levels than the weak type.

[6] The slightly simpler case in which one or both C_i are infinitely large is omitted to keep the exposition more focused.

[7] See e.g. Siegel (2009).

- No bidder will exert effort levels higher than the highest effort level of the other bidder.

- Both bidders (at least their weak types) have to mix down to effort level zero.

Thus a bidder i mixes in equilibrium on a gapless support. His strong type will mix up to $\overline{e}$ and down to some lower boundary, say e_i. His weak type plays all the effort-levels between e_i and zero. For further details, see the proof of Proposition 1.1 in the Appendix.

Let us first consider the effort interval on which the strong types of both bidders compete against each other. Which strong type has to play the more concentrated strategy, i.e. has to mix with a higher density? If (A1) is fulfilled,

$$p_1 c_1 > p_2 c_2, \tag{A1}$$

it must be the strong type of bidder 1 (which we call "strong 1" from now on).

Let us see why: Consider the strong type of bidder 2, strong 2. Increasing his effort slightly by ϵ inside the equilibrium support must not affect his payoffs. An effort increase can only pay out if strong 1 is active, which happens with probability p_1. (Winning over weak 1, the weak type of bidder 1, is certain.) Hence it has to hold:

$$p_1(F_1^c(e_2 + \epsilon) - F_1^c(e_2)) = c_2 \epsilon, \tag{1.1}$$

where the left hand side denotes the expected additional gain by increasing the effort by ϵ. The right hand side denotes the additional cost. Taking ϵ to zero we also obtain the density strong 1 has to play on the interval where strong 2 is active. It must be given by $f_1^c = \frac{c_2}{p_1}$. Yet then, $\frac{p_1}{c_2}$ denotes the length of the interval strong 1 would have to mix on if he always played against strong 2.

Let us compare the lengths of the intervals strong 1 and strong 2 would have to play on if they always played against the strong type of the other bidder:

$$\frac{p_1}{c_2} > \frac{p_2}{c_1}. \tag{A1'}$$

If (A1') holds, as both strong types mix up to the same upper boundary $\overline{e}$ (whose value we still have to determine), strong 1 has to mix down to lower effort levels than strong 2. So we see that strong 2 indeed always competes against the strong type of his competitor, in contrast to strong 1, who is left with some probability mass he has to "spend" elsewhere. Note that (A1') and (A1) are equivalent.

Let us for the rest of the section without loss of generality assume that (A1) is fulfilled. Hence we know that strong 1 has to mix over a larger interval than strong 2. We know he cannot mix over higher effort levels than strong 2. Hence he has to mix down to lower effort levels than strong 2. Will he even mix down to effort level zero? This depends on (A2):

$$p_1c_1 > p_2c_2 + (1-p_2)C_2, \tag{A2}$$

If (A2) holds, strong 1 even has to mix over a larger interval than both strong 2 and weak 2 together. Note that (A2) resembles (A1) very much: The left hand side is identical in both conditions, as in both situations it is always the strong type of bidder 1 who is active. The left hand side of (A2) is analogous to the left hand side of (A1) if we reinterpret bidder 2 as a bidder $2'$ who is always of strong type (p=1), but who has costs of $p_2c_2 + (1-p_2)C_2$. Then (A2) reads:

$$p_1c_1 > 1 \cdot (p_2c_2 + (1-p_2)C_2), \tag{A2'}$$

Note that such a reinterpretation is valid here because bidder 2 always competes against strong 1, and never against weak 1 under (A2). Let us consider the corresponding picture of the shape of the equilibrium:

Strong 1 has to mix over such a long interval that there is no room for weak 1 to mix over any positive effort levels. Hence weak 1 has to put all his probability-mass on zero.

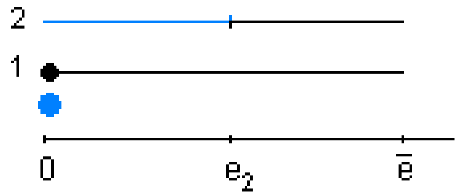

Figure 1.1: Supports of the bidders' strategies if (A1) and (A2) hold, strong types in black, weak types in blue.

Our third condition, Condition (A3),

$$p_1 c_1 + (1 - p_1)C_1 > p_2 c_2 + (1 - p_2)C_2, \tag{A3}$$

is relevant only if (A2) is not fulfilled, i.e. in the case where the weak types of both bidders exert positive efforts with some probability. Then it depends on (A3) which of the weak types puts an atom on zero. If (A3) is fulfilled, weak 1 has to play the atom in zero (and consequently earn zero profits):

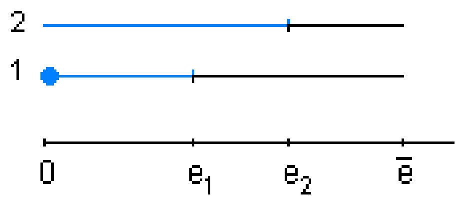

Figure 1.2: Supports of the bidders' strategies if (A1) holds, (A2) does not hold, but (A3) holds. Strong types in black, weak types in blue.

Conversely, if (A3) is violated, the weak type of bidder 2 puts an atom on zero and makes no profits.

Proposition 1.1 formally characterizes the unique equilibria for the three cases in which (A1) holds. The cases in which (A1) does not hold can of course be extracted from Proposition 1.1 by exchanging the roles of the indices 1 and 2. As one can see from the proposition, it is quite lengthy to state the equilibria explicitly. Yet note

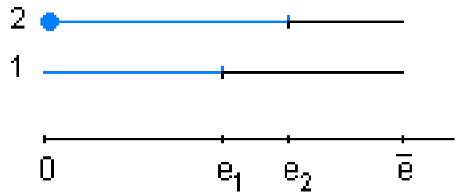

Figure 1.3: Supports of the bidders' strategies if (A1) holds, but (A2) and (A3) do not hold. Strong types in black, weak types in blue.

that the bidders just mix uniformly over the intervals specified before, with densities such that the opponent's active type would not gain or lose from marginally changing his effort level.

Proposition 1.1. *Consider the two bidder case and assume that (A1) holds. Then we have to distinguish three cases:*

1. *Assume (A2) holds. Define boundaries e_1, e_2, and $\bar{e}$ by $e_1 = 0$,*
$$e_2 = \frac{(1-p_2)C_2}{c_1 C_2} = \frac{1-p_2}{c_1}, \qquad \bar{e} = e_2 + \frac{p_2 c_2}{c_1 c_2} = \frac{1}{c_1}.$$
In the unique equilibrium, weak 1 places an atom of size 1 on 0. Strong 1 mixes over $(e_1, e_2]$ with constant density $\frac{C_2}{p_1}$ and over $(e_2, \bar{e}]$ with density $\frac{c_2}{p_1}$. Additionally, strong 1 places an atom of size
$$\frac{p_1 c_1 - p_2 c_2 - (1-p_2)C_2}{p_1 c_1}$$
on e_1. Weak 2 mixes over $(e_1, e_2]$ with density $\frac{c_1}{(1-p_2)}$. Strong 2 mixes over $(e_2, \bar{e}]$ with density $\frac{c_1}{p_2}$.

2. *Assume (A2) does not hold but (A3) does. Define boundaries e_1, e_2, and $\bar{e}$ by*
$$e_1 = \frac{p_2 c_2 + (1-p_2)C_2 - p_1 c_1}{C_1 C_2}, \qquad e_2 = e_1 + \frac{p_1 c_1 - p_2 c_2}{c_1 C_2}, \qquad \bar{e} = e_2 + \frac{p_2 c_2}{c_1 c_2}.$$
Then the unique equilibrium is given by the following strategies: Weak 1 mixes over $(0, e_1]$ with density $\frac{C_2}{1-p_1}$ and places an atom of size
$$\frac{p_1 c_1 + (1-p_1)C_1 - p_2 c_2 - (1-p_2)C_2}{(1-p_1)C_1}$$
on 0. Strong 1 mixes over $(e_1, e_2]$ with density $\frac{C_2}{p_1}$ and over $(e_2, \bar{e}]$ with density $\frac{c_2}{p_1}$. Weak 2 mixes over $(0, e_1]$ with density $\frac{C_1}{1-p_2}$ and over $(e_1, e_2]$ with density $\frac{c_1}{1-p_2}$. Strong 2 mixes over $(e_2, \bar{e}]$ with density $\frac{c_1}{p_2}$.

3. *Assume (A2) and (A3) both do not hold. Define boundaries* e_1, e_2*, and* $\overline{e}$ *by*

$$e_1 = \frac{(1-p_1)C_1}{C_1C_2}, \qquad e_2 = e_1 + \frac{p_1c_1 - p_2c_2}{c_1C_2}, \qquad \overline{e} = e_2 + \frac{p_2c_2}{c_1c_2}.$$

Then the unique equilibrium is given by the following strategies: Weak 1 *mixes over* $(0, e_1]$ *with density* $\frac{C_2}{1-p_1}$*. Strong* 1 *mixes over* $(e_1, e_2]$ *with density* $\frac{C_2}{p_1}$ *and over* $(e_2, \overline{e}]$ *with density* $\frac{c_2}{p_1}$*. Weak* 2 *mixes over* $(0, e_1]$ *with density* $\frac{C_1}{1-p_2}$ *and over* $(e_1, e_2]$ *with density* $\frac{c_1}{1-p_2}$*. Additionally, weak* 2 *places an atom of size*

$$\frac{p_2c_2 + (1-p_2)C_2 - p_1c_1 - (1-p_1)C_1}{(1-p_2)C_2}$$

on 0*. Strong* 2 *mixes over* $(e_2, \overline{e}]$ *with density* $\frac{c_1}{p_2}$*.*

From Proposition 1.1 it is easy to calculate the expected equilibrium payoffs:

Corollary 1.1. *In the setting of Proposition 1.1, the payoff of strong* i *equals* $1 - c_i\overline{e}$*. The payoff of weak* i *equals* A_{-i}*, where* A_{-i} *is the probability that* i*'s opponent exerts an effort of zero.*

Just like bidders in a complete information all-pay auction, the strong types earn the same expected payoffs if $c_1 = c_2$. Yet the same is not true for weak type bidders. Even if $C_1 = C_2$ their expected payoffs will generally differ: Due to the atom, one of them earns a positive expected payoff while the other obtains zero payoff.

To gain some more intuition, and since we work with this case in the first two applications of Section 1.5, we finish our analysis of the two bidder case with a closer look at the situation where asymmetries lie only in the probabilities, i.e. $c_1 = c_2 = c$ and $C_1 = C_2 = C$. Then, assumption (A1), i.e. $p_1 > p_2$, immediately implies that (A2) and (A3) must be violated. Hence we are then always in the third case of Proposition 1.1. Then we get the following simplified formulas for the payoffs:

Corollary 1.2. *Assume that in the setting of Proposition 1.1 it holds that* $c_1 = c_2 = c$ *and* $C_1 = C_2 = C$*. Then the atom of the opponent is given by*

$$A_{-i} = (p_i - \min(p_1, p_2))(1 - \frac{c}{C}).$$

The upper bound of supports $\overline{e}$ *is given by*

$$\overline{e} = \frac{\min(p_1, p_2)}{c} + \frac{1 - \min(p_1, p_2)}{C}.$$

Accordingly, the expected payoff of weak i is given by

$$\pi_i^w = (p_i - \min(p_1, p_2)) \left(1 - \frac{c}{C}\right).$$

The expected payoff of strong i

$$\pi_i^s = (1 - \min(p_1, p_2)) \left(1 - \frac{c}{C}\right).$$

Note that the corollary is written in a way that it holds regardless of whether (A1) is fulfilled or not.

1.4 All-Pay Auctions with Random Participation

We now turn to the case of n bidders who have a valuation of 1 for winning. We restrict attention to the case where bidders differ only in their probabilities p_i of being strong, i.e. $c_i = c$ and $C_i = C$, and where weak types have infinitely high marginal effort costs $C = \infty$. The most natural interpretation of these prohibitively high marginal costs is that weak types are not aware of the auction taking place or are unable to show up at the auction.[8] Hence we analyze an all-pay auction with asymmetric random participation. When exerting their efforts, bidders are uncertain who else will exert an effort. We assume without loss of generality that bidders are ranked according to their participation probabilities (i.e. their probabilities of being of strong type), $p_1 \geq \ldots \geq p_n$.[9]

We assume that if a weak type and a strong type exert the same effort then the strong type wins (and other ties are broken with equal probability). If we think of p_i as a participation probability, this requirement is quite natural: A bidder who is

[8] A reader who is bothered about the idea of infinite costs can - without any changes in what follows - assume that all types have marginal costs c, that strong types have a valuation of 1 and weak types have a valuation of 0. The present formulation is chosen for reasons of consistency with the previous section.

[9] Recall that we assumed $p_i \in (0, 1)$. The analysis of this section still goes through if $1 = p_1 > p_2$. With $1 = p_1 = p_2$, the game essentially degenerates to a complete information all-pay auction. Bidders with $p_i = 0$ can be viewed as non-existent.

not present rather submits no bid at all than a zero bid.[10]

Each strong type bidder i earns a strictly positive maxmin payoff (after deletion of strictly dominated strategies) of $\prod_{j\neq i}(1-p_j)$ and thus a strictly positive payoff in every equilibrium. All strong types must obtain the same payoffs in equilibrium: Assume strong j obtains a payoff of π_j and the highest effort level in the support of his strategy is $\overline{e}_j$. Then by exerting an effort level slightly above $\overline{e}_j$ any other strong type bidder i can attain a payoff slightly below π_j. Taking these observations together it follows that all strong types must earn a payoff of at least $\prod_{j\neq 1}(1-p_j)$, which is the maxmin payoff of the strong type of the ex-ante "most present" bidder 1. Furthermore, the union of the bidders' supports must go down to zero and at most one bidder puts an atom on zero. Thus at least one bidder does not earn more than his maxmin payoff. Hence we see that equilibrium payoffs of strong types must equal $\prod_{j\neq 1}(1-p_j)$ for all bidders (while weak types obtain, of course, zero payoffs). We collect this result of payoff-uniqueness together with a characterization of the sum of efforts in the following lemma:

Lemma 1.1. *In any equilibrium, all strong types expect a payoff of*

$$\pi_i^s = \prod_{j\neq 1}(1-p_j)$$

while weak types expect a payoff of 0*. Denote by* S *the (ex-ante) expected sum of the bidders' efforts. Then we have*

$$S = \frac{1}{c}\left[\left(1-\prod_{j=1}^{n}(1-p_j)\right) - \prod_{i\neq 1}(1-p_i)\sum_{k=1}^{n}p_k\right]. \tag{1.2}$$

The formula for S simply follows from the observation that the expected value of prizes awarded to strong types[11] minus the expected sum of efforts must equal the sum of the bidders' ex-ante expected payoffs:

$$\left(1-\prod_{j=1}^{n}(1-p_j)\right) - c\,S = \sum_{k=1}^{n}p_k\prod_{i\neq 1}(1-p_i)$$

[10]Proposition 1.3 below discusses the connection between tie-breaking rules and equilibrium existence in our model in detail.

[11]Note that weak types never exert effort in equilibrium. Furthermore, a weak type never gets the prize when a strong type is present in the auction. Thus prizes awarded to weak types can be treated as zero for these considerations.

An equilibrium of this auction is explicitly characterized in the following proposition:

Proposition 1.2. *Define an increasing sequence e_i of boundaries by $e_1 = e_2 = 0$,*

$$e_k = \frac{1}{c}\left[(1-p_k)^{k-1}\prod_{l=k+1}^{n}(1-p_l) - \prod_{l=2}^{n}(1-p_l)\right]$$

for $k = 3, \ldots, n$, and $e_{n+1} = \frac{1}{c}\left[1 - \prod_{l=2}^{n}(1-p_l)\right]$. Then it is a Nash equilibrium that weak i exerts effort level zero and strong i mixes according to the distribution function F_i. The support of F_i is the interval $S_i = [e_i, e_{n+1}]$. F_i is defined piecewise as follows: For $e \in [e_k, e_{k+1}] \subseteq S_i$, $2 \le k \le n$,

$$F_i(e) = \frac{1}{p_i}\left(\sqrt[k-1]{\frac{ce + \prod_{l=2}^{n}(1-p_l)}{\prod_{l=k+1}^{n}(1-p_l)}} - (1-p_i)\right).$$

Note that $F_1(0) = 1 - \frac{p_2}{p_1}$ such that bidder 1 sets an atom on zero unless $p_1 = p_2$. Figure 1.4 shows how the supports of the bidders' strategies typically look like.

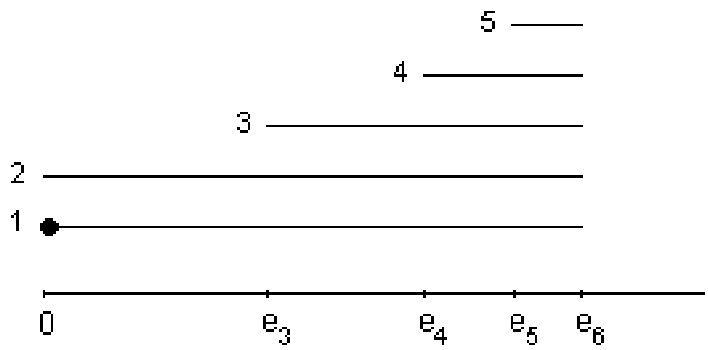

Figure 1.4: Supports of the bidders' strategies for $n = 5$

Our next result shows that the equilibrium of Proposition 1.2 is unique and characterizes the connection between tie-breaking rules and equilibrium existence.

Proposition 1.3. *For any set of tie-breaking rules, the vector of strategies in Proposition 1.2 is the unique equilibrium candidate: If it is not an equilibrium, no equilibrium exists. If it is an equilibrium, it is the unique equilibrium. If $p_1 = p_2$ so that there is no atom on zero, equilibrium exists for any tie-breaking rule. If $p_1 > p_2$, there is equilibrium existence if and only if the tie-breaking rule meets the following requirement: If bidder 1 is of strong type and all other bidders are of weak type and all bidders exert zero efforts, bidder 1 wins for sure.*

Roughly, this tie-breaking in favor of strong types is needed to ensure that payoffs are sufficiently continuous at zero. The equilibrium uniqueness sets our model apart from related models such as the complete information all-pay auction or Varian's (1980) model of sales where a multiplicity of equilibria may arise in the case of more than two bidders: For example, multiple equilibria arise in the symmetric case. The reason for this multiplicity is that several bidders can choose to stay out of competition.[12] This is in contrast to our game: In our setting, with some probability, each pair of bidders turns out to be the only bidders present in the auction. Hence two bidders never both put positive mass on zero in equilibrium. For this reason the equilibrium uniqueness of the complete information all-pay auction with two bidders transfers to our n bidder auction. To phrase the same observation differently, the model of this section differs from the complete-information all-pay auction in that in equilibrium not only the two ex-ante strongest bidders actively participate but all (strong type) bidders do.

In the two-bidder case, there is another interesting connection between our game and the standard complete information all-pay auction: One can easily see that the strategies chosen by the strong types of bidders 1 and 2 are exactly the same strategies as those of two bidders with valuations p_1 and p_2 and marginal costs c in a complete information all-pay auction. The fact that bidders in our game employ these strategies only with probabilities p_i neatly illustrates the fact that expected efforts are lower in our model due to the bidders' informational rents. The same equivalence does not translate to the case of more than two bidders since in our model all bidders make positive bids with positive probability. This stands in contrast to the complete information all-pay auction.

Note that bidders with low participation probability exert high levels of effort if they take part in the auction. The reason is the following: Bidders who rarely participate earn expected payoffs that are considerably higher than their maxmin-payoffs. This would not be possible if they exerted low effort levels: Bidders exerting low efforts win the auction if they are essentially the only ones who are participating. Yet

[12] See Baye, Kovenock and de Vries (1992, 1996).

bidders with low p are seldom the only ones participating. This is why they can earn higher profits out of high effort levels, at which competition among several bidders typically decides which bidder wins.

Additionally, the bidder who is most often participating in the auction, bidder 1, does not very often choose high effort levels, but puts considerable mass on low effort levels. This is necessary in equilibrium due to the following argument: At least two bidders have to mix down to zero in equilibrium.[13] Yet no bidder except bidder 1 would be willing to exert zero effort if bidder 1 did not play an atom in zero. Consider bidder 2: If bidder 1 played no atom, bidder 2 would only earn his maxmin-payoff from playing zero. This would be less than what he can earn in equilibrium. Hence bidder 1 has to be a soft competitor to his rivals in equilibrium.

If we put the two last observations together, we can conclude that bidders who are rarely participating in the auction have to win over-proportionally often: If they take part in the auction, they play high effort levels at which bidders who participate more often compete comparatively softly. The following corollary sums up this result:[14]

Corollary 1.3. *Assume $p_i > p_j$. Denote by P_i the probability that bidder i wins the auction. Then $P_i > P_j$, but $P_i < \frac{p_i}{p_j} P_j$.*

Like in the complete information all-pay auction, it is straightforward to generalize our results of this section (except formula (1.2) for the expected effort sum) to the case where all strong-type bidders share the same non-linear cost function. We close this section by revisiting our result of payoff-uniqueness of Lemma 1.1. This result easily generalizes to richer settings: We focus on multiple prizes and asymmetric cost functions which fulfill a local symmetry condition.

[13] This follows by standard arguments.

[14] To make the above heuristic arguments precise observe that with $p_j < p_i$ the equilibrium strategy of strong i can be seen as a mixture between strong j's strategy and a strategy over strictly lower effort levels.

Consider the above model but assume that the seller awards $k < n$ prizes of size 1. Each bidder receives at most one prize. The following result is an extension of Clark and Riis' (1998) characterization of equilibrium payoffs to the asymmetric incomplete information case. Strong i's maxmin-payoff is given by the probability that at most $k-1$ of his opponents are strong-types. Denote this probability by $\lambda_i(k)$ and denote by $\Omega_i = \{1, \ldots, n\} \setminus \{i\}$ the set of i's opponents. Then we can write $\lambda_i(k)$ as

$$\lambda_i(k) = \sum_{A \subseteq \Omega_i, |A| \leq k-1} \prod_{j \in A} p_j \prod_{l \in \Omega_i \setminus A} (1 - p_l).$$

Clearly, λ is largest for the "most present" bidder 1. It is then easy to verify that the argument behind Lemma 1.1 still goes through in this setting. This yields the following result:

Corollary 1.4. *In the model with k identical prizes, in any equilibrium, all strong types expect a payoff of*

$$\pi_i^s = \lambda_1(k)$$

while weak types expect a payoff of zero.

It is straightforward to extend this result also to multiple non-identical prizes similar to the results of Barut and Kovenock (1998) for the complete information case. A formula for the expected effort sum analogous to (1.2) can also be given.

The second extension of our payoff-uniqueness result is similar to Siegel's (2009) extension of classical results about the complete information case: For our proof of payoff-uniqueness, we only need symmetry in costs for two special effort-levels - for effort-level zero, and for the effort-level that turns out to be the upper boundary of the equilibrium supports. This yields a generalization of Lemma 1.1 where cost functions do not have to be linear, identical, nor even ranked - they just have to coincide at these two significant points:

Corollary 1.5. *Consider the basic model of this section but assume that strong i's costs are given by a continuous, increasing function c_i. Assume $c_i(0) = 0$ for all i. Furthermore assume that*

$$c_i^{-1}\left(1 - \prod_{k \neq 1}(1 - p_k)\right) = c_j^{-1}\left(1 - \prod_{k \neq 1}(1 - p_k)\right)$$

for all i and j. Then, in any equilibrium, all strong types expect a payoff of

$$\pi_i^s = \prod_{k \neq 1} (1 - p_k)$$

while weak types expect a payoff of zero.

Since an explicit characterization of the equilibrium strategies does not seem possible for the situations considered in Corollaries 1.4 and 1.5, a word on equilibrium existence in these generalized models is in order. Note that weak types will never bid anything other than zero and can thus be ignored. Now, the key observation is that the auctions are essentially equivalent to auctions in which bidders are always strong types and where each bidder submits one simultaneous bid in (finitely) many all-pay auctions: There is one auction for every subset of bidders. To this model of simultaneous all-pay auctions one can easily apply classical existence results such as Simon and Zame (1990).[15]

1.5 Applications

In the following, we consider three applications of our analysis. The first is about information-sharing in contests. The second considers all-pay auctions where the bidders choose their type-probabilities in a preliminary stage. The third applies our results to a model of competing complete-information all-pay auctions.

1.5.1 Disclosure

In this section we apply the analysis of Section 1.3 to the study of incomplete information contests where bidders have the opportunity to share some information about their type. We do not intend to give an exhaustive analysis of this problem. Instead our aim is two-fold: First, we want to contribute to a more complete understanding of information-sharing in contests. Second, we want to show how easily this problem can be studied applying the results from Section 1.3.

[15] In particular, the existence proof in Siegel (2009, p. 89) which is based on Simon and Zame's result translates to this setting with minor changes.

In a recent paper, Kovenock, Morath and Münster (2009), KMM in the following, analyze an independent private values[16] all-pay auction with two bidders where bidders are ex-ante uninformed about their valuations. Bidders decide ex-ante whether they would like to share information they learn about their valuations. Depending on the bidders' sharing decision[17], the bidders play either a complete information all-pay auction or an incomplete information all-pay auction.

As KMM, we consider two different decision frameworks for information-sharing:

1. **Industry-wide agreements:** Both bidders simultaneously cast a vote for or against sharing information. If both bidders vote for sharing, valuations are revealed and a complete information all-pay auction takes place in the final stage. Otherwise, an incomplete information all-pay auction takes place.

2. **Independent commitments:** The bidders independently make an ex-ante commitment about sharing information or not. Depending on the bidders' decisions, either both valuations, or only one, or none become common knowledge before the auction takes place.

For the case of industry-wide agreements, KMM show that the complete and incomplete information auctions yield the same payoffs to the bidders, implying that any choice of actions is a Nash equilibrium.[18] The loss in informational rents from disclosing is exactly off-set by the economic rents arising from the bidders' different strengths becoming common knowledge. For the case of independent commitments, KMM show that sharing information is strictly dominated.

[16]The second part of their paper considers common value auctions. These will not be discussed here.

[17]For the sakes of comparability and brevity, we also focus on ex-ante decisions here. Yet all our results of this section (as well as most of theirs) carry over to the situation where bidders decide after they have learned their valuations.

[18]To see this, recall that in a complete information all-pay auction the stronger bidder earns the difference in valuations while the other bidder earns zero payoffs. To see that the same is true for the incomplete information case, note that these payoffs are identical to those of an incomplete information second price auction. Thus by revenue equivalence these are also the payoffs in the incomplete information all-pay auction.

In KMM the bidders can only choose between disclosing all their private information or none. Generally, all-pay auctions where bidders can reveal parts of their private information are technically challenging: They require the analysis of asymmetric all-pay auctions with incomplete information. In Section 1.3, we derived the equilibrium of a simple class of asymmetric all-pay auctions with incomplete information. This gives us a natural starting-point for a tractable model with partial release of information.

In our framework, where partial disclosure is possible, we obtain essentially the opposite of the results of KMM. We find that the trade-off between gaining economic rents and losing informational rents is not as simple as one might think: For any partial release of information, the gain in the economic rents strictly dominates the loss in the informational rents.

We consider the following model: There are two bidders, both with a valuation of 1 for the object for sale in an all-pay auction.[19] Bidder i's marginal costs for exerting effort are c with probability p_i and C with probability $1-p_i$ where $0 < c < C < \infty$. The probabilities p_i are independent random variables drawn from a distribution F on $[0,1]$ with $E[p_i] = \mu$.[20] Ex-ante, the bidders only know F, c, and C. They know neither their realization of p_i nor the realization of their costs of bidding. The timing is as follows:

1. The bidders decide whether to share information later in the game, at stage 3. The decision game is either modeled as an industry-wide agreement or as individual commitments as described above.

2. The bidders learn their realization of p_i. Each bidder hence gets a more concrete estimate of his type.

3. Depending on the decision at stage 1, the realizations of the p_i become common knowledge or remain private information.

[19] For reasons of consistency, we focus on this case. Our results also hold when bidders are heterogeneous in valuations instead of effort costs.

[20] As will become apparent below, we could as well assume that p_1 and p_2 are independent with $E[p_1] = E[p_2]$.

4. The bidders learn the realization of their types and the all-pay auction takes place.

For the case where both bidders share their p_i, Corollary 1.2 provides the payoffs of the all-pay auction. It is easy to see that when bidder i does not share his realization of p_i, it is an equilibrium that both bidders still play the equilibrium of Corollary 1.2 but with μ in place of p_i.[21]

The following corollary - which is an immediate consequence of Corollary 1.2 - states the bidders' ex-ante expected payoffs for the different decisions about sharing information.

Corollary 1.6. *Define $\theta = 1 - \frac{c}{C}$. The ex-ante expected payoffs from the all-pay auction for the different disclosure decisions are as follows:*

1. *If both bidders decide to reveal their p_i, the ex-ante expected payoff of bidder 1 is*
$$\pi_1(1r,\ 2r) = E[p_1(1 - \min(p_1, p_2)) + (1 - p_1)(p_1 - \min(p_1, p_2))]\theta.$$
2. *If bidder 1 decides to reveal but bidder 2 not, the ex-ante expected payoff of bidder 1 is*
$$\pi_1(1r,\ 2n) = E[p_1(1 - \min(p_1, \mu)) + (1 - p_1)(p_1 - \min(p_1, \mu))]\theta.$$
3. *If bidder 1 does not reveal but bidder 2 does, the ex-ante expected payoff of bidder 1 is*
$$\pi_1(1n,\ 2r) = E[p_1(1 - \min(\mu, p_2)) + (1 - p_1)(\mu - \min(\mu, p_2))]\theta.$$
4. *If both bidders decide not to reveal, the ex-ante expected payoff of bidder 1 is*
$$\pi_1(1n,\ 2n) = E[p_1(1 - \min(\mu, \mu)) + (1 - p_1)(\mu - \min(\mu, \mu))]\theta = \mu(1 - \mu)\theta.$$

[21] There are further equilibria where the bidders utilize their realizations of p_i as a randomizing device. These equilibria are however all payoff-equivalent to the one in Corollary 1.2. The reason for the payoff-equivalence is that - as during the auction the bidders know their own types - the private information about the p_i is useless unless it is shared.

Let us now consider the bidders' disclosure decisions in the industry-wide agreements regime. Since each bidder can veto against information-sharing, it is always a weak Nash equilibrium that both bidders vote against disclosure. In the model of KMM, payoffs are the same regardless of disclosure decisions. Thus in their setting any vector of strategies is a weak Nash equilibrium. In our model we obtain instead:

Corollary 1.7. *Consider industry-wide agreements on sharing information. Assume that neither $p_i = \mu$ a.s. nor $p_i \in \{0, 1\}$ a.s.. Then $\pi_1(1r,\ 2r) > \pi_1(1n,\ 2n)$. Thus in the only strict Nash equilibrium both bidders vote for information-sharing. This equilibrium is also strictly payoff-dominant.*

In the corollary we had to exclude two cases: If p_i is deterministic, disclosure transports no information. If p_i is always in $\{0, 1\}$, disclosure is fully revealing such that we are essentially in the setting of KMM. The corollary follows immediately from observing that (except in the two excluded cases)

$$\pi_1(\text{1r, 2r}) > E[p_1(1 - p_2)]\theta = \mu(1 - \mu)\theta = \pi_1(\text{1n, 2n}).$$

Let us now turn to the game with individual commitments to share information. Under this regime, KMM show that in their setting committing to reveal information is a strictly dominated action. Accordingly, the unique Nash equilibrium is that both bidders do not disclose. This is in contrast to our model with partial revelation:

Corollary 1.8. *Consider individual commitments on sharing information. Assume that neither $p_i = \mu$ a.s. nor $p_i \in \{0, 1\}$ a.s.. Then it holds that $\pi_1(1r,\ 2n) > \pi_2(1n,\ 2n)$. Thus, given that the opponent does not reveal, it is a strictly dominant action for a bidder to reveal. Hence it is not a Nash equilibrium that both bidders withhold their private information.*

The corollary follows immediately from the observation that

$$\pi_1(\text{1r, 2n}) > E[p_1(1 - \min(p_1, \mu))]\theta > E[p_1(1 - \mu)]\theta = \pi_1(\text{1n, 2n}).$$

It depends on the distribution F whether a bidder prefers to reveal or not, given that his opponent reveals. Yet in any case, some bidder will reveal at least with some probability in equilibrium.[22]

Taking our results and those of KMM together shows that it will be difficult to settle the issue of information-sharing in contests without a model that has both a sufficiently rich type-space and a sufficiently rich model of information revelation. This is a challenging direction for further research. In the situation with industry-wide agreements, the result of KMM looks essentially like a boundary case of our result. We hence conjecture that the result of Corollary 1.7 is quite robust. The situation with individual commitments to share information is more complex. KMM rightly point out that the two-types case is an extreme case concerning individual sharing decisions: With two types, bidders are indifferent between completely revealing and not revealing regardless of the opponent's behavior. This does not carry over to a state space with more than two types. Yet we have seen that *partial* sharing is a *strictly* dominant action. It seems highly intuitive that this strict dominance will not disappear instantly, e.g., whenever a third (possibly very unlikely) type is introduced. We thus conjecture that it will depend sensitively on the distribution of types and other model parameters whether bidders want to share information or not.

Based on the results of Section 1.4, one can obtain parallel results for the n bidder case.

1.5.2 Endogenous Choice of Type-Probabilities

We now allow the bidders to choose their probabilities p_i themselves. A motivation comes from lobbying: For an interest group, it is much easier to take influence if it can get close to the decision-maker. This is e.g. why many interest groups in Europe have some representation in Brussels, close to the EU decision-makers. The better an interest group is represented in Brussels, the more likely it gets in direct contact to the decision-makers. In the following, we analyze the game where each interest

[22]For example, if F is the uniform distribution on $[0, 1]$, there are three Nash Equilibria: two equilibria where one bidder reveals for sure while the other does not, and a symmetric mixed equilibrium.

group first decides on how much representation it wants to have in Brussels. Then it decides on how much effort it wants to spend into direct lobbying activity.[23]

We hence consider the following simple two-stage auction model: There are two bidders who incur marginal costs of exerting effort c or C with $0 < c < C < \infty$. Both have a valuation of 1 for winning the (final) auction stage. In stage one, both bidders simultaneously choose their type-probabilities $p_i \in [0, 1]$. Then the marginal costs of exerting effort of bidder i realize as c with probability p_i and as C with the counter-probability. The bidders privately observe their types before the second stage takes place, where both bidders compete in an all-pay auction.

Since Corollary 1.2 uniquely determines the bidders' equilibrium payoffs from the auction stage, one can identify the SPNE of our game by finding the Nash equilibria of the reduced one-stage game in which the bidders choose their p_i. By Corollary 1.2, the payoff of bidder 1 from the auction for fixed p_1 and p_2 is given by

$$\pi_1(p_1, p_2) = (p_1(1 - \min(p_1, p_2)) + (1 - p_1)(p_1 - \min(p_1, p_2)))\theta$$

where as in the previous section $\theta = 1 - \frac{c}{C}$.

Our main result is that - somewhat counter-intuitively - bidders do not choose p_i as large as they can: It is not an equilibrium that both bidders set their probabilities to 1. We find a multiplicity of equilibria, each leading, loosely speaking, to an average probability of $\frac{3}{4}$ to become a good type.

Lemma 1.2. *Let $\{i, j\} = \{1, 2\}$. Nash equilibria of the probability setting game are given by:*

- *Bidder i sets $p_i = 1$ and bidder j sets $p_j = \frac{1}{2}$.*
- *Bidder i sets $p_i = \frac{3}{4}$ and bidder j plays either $p_j = \frac{1}{2}$ or $p_j = 1$ with equal probability.*
- *Both bidders mix uniformly over the interval $[\frac{1}{2}, 1]$.*

[23]Lobbying has been one of the oldest and most prominent applications of all-pay auctions. See, e.g., Hillman and Samet (1987), Hillman and Riley (1989) and Che and Gale (1998).

In the proof of this lemma we analyze a slightly more general problem which can be easily adapted, e.g., to the case in which the bidders have marginal costs of 1 and either valuation V or v.[24] Bidders accept the risk of becoming a weak type to soften competition in the auction.[25]

Coordination on a specific equilibrium is simple if we impose a time-structure and assume that one bidder decides first (as the incumbent) and the other second (as the entrant). In this case, the incumbent chooses 1 and the entrant $\frac{1}{2}$.

1.5.3 Competing Contests

The Basic Model

As a third and last application we analyze a situation of multiple competing contests: Assume there are $n \geq 2$ bidders who can exert effort in one (and only one) of two all-pay auctions, A and B. The bidders' marginal costs of exerting effort equal 1. Winning contest A leads to a utility of V_A and winning contest B leads to a utility of V_B. Bidders decide simultaneously in which contest they want to participate and how much effort they want to exert. For $X \in \{A, B\}$, we denote by p_X^i bidder i's probability of choosing contest X and by H_X^i the effort setting strategy of bidder i conditional on participating in X.

Once again, it is not our main objective here to carry out an exhaustive analysis of the model. Rather, we want to show that our results of, in this case, Section 1.4 contribute with minor modifications also to the literature on competing contests. In the following we characterize the symmetric equilibrium of the model. We then show that, with few bidders, a single all-pay auction may yield higher aggregate efforts than the same auction together with an additional, second all-pay auction with a smaller prize. As by-products, we show how bidders' participation strategies

[24] Note that while asymmetries in valuations and asymmetries in costs are equivalent at the interim-stage (when effort is exerted), they are not equivalent (but close to each other) at the ex-ante stage. See footnote 3 in the Appendix for details.

[25] A similar behavior is seen in models of advertising, compare Ireland (1993) and McAfee (1994). In these models, a Bertrand pricing game is played in the second stage instead of an all-pay auction.

determine their strategies in all equilibria, and that regardless of V_A and V_B the number of participants in each contest converges to $n/2$. Our results are easy to extend to more than two competing auctions, as will become clear below. In Section 1.5.3 we take a brief look at a generalized model where bidders are inhomogeneous with respect to the contests they can participate in: One ("global") contest is open to all bidders. Two further ("local") contests are each restricted to bidders from one out of two disjoint groups. We characterize the distribution of bidders across the three contests in the case of many bidders and show that it is independent of prize asymmetries.

Competing contests are a highly natural setting. Many of the participants in real-life contests choose their contest out of a number of alternatives: Researchers pick the journal to which they submit their article out of several journals. Firms make choices in which R&D contest they want to participate. Architects decide for which project they want to create a proposal. Nevertheless, these systems of contests have received relatively little attention in the literature so far: Amegashie and Wu (2006) consider a setting which differs from ours in two respects. They allow for asymmetries in the bidders' commonly known valuations, and they assume that the bidders observe their opponents' choices of contest before they choose their effort-levels. Due to the latter assumption, Amegashie and Wu only have to deal with complete information all-pay auctions at the contest stage. In contrast, in our model, mixed strategies concerning the choice of contest may create informational asymmetries (and thus informational rents) at the contest stage. Recently, large simultaneous contests, so-called crowd-sourcing contests, have been discussed in the computer science literature.[26] From this literature, DiPalatino and Vojnovic (2009) is closest to us, with the main difference that they assume incomplete information about the bidders' valuations. Yet the focus of their study is quite different: While we are mostly interested in the comparative statics of the equilibrium in the number of bidders n, their analysis focuses on a large systems limit. As we will see, there are

[26] The term "crowd-sourcing" goes back to Howe (2006). For theoretical results see Archak and Sundarajan (2009), Chilton and Horton (2010), DiPalatino and Vojnovic (2009), and the references therein. For a broader perspective on the theoretical literature, see Jain and Parkes (2009).

considerable qualitative differences between, e.g., $n = 2$, $n = 4$ and $n = 6$. Hence considering moderate numbers of bidders is an interesting task.

To see how our model of competing contests is related to the model of Section 1.4 look briefly at the case $V_A = V_B$. Assume that all opponents of bidder i mix with equal probability between participating in contests A and B. Furthermore assume that all of i's opponents use the strategies stated in Proposition 1.2 for the case $p_1 = \ldots = p_n = \frac{1}{2}$.[27] Then trivially bidder i faces the same situation in the two contests: There are $n - 1$ opponents each of whom is present with probability $\frac{1}{2}$. Hence bidder i has to choose between two instances of the game analyzed in Section 1.4. Thus bidder i does not want to deviate from playing the same strategy as his opponents and we have found a Nash equilibrium. Carrying this logic a bit further and making use of the equilibrium uniqueness we proved for Section 1.4, we can easily conclude the following necessary conditions for Nash equilibria:

Corollary 1.9. *Assume that a vector of strategies* $(p^i_A, H^i_A, H^i_B)_i$ *forms a Nash equilibrium.*

(i) For $X \in \{A, B\}$*, denote by* $I_X \subseteq \{1, \ldots, n\}$ *the set of potential particiants in* X*, i.e. the set of bidders with* $p^i_X > 0$*.*

a) If $I_X = \{i\}$*, bidder* i *exerts an effort of zero and earns a payoff of* V_X *from participating in* X*.*

b) If $|I_X| \geq 2$ *and if* $p^i_X = 1$ *for at most one bidder* i*, the strategies* H^i_X *are given by the equilibrium of Proposition 1.2 for the case of* $|I_X|$ *bidders and participation probabilities* $(p^i_X)_{i \in I_X}$*. Accordingly, equilibrium payoffs follow from Lemma 1.1.*

c) If $|I_X| \geq 2$ *and if* $p^i_X = 1$ *for two or more bidders, contest* X *is essentially a standard symmetric complete information all-pay auction with zero payoffs and the well-known non-uniqueness of equilibrium.*[28]

[27] As noted above it is straightforward to generalize the results of Section 1.4 to a prize which does not equal 1. Since we are mostly concerned with the equilibrium payoffs here (which are simply multiplied by the value of the prize), we abstract from this technicality and adapt the according results from Section 1.4 to the prizes V.

[28] See Baye, Kovenock and de Vries (1996).

(ii) *All bidders earn identical payoffs in equilibrium. Expected equilibrium payoffs also coincide for any contest that is played with positive probability in equilibrium.*

Thus we find that the p_A^i uniquely determine the equilibrium payoffs and that these equilibrium payoffs can be calculated immediately from Lemma 1.1. Part (ii) follows from the fact that if bidder i makes a payoff π in contest X then any other bidder can make a payoff arbitrarily close to π by exerting an effort marginally above the support of H_X^i.

Before we derive a symmetric mixed strategy equilibrium for the case $V_A \geq V_B$, let us note that there are also asymmetric equilibria involving coordination. For $n = 2$ and $V_A \geq V_B$, an equilibrium is characterized by $p_A^1 = 1$ and $p_A^2 = 1 - \frac{V_B}{V_A}$. Both bidders expect a payoff of V_B. In the case of $V_A = V_B$ this equilibrium simplifies to each bidder exerting zero effort in one of the contests. For $n = 3$, an equilibrium is characterized by $p_A^1 = 1$, $p_A^2 = 0$ and $p_A^3 = V_A/(V_A + V_B)$. This yields equilibrium payoffs of $V_A V_B/(V_A + V_B)$. For $n \geq 4$, Nash equilibria are given by at least two bidders playing a complete information all-pay auction in contest A and at least two bidders playing a complete information all-pay auction in contest B with certainty. In this equilibrium all bidders earn zero payoffs. In neither of these cases the payoffs of the asymmetric equilibrium coincide with those of the symmetric mixed strategy equilibrium we determine now.

Denote by p_A and p_B the participation probabilities in a symmetric mixed strategy equilibrium. By Part (ii) of Corollary 1.9 expected payoffs must be identical in contests A and B. Hence we get by Lemma 1.1 that

$$V_A(1 - p_A)^{n-1} = V_B(1 - p_B)^{n-1}$$

which implies (since $p_A = 1 - p_B$) that

$$\frac{V_A}{V_B} = \left(\frac{p_A}{p_B}\right)^{n-1}.$$

As $p_A = 1 - p_B$, this yields that in the symmetric equilibrium

$$p_A(V_A, V_B, n) = \frac{V_A^{\frac{1}{n-1}}}{V_A^{\frac{1}{n-1}} + V_B^{\frac{1}{n-1}}} \text{ and } p_B(V_A, V_B, n) = \frac{V_B^{\frac{1}{n-1}}}{V_A^{\frac{1}{n-1}} + V_B^{\frac{1}{n-1}}}.$$

We thus see that with two bidders p_A and p_B stand in the same proportions as V_A and V_B. As n increases, p_A and p_B quickly converge monotonically to $\frac{1}{2}$. Obviously, this implies that the expected numbers of participants in the respective contests are close to $\frac{n}{2}$ for large n. Note as well that p_X only depends on the ratio of V_A and V_B.

For the remainder of this section we are concerned with the aggregate efforts in the symmetric equilibrium. Like in Section 1.4 we can calculate the expected sum S_X of efforts exerted in contest X as the difference between the expected value of prizes awarded minus the sum of the bidders' expected payoffs from contest X. Accordingly,

$$S_A(V_A, V_B, n) = V_A(1 - (1 - p_A)^n) - nV_A p_A(1 - p_A)^{n-1}$$

and

$$S_B(V_A, V_B, n) = V_B(1 - (1 - p_B)^n) - nV_B p_B(1 - p_B)^{n-1},$$

where p_A and p_B depend, of course, on V_A, V_B, and n. Define the overall expected sum of payoffs as

$$S(V_A, V_B, n) = S_A(V_A, V_B, n) + S_B(V_A, V_B, n).$$

In the following we analyze the function $S(V_A, V_B, n)$. For the sake of brevity we rely for the most part on numerical instead of analytical results (which are, of course, straightforward to derive). As a preliminary step, we address the question of how to optimally split up a fixed sum of prize money over the two contests in order to maximize S. Figure 1.5 depicts $S(V, 1 - V, n)$ as a function of V for different numbers of bidders n. The dotted line corresponds to the overall prize money of 1. Not surprisingly, regardless of n it is optimal to concentrate all prize money in one contest and worst to install two prizes of equal size. For $n = 2$ aggregate effort more than doubles if one contest with prize 1 is played instead of two contests with prizes of $\frac{1}{2}$ each. For $n > 7$ there is little quantitative difference between the different ways of splitting up the prize - there are enough bidders to guarantee almost full dissipation of rents in both contests. The most interesting observation to make from the figure is that for small values of n even a slight move out of the corners leads to a substantial loss in aggregate effort.

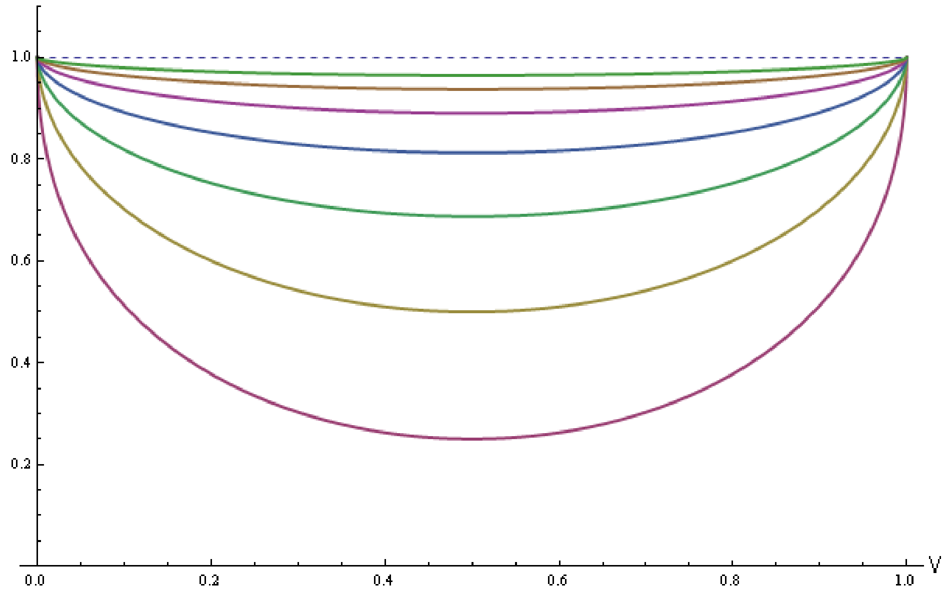

Figure 1.5: Aggregate effort when $V_A = V$ and $V_B = 1 - V$ as a function of V for (bottom to top) $n = 2, \ldots, 8$.

Let us now consider aggregate efforts when one contest offers a prize of 1 and a second contest offers a prize of $V_B \leq 1$. How does the size of V_B affect aggregate efforts? Clearly, there are two effects at work: First, more prize money makes exerting efforts more attractive. Yet second, as we already saw in the situation of Figure 1.5, introducing a second contest may be detrimental to competitiveness and may hence lead to lower aggregate effort. Our main result is that the second effect is stronger than one might initially think. This can be seen from Figure 1.6: The dotted lines correspond to the two benchmarks, i.e., to the total prize money $1 + V_B$ and to the aggregate efforts of 1 for $V_B = 0$. With two or three bidders, any prize $V_B \leq 1$ leads to lower aggregate efforts than $V_B = 0$. For intermediate n it sensitively depends on the value of V_B whether the presence of the second contest is detrimental to aggregate efforts or not. A small prize V_B is detrimental while a larger prize increases aggregate efforts. For larger n (starting with, roughly, $n = 7$), aggregate efforts are close to $1 + V_B$, i.e., aggregate efforts are almost as large as for a single prize of value $1 + V_B$.

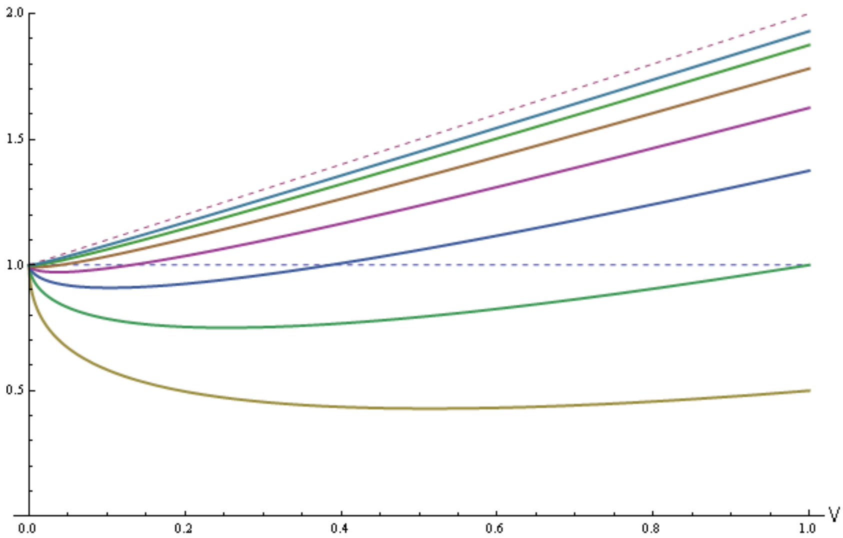

Figure 1.6: Aggregate effort when $V_A = 1$ and $V_B = V$ as a function of V for (bottom to top) $n = 2, \ldots, 8$.

Clearly, with few bidders, introducing a second contest mostly leads to less competition in the first contest and hence to lower aggregate effort. In fact, even for large n, introducing a second contest is detrimental if V_B is sufficiently small: The reason is that a contest over a very low prize is not very attractive. Hence bidders are not willing to compete over the prize. The second contest then basically works as an outside option for each bidder. To give some analytic hint at this phenomenon, consider the derivatives of $S_A(1, V_B, n)$ and $S_B(1, V_B, n)$ at $V_B = 0$. It can be shown that for all n

$$\frac{d}{dV_B} S_A(1, V_B, n)\bigg|_{V_B=0} = -\infty \text{ and } \frac{d}{dV_B} S_B(1, V_B, n)\bigg|_{V_B=0} = 0.$$

This observation implies that increasing V_B away from zero leads to a substantial drop in S_A which is not backed up by a corresponding increase in S_B. Hence we conclude that the detrimental effect of small prizes exists for all values of n even though it is hardly noticeable quantitatively for $n > 7$.

Global and Local Contests

The analysis of the preceding section is only a small part of what can be shown about models of competing contests using our techniques. To substantiate this claim we take a brief look at the following related model: There are $n = n_1 + n_2$ bidders, $n_1 \geq 2$ of them in group 1, the remaining $n_2 \geq 2$ of them in group 2. There are three contests, a global one with prize $G > 0$ and two local ones with prizes $L_1 > 0$ and $L_2 > 0$. As in the above model, bidders simultaneously decide in which contest to participate and how much effort to exert. The only difference is that bidders from group i can only compete for L_i but not for L_{-i}. The contest for the global prize G is open to both groups. For instance, the two groups can be thought of as researchers from two different fields that can either submit to the general interest journal or to the their respective field journal.

We focus on equilibria where all bidders in group i follow the same (typically mixed) strategy. For $i = \{1, 2\}$, denote by p_i the probability with which a bidder in group i exerts effort in the local contest for L_i. With $1 - p_i$, the bidder participates in the global contest. It is straightforward to see that the arguments behind Corollary 1.9 still apply here: The probabilities p_i determine the equilibrium payoffs and the strategies of effort exertion. Expected equilibrium payoffs in all three contests must be identical. Thus an equilibrium is characterized by the following system of equations:

$$L_1(1 - p_1)^{n_1 - 1} = G{p_1}^{n_1 - 1}{p_2}^{n_2 - 1} \max(p_1, p_2) \tag{1.3}$$

and

$$L_1(1 - p_1)^{n_1 - 1} = L_2(1 - p_2)^{n_2 - 1}. \tag{1.4}$$

The factor $\max(p_1, p_2)$ in the first equation stems from the payoff formula from Lemma 1.1. Solving (1.3)-(1.4) for p_1 and p_2 pins down an equilibrium of the game. While an explicit solution is difficult, a numerical solution is easy to obtain: Using (1.4) we can eliminate p_2 from (1.3). This leaves us with the problem of finding the values of p_1 which solve (1.3).[29] We will not go into the details of this analysis here. Instead, we focus on analyzing the limits of p_1 and p_2 as n gets large. This is equivalent to an analysis of how a large number of bidders is distributed over the three contests. We identify non-trivial limits $\overline{p}_1$ and $\overline{p}_2$. It turns out that these depend on the ratio of the group sizes but not on the values of the three prizes. This

[29]It is also straightforward to show that such a solution exists.

sets into perspective our previous result that in the main model of this section both contests attracted an equal number of bidders when many bidders were present.

We assume that $n_1 - 1 = \alpha(n-2)$ and accordingly $n_2 - 1 = (1-\alpha)(n-2)$.[30] Thus, for large n, α is the proportion of group 1 bidders in the population. Inserting the definition of α into equations (1.3) and (1.4) and rearranging yields, respectively,

$$\left(\frac{L_1}{G\max(p_1, p_2)}\right)^{\frac{1}{n-2}} = \frac{{p_1}^{\alpha}{p_2}^{1-\alpha}}{(1-p_1)^{\alpha}} \tag{1.5}$$

and

$$\left(\frac{L_2}{L_1}\right)^{\frac{1}{n-2}} = \frac{(1-p_1)^{\alpha}}{(1-p_2)^{1-\alpha}}. \tag{1.6}$$

Now taking n to infinity yields the following system of equations for the limits $\overline{p}_1$ and $\overline{p}_2$ of $p_1(n)$ and $p_2(n)$:

$$1 = \frac{{\overline{p}_1}^{\alpha}{\overline{p}_2}^{1-\alpha}}{(1-\overline{p}_1)^{\alpha}} \tag{1.7}$$

and

$$1 = \frac{(1-\overline{p}_1)^{\alpha}}{(1-\overline{p}_2)^{1-\alpha}}. \tag{1.8}$$

We will not go into a detailed analysis of the difference between $p_i(n)$ and $\overline{p}_i$ here. Considering the left hand sides of equations (1.5)-(1.6) suggests however that we have a quicker convergence - and thus a smaller approximation error for a fixed n - when the three prizes do not differ too much.

Through (1.8) one can eliminate $\overline{p}_2$ from (1.7). Then it is straightforward to calculate $\overline{p}_1$ from (1.7) for fixed values of α. The resulting Figure 1.7 depicts $\overline{p}_1$ as a function of α. If group 2 does not exist, i.e. for $\alpha = 1$, bidders in group 1 mix between the global and the local contest with equal probability. This replicates our result of convergence to the uniform distribution in the basic two-contest model of Section 1.5.3. If group 1 is comparatively small, i.e. for small α, the bidders from group 1 participate in the local contest with very high probability. In the symmetric case

[30] In our model, this assumption is analytically more convenient but essentially equivalent to assuming $n_1 = \alpha\, n$. A reader who is concerned with the n_i not being integers can choose $\alpha \in \mathbb{Q}$ and let n tend to infinity along a suitable subsequence of $\mathbb{N}$.

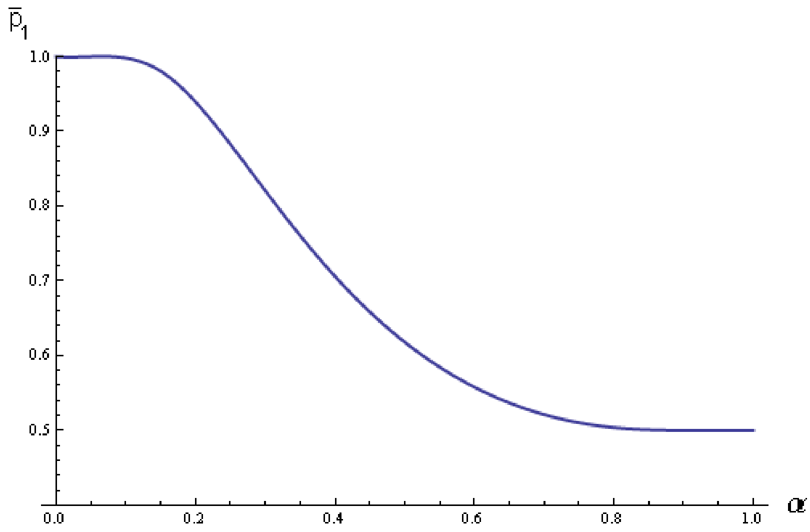

Figure 1.7: Limit proportion of group 1 bidders choosing the local contest, $\overline{p}_1$, as a function of the limit fraction of group 1 bidders, α.

$\alpha = \frac{1}{2}$, we obtain the explicit solution

$$\overline{p}_1\left(\frac{1}{2}\right) = \overline{p}_2\left(\frac{1}{2}\right) = \frac{\sqrt{5}-1}{2} \approx 0.618.$$

Thus in this case about 61.8% of the contestants participate in their local contests. Notably, the masses of the bidders in the global contest and of the bidders in the local contests stand in the celebrated Golden Ratio, i.e.,

$$\frac{\overline{p}_1\left(\frac{1}{2}\right)n}{\left(1-\overline{p}_1\left(\frac{1}{2}\right)\right)n} = \frac{1+\sqrt{5}}{2}.$$

Observe that since prize asymmetries have vanished in the limit over n, it holds that $\overline{p}_1(\alpha) = \overline{p}_2(1-\alpha)$.

Figure 1.8 depicts the limit distribution of bidders across contests as a function of α. The bidders in group 2 are depicted above the dotted line, the bidders in group 1 are below. The group members who participate in the global contest are shown around the dotted line (in the shaded area), while participants in the two local contests are, respectively, below and above that region. The overall participation in the global contest is remarkably stable: It varies between 0.5 (if only one group is present) and roughly 0.38 (for equal group sizes). Moreover, if the larger group is less than five

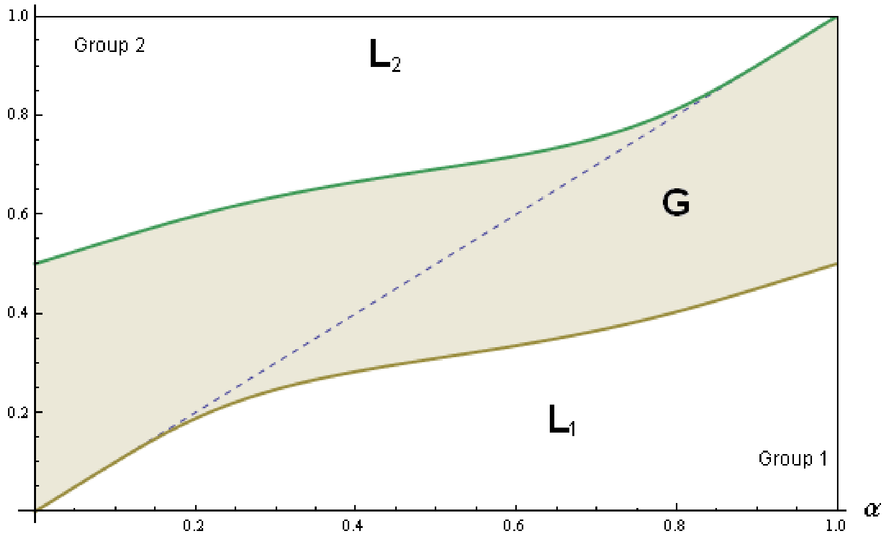

Figure 1.8: Limit distribution of bidders across contests as a function of α.

times larger than the smaller group, i.e. $\alpha \in [0.2, 0.8]$, the proportion of bidders in the global contest varies even considerably less (between 0.38 and 0.41). Closer to the corners, we observe a remarkable stickiness: As long as α smaller than 0.2 or larger than 0.8, the proportions $\overline{p}_i$ are almost constant. The smaller group stays in its local contest almost exclusively while the larger group is split up in nearly equal proportions between its local contest and the global contest. When α becomes larger than, roughly, 0.2 (or smaller than 0.8), there is a sudden change in behavior: Both groups start to participate substantially in the global contest.

1.6 Conclusion

We have analyzed an asymmetric all-pay auction with incomplete information about the different bidders' types and their different type-probabilities. The assumption of two different types enabled us to carry out an explicit analysis of equilibrium in an asymmetric auction setting when information is incomplete.

With the help of our results, one can study asymmetries in all areas in which all-pay auctions are popular models, such as lobbying, rent-seeking, R&D activities, or sport contests.[31] Our results may also serve as an easy-to-use tool for the analysis of richer models. For instance, we considered models with multiple stages, where asymmetries often arise naturally in the course of the game.

[31]For an overview, see Konrad (2009).

Chapter 2

Welfare in Markets where Consumers Rely on Word of Mouth

We analyze a market model with rational firms knowing the distributions from which their opponents' qualities are drawn. Firms engage in price competition. Following Spiegler (2006a) we assume that consumers only see the firms' prices and rely on word of mouth in order to judge the firms' different qualities.
We prove equilibrium uniqueness for the special case of complete information on the firms' side. With the help of this result, we characterize all equilibria of the general incomplete information model. Different equilibria generate identical payoffs for the firms, but different welfare results. In the monotone pricing equilibrium, welfare converges to zero in the number of firms.

2.1 Introduction

There are plenty of markets in which consumers are not fully aware of the different qualities of specialists and rely on word of mouth. Whereas consumers are unfamiliar with the market, specialists know the market well: They know how good they are, they have a good idea about their competitors' qualities, and they know how consumers search for them. Specialists act rationally.

Consumers who come into an unfamiliar market see the prices charged by the firms, but not the different firms' qualities. They rely on word of mouth to get a rough idea about the qualities. If several firms seem to offer a good quality, consumers focus on the prices as the ultimate selection device. Even though high quality firms get recommended more often, they may have to compete in prices against much lower quality firms if those get recommended as well. Anticipating this, firms play a very different pricing game than in a traditional market model with rational consumers. Competition is not necessarily beneficial in this market - as soon as some firms compete against each other, welfare may decrease substantially. This is the situation explored in this chapter.

In the following, we mostly stick to markets for health care and health insurance as leading examples. The "healers" in our model can hence be seen as specialized therapists or as providers of health care insurance. Yet our analysis applies to any market with which consumers are not familiar, such that they rely on anecdotal evidence to judge the qualities of different firms, e.g. markets for car repair or markets for financial advice. In these markets, prices are salient and easy to grasp, but quality is not.

A large body of experimental research shows that anecdotes serve as a compelling and convenient tool for transporting information and influencing behavior.[1] In the medical literature, there is broad evidence that patients rely on anecdotal reasoning.[2] Even if statistical information on different forms of therapy is available (which often is not the case, e.g., for surgical treatments)[3], patients tend to prefer to rely on personal stories. Fagerlin et al. (2005) point out identification and emotional feelings as driving factors behind this. Patients may find it much easier to identify with a "natural person" than with the "statistically average person".[4] Additionally,

[1] See, e.g., Kahneman and Tversky (1973), Borgida and Nisbett (1977).

[2] Compare, e.g., Fagerlin et al. (2005) and the references therein and Enkin and Jadad (1998).

[3] Compare Gattellari et al. (2001) and McCulloch et al. (2002). McCulloch et al. (2002) state that "treatments in general surgery are half as likely to be based on RCT [Randomised Control Trials] evidence as treatments in internal medicine" (p. 1448).

[4] Compare also Jenni and Loewenstein (1997).

in situations of uncertainty, people are often driven by emotions, and anecdotes transport more emotions than statistical results.[5]

Fagerlin et al. (2005) see another compelling characteristic about anecdotes in that "anecdotal information often provides a clear dichotomy − either an individual was cured or not" (p. 399). This kind of information may be much easier to grasp for a lay person than some statistical percentage of getting cured, and hence be much more easy to relate to. Indeed, most people have difficulties in understanding percentages and basic statistical concepts. For example, the importance of sample size is typically not recognized by untrained subjects. This has been shown in general studies as well as in medical contexts like cancer treatment or cancer screening.[6]

We assume that patients rely on word of mouth regarding quality in an otherwise standard market model: Patients only think about attending a healer if they heard some good story about him, and avoid those healers on which they heard something negative. When patients heard some favorable report about several healers, they opt for the recommended healer with the lowest price. The way we model the patients' behavior is known as the $S(1)$-rule (where $S(1)$ stands for "sampling once") and goes back to Osborne and Rubinstein (1998). It has been applied as well by Spiegler (2006a, 2006b), Rubinstein and Spiegler (2008) and Szech (2010). Of these papers, Spiegler (2006a) is closest to the present study as we discuss below. Szech (2010) complements our welfare analysis by shedding light on the case where healers choose their qualities themselves, see the discussion in Section 2.5.

Besides S(1), other related approaches for modeling boundedly rational consumer behavior are Ellison and Fudenberg's (1995) "word-of-mouth learning" and Rabin's (2002) "law of small numbers". Further related models from the literature on bounded rationality are reviewed in Spiegler (2006a). More broadly, this chapter contributes to the literature on interactions between rational firms and boundedly

[5] Compare Loewenstein et al. (2001) and Finucane et al. (2000).

[6] Compare among others Tversky and Kahneman (1971), Hamill et al. (1980), Garfield and Ahlgren (1988), Yamagishi (1997), Schwartz et al. (1997), Weinstein (1999), Lipkus et al. (2001), Weinstein et al. (2004).

rational consumers as surveyed by Ellison (2006).

We generalize the model of Spiegler (2006a). He analyzes a market of quacks who all have the same qualities and do not succeed better than some costless outside option the patients could choose instead. In our model, healers have true healing powers, but may differ strongly in their healing qualities. Additionally, we assume that the healers do not know the qualities of their competitors perfectly: Healers only know the distributions from which the qualities of their competitors are drawn. Spiegler shows different types of market failure, e.g. that patients' surplus may fall in the number of quacks for a low overall number of quacks in the market. Yet this negative effect of competition disappears if the number of quacks gets larger. Harsh competition among many quacks drives the prices down. As the quacks offer identical (low) qualities, patients fare better as competition becomes strong. In contrast, in our model, patients' surplus typically increases for low numbers of healers, but starts to decrease substantially when too many healers are active and even goes to zero when healers employ monotone pricing strategies. This negative effect of competition is in stark contrast to the predictions made by standard market models.

While our way of modeling anecdotal reasoning follows Spiegler, the logic behind our results is novel: Through the introduction of incomplete information we obtain pure equilibria which differ markedly from the mixed equilibria analyzed in Spiegler (2006a) and Szech (2010). Even a tiny amount of uncertainty in the quality realizations allows for pure price strategies where better healers charge higher prices. In this monotone equilibrium, patients who cannot properly distinguish between qualities are naturally driven to low quality healers: Patients pick the healer with the lowest price among all recommended healers. Thus they end up with the worst healer among the recommended ones if prices are monotone in quality.

Let us at this point turn briefly to the political debate about the performance of the US health insurance systems: In light of our results, it is not surprising that recent research revealed that the Veterans Health Administration (VHA) often offers better quality treatments than competitive health insurers in the US:[7] In contrast to most other medical insurers in the US, like Health Maintenance Organizations (HMOs), the VHA does not stand in competition to other insurers. The patients of the VHA typically stick with the institution for the rest of their lives. In contrast, the customers of most other US insurance systems typically switch their health plans on a regular basis. Hence competition plays a big role for most US health plans, but not for the VHA.

Our model gives an intuition for why competition can be detrimental to welfare in the health insurance market: For patients, seeing prices of medical plans is easy. Comparing myriads of different care plans for various health problems is difficult. When choosing their insurance, patients may therefore screen the different medical plans only with regard to coverage of a small sample of conditions. Alternatively, they may rely on recommendations by other insured and their limited experiences. Our analysis shows that in such a market, the choice among many insurers may lead to patients ending up with medical plans of low quality.

That in the complex insurance market consumers may indeed focus on prices too much is underpinned by the recent decision in Germany to legally cut back the price competition among social health insurers to a minimum. The idea behind was to force people to put their attention away from price differences to quality differences.[8]

Our model also contributes to explaining why many therapies that lack evidence of

---

[7] See Brooks (2008) and Longman (2010). Based on 294 indicators of quality, Asch et al. (2004) find that the VHA scores higher than all other sectors of American health care. Patients inside the VHA receive significantly better adjusted overall quality, better chronic disease treatment and preventive care.

[8] Compare the following statement by the German Federal Ministry of Health (2009), translated: "The uniform insurance fee ends the unfair competition for the cheapest fee. Instead it opens a fair competition for the best service and additional benefits to the insured." Clearly, such "fairness" considerations would be pointless with perfectly rational patients.

therapeutical advantage compared to simpler therapies[9] or placebos[10] survive in the health market. Even if a doctor has the best intentions, if his therapeutical method does not help patients in the best possible way, patients should better go elsewhere. To maximize overall well-being (welfare), only the best therapies should survive in the market. Our model shows that even strong competition over patients who rely on word of mouth does not drive out therapists of poor quality.

From a theoretical point of view, Ireland (1993) and McAfee (1994) analyze a closely related game with a different interpretation, namely advertising, in mind: Competition over consumers is modeled as in our study, yet firms know each others' qualities (or rather advertising intensities in their specification) for sure. We add to the analysis of these two papers the equilibrium uniqueness in the pricing stage. The question of equilibrium uniqueness had been pointed out by both authors as an open problem. The uniqueness result stands in an interesting contrast to the multiplicity of equilibria in related models of price dispersion such as Varian's model of sales (1980) or the complete information all-pay auction.[11] Equilibrium uniqueness in the complete information case is a crucial step for characterizing all equilibria of our general game with incomplete information.

Finally, let us point out that the natural definition of welfare is fundamentally different for the market models studied by McAfee and Ireland, as there are no differences in the firms' service-qualities, but only in the firms' advertising activities. Thus, in these models, welfare increases in the number of firms.

The chapter is structured as follows: Section 2.2 presents the model and describes in detail the behavioral $S(1)$ rule our patients follow. In Section 2.3, we characterize all equilibria of the model where each healer's quality may be drawn from a different

[9] For example, arthroscopic surgery, one of the most often performed surgeries with the aim of lowering pain in arthritic knees, was only recently questioned by Kirkley et al. (2008), who doubt the efficacy of this therapy. Kallmes et al. (2009) find that vertebroplasty, a commonly performed spinal surgery to treat osteoporotic compression fractures, leads to no improvements in pain and pain-related disability.

[10] Compare Fontanarosa et al. (1998).

[11] See Baye, Kovenock and de Vries (1992, 1996).

distribution function. As a by-product, we show equilibrium uniqueness for the pricing game of Ireland (1993) and McAfee (1994). In Section 2.4, we assume that the healers' qualities are independently drawn from the same distribution F. We characterize the equilibrium in monotone price setting strategies. We show that as the number of healers gets large, overall welfare goes to zero. As an example, we assume qualities to be uniformly distributed: Welfare starts to decrease (and decreases substantially) in the number of healers as soon as there are more than three healers in the market. Section 2.5 discusses the robustness of our results. Section 2.6 concludes.

2.2 The Model

We consider a market with n rational healers and a continuum of mass 1 of boundedly rational patients. The quality α_i of healer i is drawn from a distribution function F_i. The F_i are commonly known by all healers, but not by the patients. The supports of all F_i are assumed to be subsets of $[0, 1]$. Furthermore, we assume that the expected quality $\overline{\alpha_i}$ of each healer i satisfies $0 < \overline{\alpha_i} < 1$. Without loss of generality, healers are sorted by expected qualities, i.e., $\overline{\alpha}_i \leq \overline{\alpha}_j$ for $i < j$. Let $E[\cdot]$ denote the expectation with respect to the α_i. Initially, the patients are ill. They have a utility of 0 from staying ill, and a utility of 1 from getting cured. A healer with quality α cures each of his patients with probability α independently of each other.

The **timing** is as follows:

1. Each healer learns his personal quality realization α_i. This information is private.

2. The healers set their prices P_i simultaneously.

3. The patients decide whether to attend a healer and if so, which one.

4. Patients who consult healer i get cured with probability α_i.

In Step 3, patients decide according to the behavioral rule **S(1)** as introduced by Osborne and Rubinstein (1998), and as utilized in Spiegler (2006a). This rule works as follows:

- Each patient independently receives a signal on each healer.
- With probability α_i, a patient receives a positive signal $S_i = 1$ on healer i ("**a recommendation**").
- With probability $1 - \alpha_i$, a patient receives a negative signal $S_i = 0$ on healer i ("**no recommendation**").
- A patient attends the healer with the highest $S_i - P_i$, unless $\max_i S_i - P_i < 0$. In that case the patient stays out of the market and expects a utility of 0 at a price of 0.

The last point implicitly contains a tie-breaking rule: If a patient has to choose between consulting a recommended healer at a price of one or staying at home, the patient opts for the healer. It can be shown that no equilibrium exists if we depart from this assumption. All other ties can be broken arbitrarily.

Note that patients rely far too much on the signal they get - they over-infer from their sample. The idea behind the S(1) rule is to capture reliance on anecdotal evidence in a simple way: Each patient independently asks some "former" client of each healer.[12] A client of healer i got cured with probability α_i. Thus, with probability α_i, he recommends healer i to the patient. The patient perfectly trusts this report - he either thinks the healer can cure him for sure or not at all.

Note that if a patient consults healer i his utility is $1 - P_i$ with probability α_i and $-P_i$ otherwise. The $S(1)$ rule is supposed to capture the idea that patients are not familiar with how the market works in detail. In particular patients are not aware of the healers' qualities α_i: Patients act as if some healers were always successful and others never.

[12] Of course, this dynamic motivation is only for intuition, as we are in a static model here.

Finally, an alternative interpretation of the S(1) rule is as follows: Assume the healer is a health insurer and the quality of the insurer is given by the proportion of medical problems covered by his insurance plan. Patients just sample each insurer with regard to one random medical condition. Hence with probability α_i they receive the positive signal that the condition considered is covered by the insurance plan. They then think this plan is a good one, in contrast to the plans on which they received a signal of non-coverage. The medical condition the patient faces in the future is independently drawn by nature.

2.3 Characterization of All Equilibria

In this section, we determine all equilibria of the model. For the analysis, it is helpful to consider the model with deterministic qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ as well. This is the special case of our model where the healers hold complete information about each others' qualities. This model has been analyzed by Ireland (1993) and McAfee (1994) in the context of advertising. We add the uniqueness to their characterization of equilibrium. This result is a crucial step towards the characterization of all equilibria of the incomplete information game.

To put the equilibria we find into perspective, note that if the healers know each others' qualities perfectly, and if qualities are strictly between 0 and 1, standard arguments yield that there cannot be an equilibrium in pure strategies: Each healer has the possibility to earn a positive expected payoff, as with some probability he is the only recommended healer in the market.[13] Hence each healer chooses a price strictly higher than zero. Thus pure pricing strategies cannot constitute an equilibrium, as there would always be a healer who would like to attract more patients by deviating to a slightly lower price. This is why in the game with complete information about qualities, the equilibrium must be in mixed strategies. The unique mixed equilibrium of this game is given by Proposition 2.1.

[13]Recall that patients never attend healers that are not recommended, as they expect a negative utility of $0 - P_i$ from attending them.

Proposition 2.1. *Define a sequence of prices* $p_0, \ldots, p_n$ *by*

$$p_i = \frac{(1-\overline{\alpha}_{i+1}) \cdot \ldots \cdot (1-\overline{\alpha}_{n-1})}{(1-\overline{\alpha}_i)^{n-i-1}}$$

for $1 \leq i \leq n-2$,

$$p_0 = \prod_{i=1}^{n-1}(1-\overline{\alpha}_i)$$

and $p_{n-1} = p_n = 1$. *Then the unique Nash equilibrium of the complete information game with qualities* $\overline{\alpha}_1, ..., \overline{\alpha}_n$ *is the following: Each healer* i *mixes over the interval* $[p_0, p_i]$ *using the distribution function* H_i *defined as*

$$H_i(p) = \frac{1}{\overline{\alpha}_i}\left(1 - \sqrt[n-j]{\frac{(1-\overline{\alpha}_j) \cdot \ldots \cdot (1-\overline{\alpha}_{n-1})}{p}}\right) \tag{2.1}$$

for $p \in [p_{j-1}, p_j] \subset [p_0, p_i]$ *with* $1 \leq j \leq n-1$. *On* $[0, p_0]$, *define* $H_i = 0$ *and, on* $[p_i, 1]$, $H_i = 1$. H_n *places an atom of size* $1 - \frac{\overline{\alpha}_{n-1}}{\overline{\alpha}_n}$ *on* 1.

The question of uniqueness of equilibrium in the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ had been pointed out as an open problem by Ireland (1993) and McAfee (1994).[14] Spiegler (2006a) proves uniqueness of equilibrium for the special cases where all healers offer the same quality and where all but one healers offer the same low quality and one healer a higher quality.

In the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$, the payoff of healer i from playing some price p while the other healers use the mixed strategies H_j is given by

$$\pi_i(p) = p\overline{\alpha}_i \prod_{j \neq i}(1 - \overline{\alpha}_j H_j(p)). \tag{2.2}$$

The intuition is as follows: In order to attract a patient and earn p, healer i has to be recommended (which happens with probability $\overline{\alpha}_i$) and has to be the cheapest healer among those who are recommended. (The probability that a competitor j is not both recommended and cheaper than p is $1 - \overline{\alpha}_j H_j(p)$.) We insert the explicit formulas for the distribution functions H_j from Proposition 2.1 into (2.2). Then we can calculate that the expected equilibrium payoff of healer i is given by

$$\pi_i(p) = \overline{\alpha}_i \prod_{j \neq n}(1 - \overline{\alpha}_j) \text{ for all } p \in [p_0, p_i].$$

[14]It is straightforward to generalize our uniqueness result to the more general demand functions considered in McAfee (1994).

From this we can deduce that the distribution functions H_i also form a Nash equilibrium in the incomplete information game:

Proposition 2.2. *The distribution functions H_i defined in (2.1) form a Nash equilibrium in the incomplete information game with qualities $\alpha_1, ..., \alpha_n$. In this equilibrium, the payoff of healer i is given by*

$$\pi_i = \alpha_i \prod_{j \neq n} (1 - \overline{\alpha}_j).$$

Note that the equilibrium strategies of the healers do not depend on the realizations of their qualities α_i. Thus the healers do not make use of their private information. The intuition is as follows: Once he got recommended to a patient, it does not matter anymore for a healer how good or bad his quality actually is. For his competitors, the exact quality of the healer plays no role either, as they do not know it: The competitors can base their strategies only on the healer's expected quality. Yet most healers (that is to say all healers $i \neq n$ if $\overline{\alpha}_n > \overline{\alpha}_{n-1}$) do incorporate their expected quality $\overline{\alpha}_i$ into their equilibrium strategies.

The next proposition establishes that the expected healers' payoffs are the same in all equilibria of the incomplete information game. Furthermore, all the equilibria are interchangeable[15]: Any two equilibria A and A' can be combined to form another equilibrium A'' by assigning to the healers their respective strategies from A or A' in an arbitrary way. This is the extent to which the equilibrium uniqueness from the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ carries over.

Proposition 2.3. *A profile of strategies $((G_1^{\alpha_1})_{\alpha_1}, \ldots, (G_n^{\alpha_n})_{\alpha_n})$ is a Nash equilibrium if and only if:*

$$E[\alpha_i G_i^{\alpha_i}(p)] = \overline{\alpha}_i H_i(p) \text{ for all } i \tag{2.3}$$

and all $G_i^{\alpha_i}$ have their support in $[p_0, p_i]$.

Furthermore, the healers expect the same payoffs in all equilibria.

Obviously, there are infinitely many equilibria since each healer i can make his

[15]See Osborne and Rubinstein (1994), p. 23.

strategy dependent on α_i in an arbitrary way as long as (2.3) is satisfied. Notably, if the noise in the α_i is rich enough, pure price strategies are possible in equilibrium.

Proposition 2.4. *Assume that the distribution functions F_i have continuous and strictly positive densities f_i on $[0,1]$. Furthermore, assume that $\overline{\alpha}_n = \overline{\alpha}_{n-1}$. Then there is a unique pure strategy equilibrium with strictly increasing price setting functions $\bar{P}_i(\alpha_i)$. Moreover,*

$$\bar{P}_i(\alpha_i) = H_i^{-1}\left(\frac{\int_0^{\alpha_i} \beta f_i(\beta) d\beta}{\overline{\alpha}_i}\right).$$

The randomness in the quality realizations allows the healers to remain unpredictable competitors even if they choose a pure pricing strategy. Note that Proposition 2.4 does not demand for much noise in the sense of a large variance. The main requirement in Proposition 2.4 is that the qualities are drawn from atomless distributions. (An atom would force the healers to mix over prices in equilibrium.) For the sake of brevity, we exclude the case of $\overline{\alpha}_n > \overline{\alpha}_{n-1}$. In that case all our arguments still go through but, because of the atom in H_n, $\bar{P}_n$ reaches the value 1 already for some $\alpha_n < 1$ and then stays constant. Intuition clearly favors the monotone price equilibrium over the other equilibria: It is natural to assume that a healer with a higher quality charges a higher price. The same selection criterion among equilibria is applied in the standard literature on auctions.

2.4 Welfare

In this section we focus on the symmetric case where the qualities of all healers i are independently drawn from the same distribution $F_i = F$. We assume that F has a continuous and strictly positive density f. Denote by $\overline{\alpha}$ the mean of F. Hence $\overline{\alpha}$ is also the expected quality of a randomly drawn healer. We study welfare in the monotone strategy equilibrium and show that it deteriorates an n gets large. At the end, we compare our welfare results to those of the standard model where patients act rationally, but hold only incomplete information about the healers' qualities.

By the results of the previous section we have the following monotone pricing equilibrium:

Proposition 2.5. *There is a unique equilibrium in monotonically increasing price strategies. In this equilibrium, each healer i uses the price setting function*

$$\bar{P}(\alpha_i) = \left(\frac{1-\overline{\alpha}}{1-\int_0^{\alpha_i} \beta f(\beta) d\beta} \right)^{n-1}.$$

The expected equilibrium payoff of healer i is given by

$$\pi_i = \alpha_i (1-\overline{\alpha})^{n-1}.$$

Conditional on his quality realization, each healer plays a pure pricing strategy. Looking at the payoffs, we see that each healer i earns in expectation only his expected maxmin payoff – the payoff he can earn for sure no matter what prices his competitors choose: With an expected probability of $(1-\overline{\alpha})^{n-1}$, all his competitors are not recommended. With a probability of α_i, healer i is recommended. Hence with an expected probability of $\alpha_i(1-\overline{\alpha})^{n-1}$, healer i is the only one who is recommended. Thus by charging a price of 1, healer i can secure an expected payoff of π_i to himself.

Denote by θ_n the healers' aggregate payoffs in a market with n healers:

$$\theta_n = \sum_i E[\pi_i] = n\overline{\alpha}(1-\overline{\alpha})^{n-1}. \quad (2.4)$$

As one easily sees from (2.4), the healers' aggregate payoffs may initially increase in n if $\overline{\alpha}$ is not too large but eventually converge to zero as n increases. The intuition is as follows: With few healers and low qualities, competition is soft. Only few patients are attracted by several healers, and many patients do not get recommended to any healer at all. A new healer entering the market may attract most of his patients from the group of patients that would otherwise stay at home. Thus the new healer does not strengthen competition much. Yet if more and more healers enter the market, even with low qualities, more and more patients get recommended to several healers. Then price competition gets more and more severe, driving the healers' payoffs down.

We have found that in expectation the healers' payoffs go to zero as n gets large. But does that mean that the patients are better off the more healers enter the market?

At least, in the limit, patients do not have to pay anything for the healers' services. Yet it turns out that also patients fare badly: Recall that a patient always consults the cheapest healer who is recommended to him because he thinks all recommended healers are of the same high quality and just differ in prices. Yet as the healers apply monotone price-setting strategies, by picking the cheapest the patient also picks the worst of all recommended healers. Since the distribution function of qualities F has support on the whole interval $[0, 1]$, as many healers enter the market there is - with high probability - also a considerable amount of very low quality healers, some of which get recommended. Making this reasoning precise, one can see that as n gets large, overall welfare converges to zero: No patient gets cured in the limit.

Proposition 2.6. *Denote by γ_n the expected social welfare, i.e. the expected proportion of patients cured, in the monotone equilibrium with n healers. Then*

$$\gamma_n = n \int_0^1 \alpha^2 \left(1 - \int_0^\alpha \beta f(\beta) d\beta\right)^{n-1} f(\alpha) d\alpha \tag{2.5}$$

and

$$\lim_{n\to\infty} \gamma_n = 0.$$

Like the healers' payoff θ_n, welfare γ_n is an expected value taken over the qualities α_i. The intuition behind formula (2.5) is the following: Consider the expected quality of the treatment chosen by a patient: A quality α is chosen if a) the healer offering that quality is recommended and b) all his competitors are either not recommended or are charging a higher price. The probability of a) is α. The probability of b) is $\left(1 - \int_0^\alpha \beta f(\beta) d\beta\right)^{n-1}$ as charging a higher price is equivalent to offering a higher quality in the monotone equilibrium.

Welfare may already decrease for a quite low number of healers. This is relevant as it seems much more natural to think of a patient receiving an anecdote on each healer if the market is not too large. However, as pointed out in Spiegler (2006a), n can also be interpreted as the number of healers a patient gets a report on in a market with very many healers. Then n would be a measure of patients' awareness.[16] The following example demonstrates that the effects at work are not only limit results

[16]We discuss this further in Section 2.5.

but already play an important role already for moderate numbers of healers. Figure 2.1 depicts welfare for the special case that qualities are uniformly distributed on $[0, 1]$, i.e. $F(\alpha) = \alpha$:

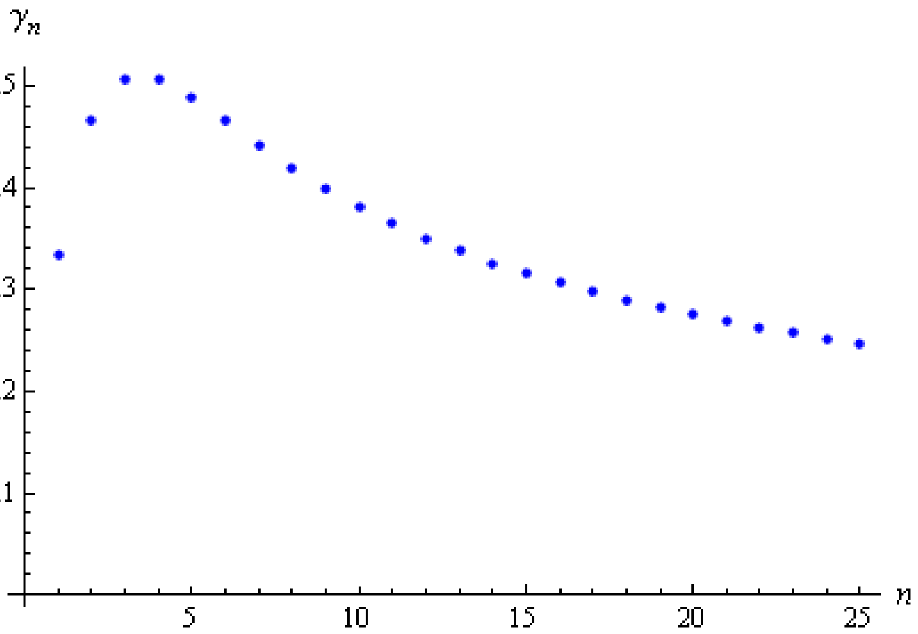

Figure 2.1: Expected social welfare γ_n for $\alpha_i \sim U[0, 1]$

Figure 2.1 shows that welfare is maximized for $n = 3$ healers. With three or four healers, the average quality a patient receives is slightly larger than the quality of the average healer $\overline{\alpha} = 1/2$, as better healers are more often recommended. Afterwards, less and less patients get cured, as most patients get recommended to several healers and then, led by price-comparison, end up with a low quality healer.

Figure 2.2 depicts the sum of the n healers' expected payoffs θ_n. θ_n is maximized with one or two healers where it equals $1/2$. Afterwards θ_n decreases quickly.

Figure 2.3 depicts patients' aggregate surplus which is the difference between overall welfare γ_n and healers' surplus θ_n. The healers' surplus θ_n decreases considerably faster than the proportion of patients cured γ_n. Hence the patients' surplus is largest

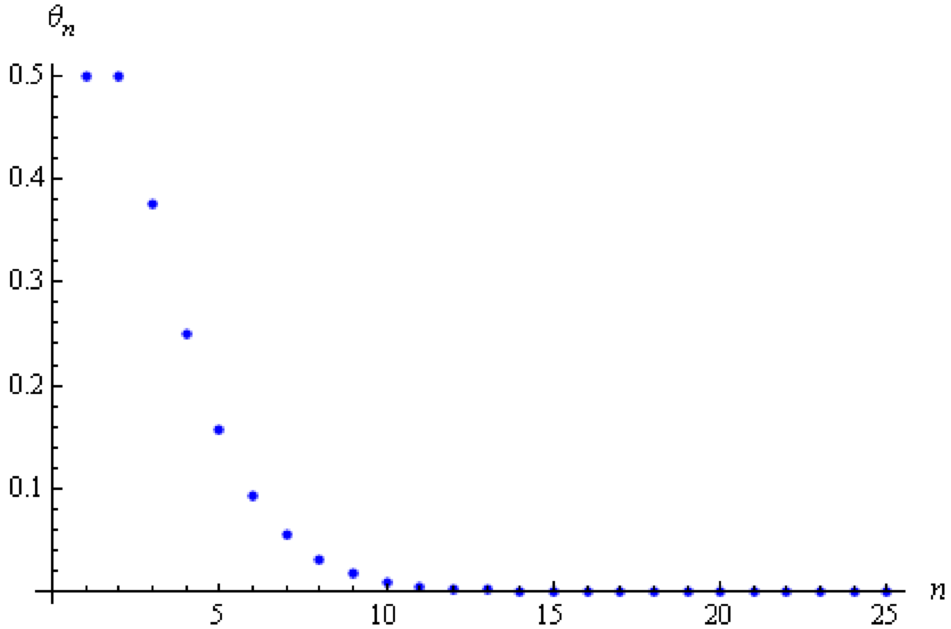

Figure 2.2: Expected healers' aggregate surplus θ_n for $E[\alpha_i] = \frac{1}{2}$

at an intermediate market size of $n = 8$. As n gets larger, the patients' surplus decreases, driven by the decreasing average quality received γ_n. Note that with one or two healers the patients' surplus is negative: In monopoly, the healer attracts all patients to whom he is recommended. He then charges a price of 1 for a treatment of expected quality $\overline{\alpha} = \frac{1}{2}$. In duopoly, patients are more often recommended to the healer with the higher quality (who offers in expectation a healing probability of $\frac{2}{3}$). Yet, as competition is weak, prices are still quite high. Patients' surplus increases, but remains negative.

From Proposition 2.3 it is clear that for the healers' expected payoffs it does not matter which equilibrium they play. The patients' health, however, varies across equilibria. We close this section by considering welfare in the equilibrium from Proposition 2.2 where all healers play the same mixed strategy H (independent of their quality realization). In this equilibrium, welfare converges to some value strictly above $\overline{\alpha}$. The intuition is as follows: If all healers play the same price

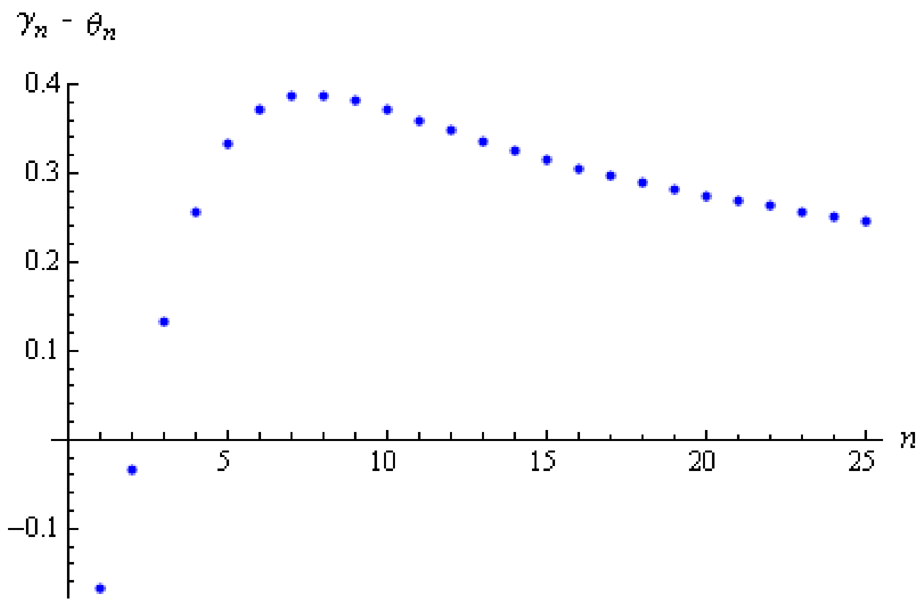

Figure 2.3: Expected patients' aggregate surplus $\gamma_n - \theta_n$ for $E[\alpha_i] = \frac{1}{2}$

strategy, they all have the same probability that they are the cheapest and thus attended by the patients - given that they are recommended. But since better healers are recommended more often a patient will get an above average treatment in expectation when there are sufficiently many healers.

Remark 2.1. *Consider the symmetric case where all α_i are drawn from the same distribution F. Consider the mixed equilibrium where all healers i apply the same price strategy H as defined in Proposition 2.1. Let $\widetilde{\gamma}_n$ denote the expected welfare in this equilibrium. Then*

$$\widetilde{\gamma}_n = \frac{E[\alpha^2]}{E[\alpha]}(1 - (1 - E[\alpha])^n) \tag{2.6}$$

where α has distribution F.

By Jensen's inequality the limit for n to infinity, $E[\alpha^2]/E[\alpha]$, is strictly greater than $E[\alpha] = \overline{\alpha}$ (unless F is deterministic). The second factor in (2.6), $1 - (1 - E[\alpha])^n$, is the probability that a patient gets at least one recommendation and turns to the market for healers. The first factor of (2.6), $\frac{E[\alpha^2]}{E[\alpha]}$, is the expected quality a patient receives given he does not stay at home. This factor does not depend on the number of healers n as prices do not reveal anything about healers' qualities. Yet it depends on the variance of F:

$$\frac{E[\alpha^2]}{E[\alpha]} = \frac{Var[\alpha]}{E[\alpha]} + E[\alpha].$$

A higher variance is beneficial, as it increases the average quality of a recommended set of healers.

Note that $\widetilde{\gamma}_n$ describes also the proportion of patients cured in a standard market model with rational, incompletely informed patients. In such a model, prices cannot transmit any information so that there must be pooling over the prices in the equilibrium of the pricing game: It is rational for patients to use the word of mouth as a selection device, i.e. to update in a Bayesian way the initial beliefs about the healers' qualities. But it would be irrational to take the prices into account. Hence we obtain the same welfare result $\widetilde{\gamma}_n$.

Finally, note that $\widetilde{\gamma}_n$ also describes the proportion of patients cured in a situation where patients utilize anecdotal reasoning but where a fixed price is exogenously prescribed to the healers.

2.5 Discussion

In this section we discuss several directions for extensions and show robustness of our results.

Awareness of patients:

Throughout the analysis, we assumed that n is the number of healers in the market. Yet as we outlined before[17], we can reinterpret n as the number of healers a patient samples, hence as a measure of patients' awareness. Especially in a large market, patients might only sample a fraction of all healers − and sampling intensity may

[17] Compare also Spiegler (2006a).

be heterogeneous among the patients. We can incorporate this heterogeneity into our model by taking n as an integer-valued random variable. Our results prove to be robust to this extension: For simplicity, focus on the symmetric case $F_i = F$. Let us define

$$\rho = E[(1-\overline{\alpha})^{n-1}],$$

where the expectation is taken over n. With the same argument as in Section 2.3, healer i's expected equilibrium payoff is then given by

$$\pi_i = \alpha_i \rho.$$

There exists again a symmetric mixed strategy equilibrium H which does not depend on the realizations of the α_i. The support of H is the interval $[\rho, 1]$ and H is given implicitly as the solution of

$$\rho = pE[(1-\overline{\alpha}H(p))^{n-1}].$$

From this mixed strategy equilibrium it is straightforward to construct a monotone strategy equilibrium the same way as for deterministic n. Hence again, too much competition turns out to be detrimental to welfare.[18]

Several anecdotes:

Another plausible variation of our model would be a model where patients gather several anecdotes on each healer. For example, this might be modeled in a way that more anecdotes are gathered about more popular healers. While the equilibria of such generalized models usually preclude an explicit solution, it is mostly straightforward to see that the main arguments of our analysis carry over: As long as each healer has a certain chance of being perceived as the very best healer by some patients, healers will be unwilling to engage in harsh price competition leading to the type of pricing behavior studied in the previous sections.[19]

[18]As an aside, from applying Jensen's inequality we see that $\rho > (1-\overline{\alpha})^{E[n]}$. The healers' equilibrium payoff is thus higher than in the model with a deterministic number of $E[n]$ healers. In this sense, the heterogeneity in n is beneficial for the healers.

[19]See also the discussions in Spiegler (2006a) and Szech (2010).

Note that with one anecdote per healer, patients face a clear dichotomy: Healers get subdivided into good and bad. This simplicity has been pointed out by Fagerlin et al. (2005) as one of the factors that make anecdotes so attractive to rely on.

Endogenous qualities:

One assumption of our model which may seem rather strong is that the healers' qualities are exogenous random variables. It is plausible that while healers may not be able to fully control their qualities, they can influence them at least to some extent. With the application to advertising in mind, Ireland (1993) and McAfee (1994) analyze extended games where the healers choose their qualities themselves before pricing takes place. They focus on the case of complete information about qualities. Szech (2010) adds the welfare analysis under the S(1) interpretation.[20]

It is straightforward to combine the analysis of endogenous qualities for the complete information case with our incomplete information model: Recall that ex ante, a healer's expected payoff only depends on his expected quality (and not on any further properties of the distribution function). Thus the analysis of the complete information case transfers immediately to a model where healers choose between several distributions from which their quality is drawn (in the sense of, e.g., choosing between different specializations): When choosing between different distribution functions, healers only take into account the expected quality.

From the analysis of the complete information case it is known that healers typically choose much lower than socially optimal qualities, to make competition softer and hence raise revenues.[21] Also, if the best possible qualities are not too low, there will

[20] Welfare results differ markedly from welfare in the advertising interpretation. Under the interpretation of advertising, firms differ only in advertising activities, but not in service qualities. α_i describes solely the probability with which a consumer gets aware of firm i. Hence any consumer ending up at a firm receives the same gross utility.

[21] See Ireland (1993) and McAfee (1994) for the explicit form of the pure quality-setting equilibria, respectively, with and without costs of quality-setting. Szech (2010) adds a more detailed analysis of mixed strategy equilibria and shows that welfare decreases for larger numbers of healers under the interpretation of anecdotal reasoning.

be much difference in the qualities offered by the best and the worst healers. Hence endogenous quality choice creates the situations of varying qualities to which our bad welfare results apply.

2.6 Conclusion

"It is unwise to pay too much, but it's worse to pay too little. When you pay too little, you sometimes lose everything because the thing you bought was incapable of doing the thing you bought it to do." This recommendation is attributed to the social thinker John Ruskin (1819-1900). Indeed, in markets where qualities are not easy to grasp, competition among firms may lead to consumers ending up with poor qualities, as they focused too much on price differences. Our model shows that even if patients try to get an idea about the qualities in a market, they likely end up with a bad quality.

The assumption that consumers rely on word of mouth captures empirical findings from the psychological and economic literature. Recently, also medical research puts a lot of attention to the phenomenon that lay people tend to prefer to rely on anecdotes even if statistical evidence is available and presented in an appealing, easy-to-grasp way. This fact has even led to recommendations of incorporating personal stories into evidence-based results, such that patients may be more willing to adhere to statistical recommendations.[22]

Assuming that consumers apply anecdotal reasoning, our model generates very different predictions than those made by standard market models. Stronger competition turns out to be detrimental to welfare. Recent surprising results from the US medical system support this conclusion, showing that the non-competing Veterans Health Administration often provides higher quality services than the competitive health systems prevalent in the US.

[22] Compare e.g. Glenton et al. (2006).

Generally, we believe that more research is needed to explore the interplay between perfectly rational firms and boundedly rational consumers following behavioral, possibly market-specific rules instead of perfectly rational thinking.

Chapter 3

Rebates in a Bertrand Game

We study a price competition game in which customers are heterogeneous in the rebates they get from either of two firms. We characterize the transition between competitive pricing (without rebates), mixed strategy equilibrium (for intermediate rebates) and monopoly pricing (for larger rebates).
In the mixed equilibrium, a firm's support consists of two parts: (i) aggressive prices that can steal away customers from the other firm; (ii) defensive prices that can only attract customers who get the rebate. Both firms earn positive expected profits.
We show that counter-intuitively, for intermediate rebates, market segmentation decreases in rebates.

3.1 Introduction

The presumably simplest and in this sense most fundamental model on rebates is not yet fully analyzed. Klemperer (1987a, Section 2) studies the situation where two firms with equal and constant marginal costs compete in prices. He frames the example as one of the airline industry where rebates are given. Each customer has to pay the full price at one firm if he buys there, but only the reduced price if he buys from the other firm. Klemperer shows that in equilibrium each customer buys from the firm where he can get the rebate and the reduced price equals the monopoly price. Therefore, firms yield monopoly profits in their segments.

Despite its simplicity the model is not yet fully analyzed: unless rebates are sufficiently high, an equilibrium in pure strategies fails to exist. This is one of the reasons why the literature, starting with Klemperer (1987a), has attached further components to the model to guarantee existence of pure strategy equilibria.[1,2]

We analyze the "innocent" model without any restriction on the size of the rebates. We show that when customers differ in the rebates they can get, both firms earn positive expected profits. The intuition is as follows: first, a firm which offers a rebate will not charge a price near or below marginal costs, because a very low price would at least attract customers who get a rebate and pay a negative net price. Due to the low price, the loss from this customer group cannot be compensated by customers who do not get a rebate from this firm. Second, given that the former firm charges only prices well above marginal costs, the other firm can set a price so that it yields a profit. Third, because the other firm earns in expectation a profit, it must charge prices well above marginal costs, too. This also enables the former firm to earn a profit. Note that when all customers get the same rebate at the same firm or when no rebates are set, the well-known Bertrand paradox arises: both firms set net prices equal to marginal costs and earn zero profits.

In the main part of our analysis, we focus on unit-demand. The equilibrium is characterized by three different regimes: first, when rebates are small, the Nash equilibrium is in mixed strategies without mass points. Second, for intermediate levels of rebates the equilibrium is still in mixed strategies but there is a mass point at the upper end of the support. Third, when rebates are high, the equilibrium is in pure strategies, just as in Klemperer (1987a). In the first two regimes firms mix between two types of strategies: an aggressive one and a defensive one. Either a firm charges low prices, attracts all customers of its home base for sure and with some probability attracts the other customers as well. Or a firm charges high prices, thus risking to lose the customers of its home base, but earns a high payoff if it

[1]Klemperer (1995, footnote 7): "Pure-strategy equilibrium can be restored either by incorporating some real (functional) differentiation between products (Klemperer (1987b)), or by modelling switching costs as continuously distributed on a range including zero (...) (Klemperer (1987a))." Banerjee and Summers (1987) consider a sequential price setting to circumvent mixed strategies. Also Caminal and Matutes (1990) analyze a setting with real differentiation.

[2]Mixed strategy equilibria often arise in oligopoly pricing models. For example, in Padilla's (1992) dynamic setting with myopic customers; in Deneckere, Kovenock, and Lee (1992) who analyze a game with loyal customer and without rebates; in Allen and Hellwig's (1986, 1989, 1993) Bertrand-Edgeworth models, where capacity-constrained firms choose prices.

still attracts them. For the case where firms mix without atoms we show that the probabilities of attacking and defending stand in the celebrated golden ratio.

Furthermore, we show that – counter-intuitively at first sight – market segmentation may decrease in rebates. This happens when rebates reach an intermediate level where the customers' limited willingness to pay starts to affect the firms' pricing behavior. From this level of rebates on, firms have to concentrate some mass of their pricing strategy into an atom at the upper end of their price interval. At that price, the firm can only attract its home-base if the other firm does not attack. Yet as the other firm still attacks with some (though diminishing) probability, market segmentation has to decrease. Then, for larger rebates, the attack probability of the other firm gets so low that market segmentation increases again and converges to full segmentation of the market.

We also study the normative aspects of our model. As discussed above, rebates increase firms' profits when customers differ in the rebates they can get. We show that under more general demand functions, rebates also lead to lower customer and total welfare results. Note that customers face a coordination problem: customers are collectively worse off when there are rebate systems, but individually they are better off when they participate in a system than when they do not.

In our analysis we take the rebates customers get as given and concentrate on the price setting behavior of the firms. This can be justified for the following reasons. First, this approach may be a good description of the short run behavior of firms where the rebate system is established and cannot be overturned. Second, in some lines of industries, like the airline industry, several firms have a common rebate system. Then a firm can hardly change rebates when it decides about its prices. Third, note that we obtain the result that with rebates and when customers differ in the rebates they can get, firms earn positive profits instead of zero profits. Therefore, even though we do not model how firms set rebates, we predict that firms have an incentives to offer rebates and set them in a way that customers are heterogeneous in the aforementioned sense.

Bester and Petrakis (1996) study the effects of coupons/rebates on price setting in a one period model where firms can target certain customers. In equilibrium, each firm sends coupons to customers who live in the "other city". Therefore, unlike in our model, coupons reduce the firms' profits. For similar models, see Shaffer and Zhang (2000) and Chen (1997). Note that despite some similarities our model is not a reinterpreted model of spatial competition: in our model, firms care which customers buy from them because customers pay different net prices.

The chapter proceeds as follows. In Section 3.2, we introduce the model. In Section 3.3, we solve the equilibrium explicitly for the case of unit-demand. In Section 3.4, we characterize the equilibrium for a large variety of demand functions. In Section 3.5 we study the welfare effects and explore the customers' coordination problem. In Section 3.7, we offer a concluding discussion. The proofs are relegated to the appendix.

3.2 The Model

We analyze a market with two firms and a continuum of customers. The customers are of one of two types: a mass m_1 of customers gets a fixed rebate $r_1 \geq 0$ at firm 1 and no rebate at firm 2. We call this group of customers the "home base" of firm 1. A mass m_2 of customers gets no rebate at firm 1 and a fixed rebate $r_2 \geq 0$ at firm 2. Each customer wants to buy exactly one object, for which his valuation is $\overline{p}$. Both firms produce these objects at costs which are normalized to zero. Firms engage in price competition: customers buy from the firm where they have to pay the lower net price (i.e., price minus rebate), provided that this net price is below the valuation. In Section 3.4, we will extend our analysis to much more general demand functions and to situations where not all customers get rebates.

Let us start with an intuition why in this game the Bertrand Paradox breaks down, i.e., why firms must earn positive profits. When a firm offers a rebate, it has to charge gross prices well above zero to yield no loss. This enables the other firm to earn a positive profit. Hence, the other firm also charges in equilibrium prices well above zero which in turn allows the former firm to earn a positive profit, too.

Klemperer (1987a) obtains essentially the following partial result:

Proposition 3.1. *Suppose $m_1 > 0$ and $m_2 > 0$. Then, if r_1 and r_2 are sufficiently large, each firm earns monopoly profits in its market segment.*

In the next sections, we explore what happens if the rebates are not that high, such that the above pure strategy equilibrium does not exist.

3.3 Characterization of Equilibria

In the following, we see that if rebates are moderate, a mixed strategy equilibrium arises. Section 3.3.1 characterizes the mixed strategy equilibrium for the case where the rebates are small enough to ensure that $\overline{p}$ does not interfere with the firms' pricing strategies: in this case, each firm i mixes over strictly positive prices that are strictly lower than $\overline{p} + r_i$. Section 3.3.2 gives a complete characterization of the transition between pure and mixed strategy equilibrium for the symmetric case $r_i = r_j$ and $m_i = m_j$.

3.3.1 Atomless Pricing for Moderate Rebates

Denote by F_i the distribution function underlying the mixed price-setting strategy of firm i, and let π_i be firm i's equilibrium payoff. Then in equilibrium it must hold that for all $p \in suppF_i$

$$\pi_i = m_i(p - r_i)(1 - F_j(p - r_i)) + m_j p(1 - F_j(p + r_j)). \tag{3.1}$$

The equilibrium distributions we identify are characterized as follows: firms mix between two types of strategies – an aggressive one and a defensive one. Either a firm charges low prices, attracts all customers of its home base for sure and with some probability attracts the other customers as well. Or a firm charges high prices, thus risking to lose the customers of its home base, but earns a high payoff if it still attracts them. Formally, F_i can be written as $q_i A_i + (1 - q_i) D_i$ where A_i and D_i are distribution functions and $q_i \in [0, 1]$. We call $q_i \in [0, 1]$ the "attack probability", as only a firm playing the aggressive strategy may steal away customers of the other firm's home base: A_i (the aggressive strategy) and D_i (the defensive strategy) have distinct supports $[\underline{a}_i, \overline{a}_i]$ and $[\underline{d}_i, \overline{d}_i]$ with $\overline{a}_i \leq \underline{d}_i$.

Figure 3.1 schematically depicts the supports of the two firms' strategies in an example with $r_i > r_j$. Given this decomposition of the firms' strategies, (3.1) becomes for small p, that is, for $p \in [\underline{a}_i, \overline{a}_i]$,

$$\pi_i = m_i(p - r_i) + m_j p(1 - q_j)(1 - D_j(p + r_j)) \tag{3.2}$$

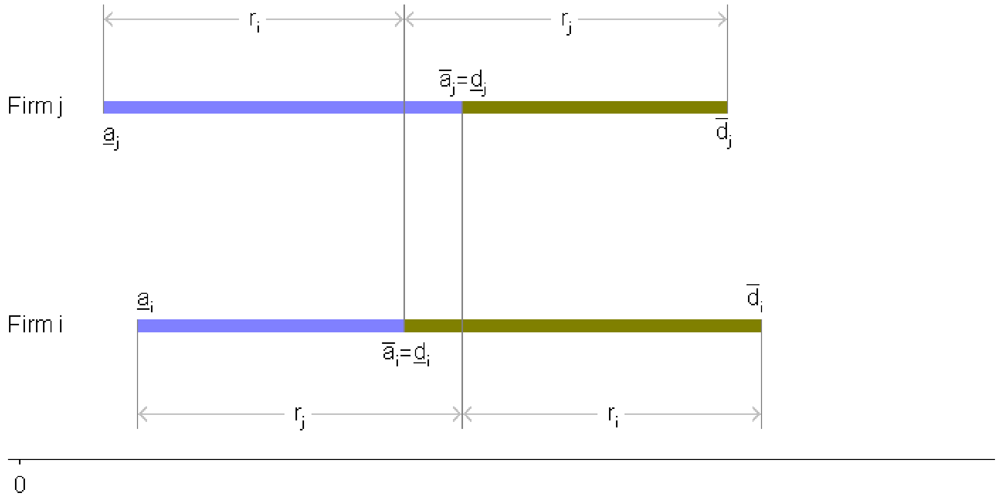

Figure 3.1: The boundaries of the strategy supports.

and for larger p, that is, for $p \in [\underline{d}_i, \overline{d}_i]$,

$$\pi_i = m_i(p - r_i)(1 - q_j A_j(p - r_i)). \tag{3.3}$$

Our first main result explicitly gives an equilibrium for the case in which the maximal willingness to pay, $\overline{p}$, is sufficiently large not to interfere with the firms' pricing strategies.

Proposition 3.2. *Assume that $\overline{p}$ is sufficiently large, i.e., $\overline{p} > \max\{\overline{d}_1 + r_1, \overline{d}_2 + r_2\}$ where $\overline{d}_j$ is defined below. Then an equilibrium is given as follows: equilibrium attack probabilities q_j and equilibrium payoffs π_j are*

$$q_j = \frac{m_i^2 + m_i m_j + m_j^2 - \psi(m_i, m_j)(m_i^2 - m_i m_j + m_j^2)}{2m_j^2} \tag{3.4}$$

and

$$\pi_j = \frac{(\psi(m_i, m_j) + 1)m_i m_j - (\psi(m_i, m_j) - 1)m_j^2}{2m_i} r_i + \frac{(\psi(m_i, m_j) - 1)m_j}{2} r_j,$$

where

$$\psi(m_i, m_j) = \sqrt{\frac{m_i^2 + 3m_i m_j + m_j^2}{m_i^2 - m_i m_j + m_j^2}}.$$

The equilibrium strategies are compositions of the defensive strategy

$$D_j(p) = 1 - \frac{\pi_i - m_i(p - r_i - r_j)}{m_j(p - r_j)(1 - q_j)} \tag{3.5}$$

and the aggressive strategy

$$A_j(p) = \frac{1}{q_j}\left(1 - \frac{\pi_i}{m_i p}\right), \tag{3.6}$$

with supports given by

$$\underline{d}_j = \frac{\pi_i + m_i(r_i + r_j) + r_j m_j(1 - q_j)}{m_i + m_j(1 - q_j)}, \quad \overline{d}_j = \frac{\pi_i}{m_i} + r_i + r_j$$

and

$$\underline{a}_j = \frac{\pi_i}{m_i}, \quad \overline{a}_j = \frac{\pi_i}{m_i(1 - q_j)}.$$

Furthermore, supports of the equilibrium strategies are connected, i.e., $\underline{d}_j = \overline{a}_j$*. The defensive strategy of firm* i *is a downward shift by* r_j *of the defensive strategy of firm* j*, i.e.,* $D_j(p + r_j) = A_i(p)$*.*

Note that one can immediately calculate these functions D_j and A_j using (3.2) and (3.3). The fact that the aggressive strategy of player i is identical, up to a shift by r_j, to the defensive strategy of player j, has the following consequence: given that firm i attacks and firm j defends, there is a probability of 1/2 that all customers end up at firm i. With the counter-probability, all customers buy at their home firm.

While the dependence of the equilibrium on the group sizes m_i and m_j is a bit more complex, the dependence on the rebates is very simple: the attack probabilities q_j are independent of the rebates. The equilibrium payoffs are linearly increasing in both rebates. The function ψ which determines equilibrium payoffs and attack probabilities is a symmetric function which only depends on the ratio of m_i and m_j. It takes its maximum value of $\sqrt{5}$ for $m_i = m_j$ and decreases to the value 1 as m_i/m_j goes to 0 or ∞.

To see how asymmetries in the attack probabilities are linked to asymmetries in group sizes observe from (3.4) that the following relation holds:

$$q_i m_i^2 = q_j m_j^2.$$

Intuitively, a firm who gives rebates only to few customers is more inclined to set small prices targeting customers who get a rebate from the other firm.

To illustrate the proposition, consider the case $m_i = m_j = 1$. Then the equilibrium is given by

$$q_i = q = \frac{3 - \sqrt{5}}{2} \approx 0.382 \text{ and } \pi_i = r_j + (1 - q)r_i.$$

Note that this implies that the probabilities of attacking and defending stand in the celebrated golden ratio, i.e.,

$$\frac{1-q}{q} = \frac{1+\sqrt{5}}{2}.$$

To get some intuition for the equilibrium – and also for the occurrence of the golden ratio – let us consider the special case $r_i = r_j = r$. Let us assume that in equilibrium both players mix with some atomless strategy over an interval of length $2r$, i.e., $[\underline{a}, \underline{a} + 2r]$. Let q be the equilibrium attack probability, i.e., the probability mass in the lower half $[\underline{a}, \underline{a} + r]$.

We demonstrate now how these assumptions uniquely determine equilibrium values of $\underline{a}$ and q and equilibrium payoffs. Let us compare the firms' expected payoffs from playing prices $\underline{a}$, $\underline{a} + r$ and $\underline{a} + 2r$ which in equilibrium must be identical. Note first that by playing a price of $\underline{a} + r$, a firm attracts all customers from its home base, but no customers from the opponent's home base. Thus

$$\pi(\underline{a} + r) = \underline{a} + r - r = \underline{a}.$$

Compare to this playing a price of $\underline{a}$. Then our firm still attracts its home base with certainty but payments from the home base decrease by r. Yet unlike before, our firm receives $\underline{a}$ from the customers in the other firm's home base as well, provided that the other firm plays a price above $\underline{a} + r$ which happens with probability $1 - q$. Thus from $\pi(\underline{a}+r) = \pi(\underline{a})$ we can conclude that advantages and disadvantages from switching from $\underline{a} + r$ to $\underline{a}$ must cancel out in equilibrium, i.e.,

$$r = (1 - q)\underline{a}. \tag{3.7}$$

Now consider the payoff from playing a price of $\underline{a} + 2r$. In this case our firm attracts its home base only if the other firm plays a price above $\underline{a} + r$ which happens with probability $1 - q$. We hence get

$$\pi(\underline{a} + 2r) = (1 - q)(\underline{a} + 2r - r) = (1 - q)(\underline{a} + r).$$

As $\pi(\underline{a} + 2r)$ and $\pi(\underline{a} + r)$ must be identical in equilibrium, we get

$$\underline{a} = (1 - q)(\underline{a} + r). \tag{3.8}$$

Now let us compare (3.7) and (3.8). From these two equations we see that the ratio between r and $\underline{a}$ is the same as the ratio between $\underline{a}$ and $\underline{a} + r$. This is exactly the defining property of the golden ratio, implying that

$$\frac{\underline{a}}{r} = \frac{1+\sqrt{5}}{2}$$

and thus by (3.7)

$$q = \frac{3-\sqrt{5}}{2}.$$

3.3.2 From Bertrand to Monopoly

So far we have analyzed the cases of sufficiently large and of sufficiently small rebates, giving rise to, respectively, a pure strategy equilibrium in $\overline{p}+r$ or a mixed strategy equilibrium. For the symmetric case, we now round out the analysis by characterizing the equilibrium also for intermediate values of r. This equilibrium is composed of an atom in $\overline{p}+r$ and mixing below this price. A gap arises between the supports of the aggressive and the defensive strategies. The transition between the different types of equilibria is continuous in r:

Proposition 3.3. *Assume $m_i = m_j = 1$, $\overline{p} = 1$ and $r_1 = r_2 = r$.*

(i) For $r \leq r^* := \frac{3-\sqrt{5}}{2}$, Proposition 3.2 characterizes an equilibrium with $q = \frac{3-\sqrt{5}}{2}$ and $\pi = (2-q)r$.

(ii) If $r^* \leq r \leq 1$, an equilibrium is given as follows: both firms play the aggressive strategy $A(p)$ with probability q^A, the defensive strategy $D(p)$ with probability q^D and a price of $1+r$ with the remaining probability. The probabilities q^A and q^D and the equilibrium payoffs π are given by

$$q^A = 1-\sqrt{r}, \quad q^D = 1-r \text{ and } \pi = \sqrt{r}.$$

The distribution functions A and D are given by

$$A(p) = \frac{1}{q^A}\left(1-\frac{1-q^A}{p}\right) \text{ and } D(p) = \frac{1}{q^D}\left(1-q^A-\frac{1-q^A-p+2r}{p-r}\right).$$

The supports of A and D are defined through

$$\underline{a}_j = \sqrt{r}, \qquad \overline{a}_j = 1,$$

and

$$\underline{d}_j = \sqrt{r}+r, \qquad \overline{d}_j = 1+r.$$

(iii) If $r \geq 1$, a pure strategy equilibrium arises where both firms set a price of $1+r$. Each firm earns an equilibrium payoff of 1.

It is straightforward to generalize Proposition 3.3 to $m_i = m_j \neq 1$ and $\overline{p} \neq 1$. Furthermore, it is easy to verify that Cases (i) and (ii) coincide for $r = \frac{3-\sqrt{5}}{2}$. Likewise, for $r = 1$, the equilibrium of Case (ii) degenerates to an atom in $1+r=2$.

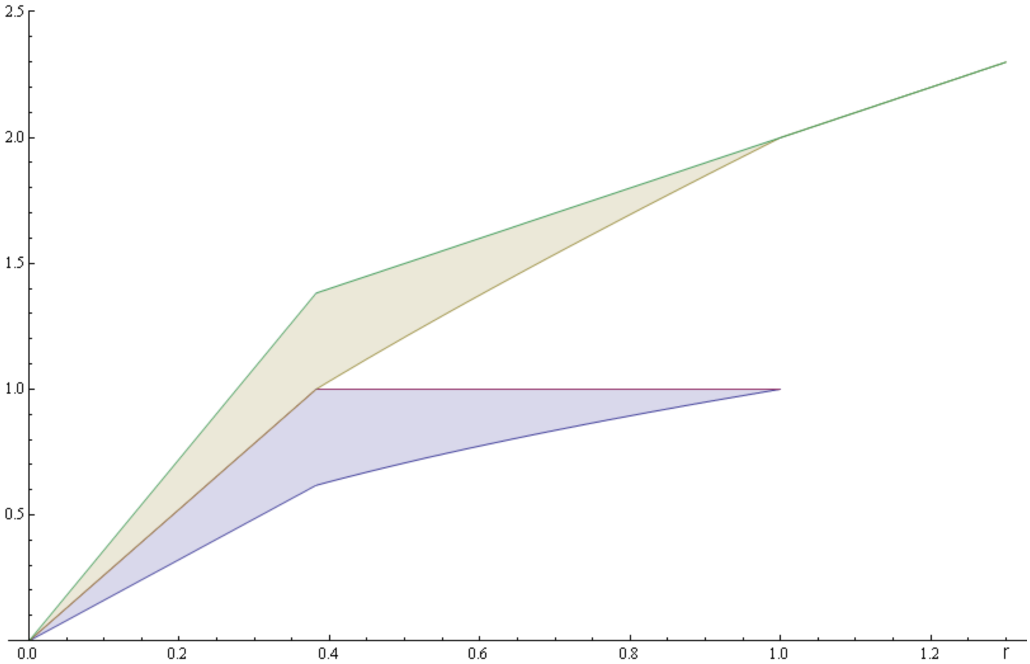

Figure 3.2: The strategy supports $\overline{d} \geq \underline{d} \geq \overline{a} \geq \underline{a}$ as functions of r.

Figures 3.2 and 3.3 illustrate Proposition 3.3. The upper quadrangle in Figure 3.2 pictures the support of the firms' defensive strategy in dependence on r. The upper bound corresponds to $\overline{d}$, the lower bound to $\underline{d}$. The lower quadrangle depicts the support of the aggressive strategy, where the upper and lower bound correspond to $\overline{a}$ and $\underline{a}$, respectively. Up to $r^* \approx 0.382$, the curves are the same as in the case of unrestricted willingness to pay. Yet once the curve $\overline{d}$ reaches the value $1 + r^*$, the limited willingness to pay of the customers kicks in: from there on, $\overline{d}$ increases less, and stays always equal to $1 + r$, the maximal willingness to pay of the home base customers. Firms put an atom on $\overline{d}$ from the kink onwards. The distance between $\overline{a}$ and $\overline{d}$ is always r, as is the distance between $\underline{a}$ and $\underline{d}$. That is, r is the maximal markup a firm can charge from its home base. The pricing strategies converge to the case of a segmented market with monopolistic prices as r approaches 1.

Figure 3.3 shows the distribution functions of the firms' pricing strategies for different values of r ($r = 0, 0.2, 0.4, \dots, 1$). We see the interpolation between competitive pricing ($r = 0$), where firms set prices of 0, and full segmentation (for $r = 1$), where both firms set a price of $1 + r = 2$ with certainty. For $r > r^*$, the pricing strategies have a gap between the aggressive and the defensive prices, corresponding to the constant part in the distribution functions. The mass of the atom corresponds to the size of the jump in the distribution functions. For $r = 0.2 < r^*$, the kink in the curve marks the boundary between aggressive and defensive pricing.

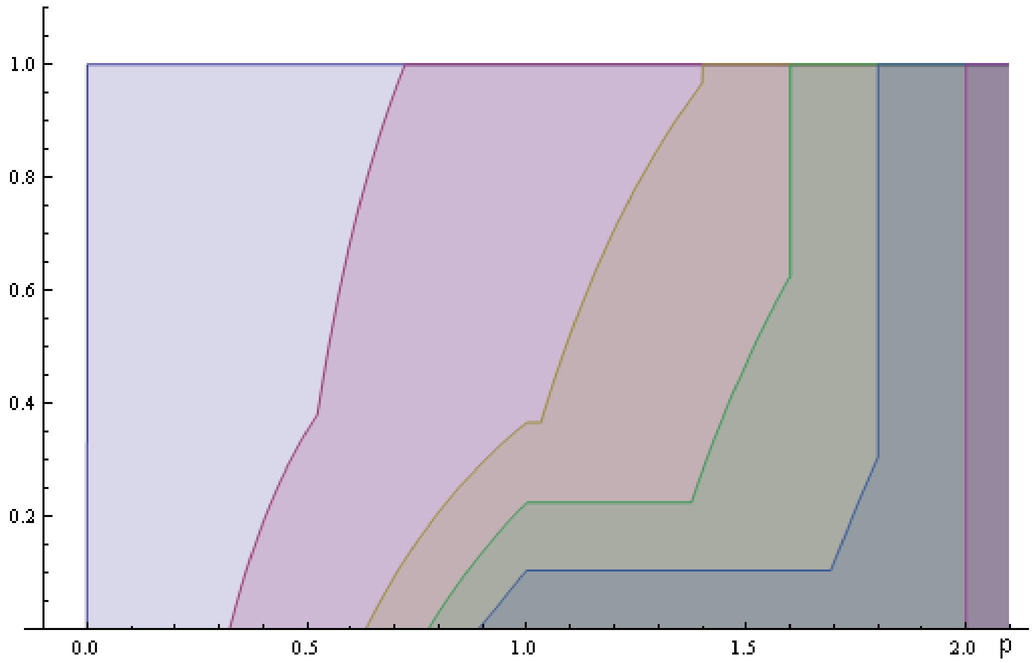

Figure 3.3: The pricing strategy $F(p)$ for $r = 0, 0.2, 0.4, \ldots, 1$.

The firms' profits increase linearly in r for r low and sub-linearly for intermediate r. When $r \geq 1$, profits stay constant in r. Intuitively, once the market is fully segmented, firms cannot earn more than monopoly profits, hence they do not gain from higher rebates.

Figure 3.4 shows the segmentation probability, i.e., the probability that all customers buy where they get the rebate, as a function of r. Note first that even arbitrarily small rebates are sufficient to generate a high segmentation probability. Interestingly, the probability that the market is segmented is not monotonically increasing in r. Rather, the segmentation probability is constant until $r = r^*$, then decreases for some interval until it increases again, reaching the value 1 for $r \geq 1$. To get an intuition for this behavior, note first that the probability of no segmentation is the same as the probability of a successful attack. Now in Cases (i) and (ii) of Proposition 3.3 we can argue as in the proof of Proposition 3.2 that $A(p) = D(p + r)$.

Therefore, given that one firm attacks and the other defends, the probability of a successful attack is 1/2. Observe also that playing an atom in $\overline{d}$ can be interpreted as deciding not to defend but to rely on the cases where the opponent does not attack. We thus get the following: for $r < r^*$, the segmentation probability is independent of r, as it only depends on q which is independent of r. For $r \geq r^*$, the firms set an atom in $\overline{d}$, which implies that the probability of success of an attack increases. This effect drives the segmentation probability down. Yet as r further approaches 1, the fact that attacks become increasingly rare takes over and the segmentation probability approaches 1.

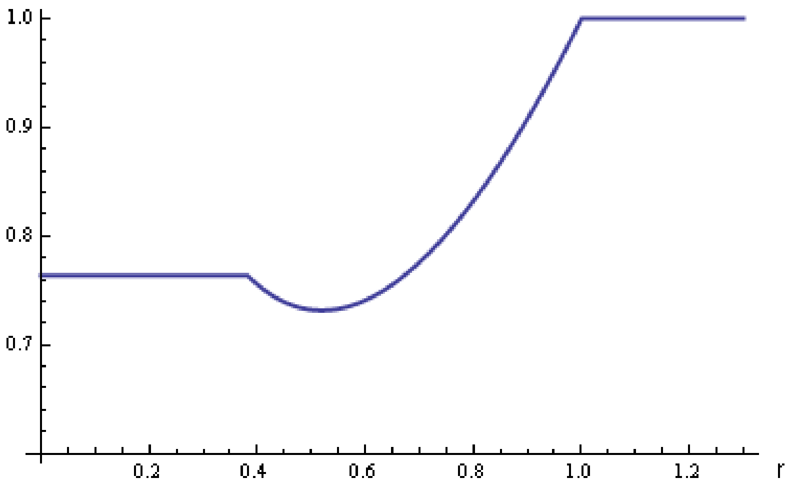

Figure 3.4: The probability of market segmentation as a function of r.

3.4 The Generalized Model

We generalize the model by introducing a mass m_0 of customers who do not get a rebate from any of the firms. We also allow for a more general demand function. A customer's demand depends on the lowest net price which he has to pay at either of the firms and is denoted by $X(\cdot)$. We impose the following assumptions on X: it is positive at least for small positive net prices and continuous and non-increasing in the net price. We also assume that the monopoly profits are bounded.[3] We next

[3]This rules out equilibria à la Baye and Morgan (1999). They show (in a model without rebates) that when the monopoly profits are unbounded "any positive (but finite) payoff vector can be achieved in a symmetric mixed-strategy Nash equilibrium" (p. 59).

distinguish two cases: in the first, all customers are homogeneous in the sense that all have the same rebate opportunities; in the second, customers are heterogeneous, i.e., they have different rebate opportunities.

3.4.1 Homogeneous Customers

Assume customers are homogeneous, i.e., $m_i > 0$ for exactly one $i \in \{0, 1, 2\}$. Then there is perfect competition in net prices and hence the Bertrand paradox arises: two firms are sufficient to yield the competitive outcome.

Proposition 3.4. *Suppose that customers are homogeneous, then both firms earn zero profits.*

Next we show that this is no longer true when customers are heterogeneous.

3.4.2 Heterogeneous Customers

Assume customers are heterogeneous, i.e., $m_i = 0$ for at most one $i \in \{0, 1, 2\}$. This implies that customers differ in the net prices they face. The next lemma states that in equilibrium no firm will charge a negative price. Loosely speaking, the reason is that a negative price leads to losses once something is sold. For a firm which offers a rebate we get a stronger condition.

Lemma 3.1. *In any Nash equilibrium, no firm charges negative prices. A firm which offers a rebate charges prices well above zero.*

We next show that the Bertrand paradox does no longer arise.

Proposition 3.5. *In any Nash equilibrium, both firms earn positive expected profits.*

That is, when customers are heterogeneous, competition is relaxed and firms earn positive expected profits. This also holds when only one firm offers a rebate. Generally, rebates make switching less attractive for customers. This segments the market and allows firms to earn profits. In contrast, without rebates or with rebates which can be used by all customers the market does not get segmented and firms yield zero profits; see Proposition 3.4.

When only one firm offers a rebate, its position in the price competition seems to be weak: when it attracts customers, it has to charge a sufficiently positive gross price to make no loss. In contrast, the competitor also makes no loss when it charges a price of zero. So why should a firm offer a rebate to some customers? The reason is that the competitor knows about the "weakness" of the rebate offering firm and therefore sets a positive price in equilibrium. But given this, the rebate offering firm can target the potential rebate receiving customers and yield a positive expected profit.

So far we have derived characteristics of any Nash equilibrium. Yet we were silent about equilibrium existence. Before we turn to this, we make an assumption which guarantees that playing very high prices is dominated.

Assumption 3.1. *The demand function is elastic above a threshold price. Technically, $X(p)$ is such that there exists a $\hat{p}$ so that $\varepsilon_{x,p} := -\frac{X'(p)}{X(p)/p} > 1 \ \forall p > \hat{p}$.*

Sufficient conditions for Assumption 3.1 to hold are that for some price the demand function is elastic ($\varepsilon_{x,p} > 1$) and that the demand is log concave (this implies, see Hermalin (2009), that $\varepsilon_{x,p}$ is increasing in p).

Lemma 3.2. *Under Assumption 3.1 playing prices above $\hat{p} + r_j$ is dominated for firm j in any Nash equilibrium.*

With the help of Lemma 3.2 we can establish the existence of a Nash equilibrium.

Proposition 3.6. *Under Assumption 3.1, for any tie-breaking rule a Nash equilibrium exists.*

There is an alternative assumption to Assumption 3.1 which yields Lemma 3.2 and also Proposition 3.6. There is a choke price: $X(p) = 0 \ \forall p \geq \tilde{p}$. Then prices above $\tilde{p} + r_i$ are dominated for firm i.

Klemperer (1987a, Section 2) shows for an example that firms yield monopoly profits in their market segment. This result holds more generally.[4]

[4] Existence of a monopoly price is assumed in Proposition 3.7. One can easily show that Assumption 3.1 is sufficient for existence.

Proposition 3.7. *Suppose $m_0 = 0$, $m_1, m_2 > 0$, and there exists a monopoly price p^M. When the rebates r_1 and r_2 are sufficiently large, both firms earn monopoly profits in their market segment in equilibrium. An equilibrium in pure strategies supports this outcome. The same is true when there exists a choke price $\tilde{p}$ and $m_0, m_1, m_2 > 0$.*

Intuitively, when the rebates are high no firm wants to attack the customers in the other firm's home base. The reason is that such an attack would require setting a gross price which is low compared to the rebate the customers in the own home base get. Therefore, attacking would lead to a loss. This gives both firms the freedom to set gross prices such that customers pay net prices equal to the monopoly price. Thus the home base of firm i buys at firm i and both firms yield monopoly profits in their market segment.

When there is a choke price which is low compared to the respective rebates, even the existence of customers who do not get rebates does not affect this result: firms still target only their home bases, because the high rebates make lower prices unattractive. Hence customers without rebate opportunities end up buying no product.

3.5 Welfare and Customers' Coordination Problem

Without rebates or with homogeneous customers total welfare is maximized because net prices equal marginal costs. With rebates and heterogeneous customers at least some customers buy for positive net prices. Hence, given a standard downward sloping demand function, total welfare is no longer maximized.[5] Note that firms are in expectation better off (see Propositions 3.4 vs. 3.5). Taken together, this implies that rebates deteriorate the customer welfare.

Customers face a coordination problem. They would collectively be better off when there are no rebates. This type of coordination is, however, not credible when there are many customers who cannot write contracts on whether or not they participate

[5] For the case in which there is constant demand, total welfare is constant for all prices for which customers buy. Nonetheless, rebates deteriorate the customer welfare.

in rebate systems. First note that when a customer has no mass, then he does not change the firms' pricing policies by participating or not participating in a rebate system. When he participates, he has the option to use the rebate and is therefore weakly better off than we he does not participate. There are cases where he is strictly better off. Therefore, the customer is in expectation strictly better off when he participates in a rebate program.

3.6 Endogenous Rebates

Up to now we have concentrated on the price setting of the firms when the rebates are given. This approach may be a good description of the short run behavior of firms where the rebate system is established and cannot be overturned. Additionally, in some lines of industries, like the airline industry, several firms have a common rebate system. Then a firm can hardly change rebates when it decides about its prices.

Next, we endogenize the rebates. We hold the analysis brief and non-technical. This should allow the reader to gain some intuition which does not rely on some specific – and to some extent arbitrary – modelling of the rebate setting stage.

Suppose that firms first set rebates simultaneously before they compete in prices. From Proposition 3.4 the following result is immediate.

Proposition 3.8. *That both firms set no rebate is not a subgame perfect Nash equilibrium. It is also not subgame perfect that both firms offer rebates to all customers.*

When both firms set no rebate then both firms will earn zero profits. This cannot be optimal because by offering a rebate to some customers a firm can yield a positive expected profit; see Proposition 3.5. The same arguments apply when firms offer rebates to all customers.

The equilibrium of the rebate setting stage may be in pure or mixed strategies. When it is in pure strategies, then firms offer rebates to a subset of customers which implies that both firms will earn positive expected profits. When the equilibrium is in mixed strategies, firms will still earn positive expected profits. The reason is that a firm may use the pure strategy and offer a rebate to a subset of customers which guarantees the firm a positive expected profit.

We close with a simple example. Suppose a mass of 1/3 of the customers participate in the rebate program of each firm, while the remaining 1/3 does not participate in any program. Technically, $m_0 = m_1 = m_2 = 1/3$. Suppose that the customer's choke price is 1. For concreteness assume that each firm can choose one of the following rebates: $\{0, \underline{r}, \bar{r}\}$, where $0 < \underline{r} < 1$ and $\bar{r} > 3$. When both firms choose the high rebate, both firms earn monopoly profits in their market segment in equilibrium; cf. Proposition 3.7. Each firm's profit is then 1/3. To see that this cannot be an equilibrium observe that a firm which chooses a rebate of zero earns 2/3, given that the other firms will choose the high rebate. The reason therefore is that the other firm would yield a loss when it offers a gross price which is lower than 1. Therefore, the firms will not choose rebates which are very high and yield to the monopolization of their market segments. This is due to the customers which do not participate in the rebate program. It is too tempting to sell to those customers. Additionally, it can be no equilibrium that both firms offer zero rebates; cf. Proposition 3.8.

3.7 Concluding Discussion

We showed that in a Bertrand game rebates lead to a segmentation of the market when customers are heterogeneous in the rebates they can get. This segmentation has the effect that both firms earn positive expected profits. We close with a discussion.

Entry.— Rebates lead to positive profits for firms when customers are heterogeneous. Therefore, when entry costs are positive, rebates may lead to entry into a market into which otherwise there would be no entry. In this sense, rebates may increase competition in a market.[6]

Heterogeneous Demand.— Note that the results obtained in Section 3.4 also hold when customer types have different demand functions: all proofs can be modified so that the demand function is type-dependent as long as the demand functions fulfill the assumptions we made.

[6] An argument along these lines is already made, for example, by Beggs and Klemperer (1992). For a model on entry deterrence in case of switching costs, see Klemperer (1987c).

Discrimination.— We assumed that firms cannot price discriminate. Technically, each firm has to offer a single gross price to all customers. Suppose now that firms can perfectly price discriminate. Then firms know what rebates a customer can get and are able to offer customer-specific gross prices. Then each customer can be thought of as an own, separate market. Because there is competition in prices, both firms will in equilibrium earn zero profits on each market. More specifically, in equilibrium both firms offer each customer a gross price so that the net price equals the marginal production costs. Therefore, for the effectiveness of rebates it is crucial that firms cannot discriminate perfectly.

More Than Two Firms.— Suppose there are $N > 2$ firms. When customers are homogeneous or at least two firms set no rebates the Bertrand paradox arises: it is an equilibrium that all firms set prices equal to their rebate and all firms yield zero profits. Otherwise, the logic of Proposition 3.5 applies and all firms earn positive expected profits in equilibrium.

Customers Who Can Get Rebates From Both Firms.— Suppose there is a mass m_3 of customers who can get rebates from both firms. Suppose $m_0, m_1, m_2, m_3 > 0$. This case arises, e.g., when customers randomly receive rebate coupons: some might receive coupons from both firms, some from one firm, and others from no firm. Then both firms must still earn positive expected profits in equilibrium. The line of argument is as before: first, both firms will only charge prices well above zero. Second, this gives both firms the opportunity to earn a positive profit by charging gross prices which are higher than their rebates.

Chapter 4

Second-Price Auctions with Information Release

This chapter studies the optimal release of advertising and information in independent private values second-price auctions.
We develop mild but sharp conditions under which the seller decides to allocate his costly advertising or informational efforts among the potential bidders as concentrated as possible. The seller overinvests in advertising if the valuations of the bidders are drawn from a distribution with increasing failure rate. He underinvests if the distribution has decreasing failure rate. The overall level of advertising is higher under distributions that are more dispersed in terms of the excess wealth order.

4.1 Introduction

"The outcome of an auction never depends on the number of participating bidders." This (translated) citation from Matthias Kurth, head of the federal agency in charge of the German Spectrum Auction 2010, makes one stumble at first sight. What he probably wanted to say is that in some situations it may well make sense to include less bidders into an auction than would be possible.[1]

Three explanations are typically given for this:[2] First, it may be costly for bid-

[1]Compare Financial Times Deutschland, 9 April 2010.

[2]Compare Milgrom (2004), Chapter 6.

ders to participate in an auction. Then bidders may shy away from taking part in the auction, fearing that so many other bidders may participate that entering the auction does not pay out. Second, if bidders enter one after another and entry is again costly, if some bidders already decided to enter, even high type bidders may prefer not to enter the auction. A third reason comes from practical considerations: Confidentiality of information transmitted between seller and bidder may become more insecure the more bidders get involved. Hence only a small number of bidders may get invited to participate. This chapter provides another reason why sellers decide to include only a selection of bidders into a second-price auction, why they try to eliminate all randomness in the number of participating bidders, and why they equip those bidders involved with as much information as possible.

Our argument for concentrated advertising and concentrated information release in second-price auctions works without entry-costs, without ex-ante asymmetric bidders, and it works for welfare maximization as well as for maximizing the seller's revenues. We assume that initially, the bidders are not informed (or not well-informed) about the auction taking place. In the first model we analyze, the bidders need to receive advertising from the seller to become aware of the second-price auction. Advertising efforts are costly for the seller. We find that the seller generally gives out advertising as concentrated as possible, informing essentially each potential bidder completely or not at all. For many situations, like technically demanding procurement auctions or auctions of fine art or jewelry, it is realistic to assume that bidders may not be aware of the auction unless they receive an invitation from the seller. Also, it is typically observed that advertising or invitations to the bidding process are given out to a small, selected group of bidders: In military procurement auctions, the number of invited bidders often does not exceed the number of two, even if there may be other firms meeting high research and quality standards.

Our study derives a mild (but sharp) condition, fulfilled by many distribution functions F from which the private valuations of the bidders may be drawn, under which the seller's gross payoffs, i.e. the second order statistics of valuations, are essentially concave in the number of participating bidders. Under concavity, the seller wants to concentrate advertising in order to eliminate the randomness in the number of active bidders. By this, he maximizes his expected revenue.

The condition we derive for the concavity of payoffs makes use of techniques from reliability theory.[3] We find that if

$$\frac{f(x)}{(1-F(x))^2}$$

is an increasing function of x, the sequence of expected second order statistics is concave in the number of participating bidders. In the words of reliability theory, this condition says that the instantaneous probability that an object A fails, given that two objects A and B have lasted until now, must be increasing. The link to second order statistics is hence not surprising.

As a second related model, we consider a private values second-price auction where the bidders are aware of the auction, but need further information from the seller to infer their valuations. Uninformed bidders just bid their best ex-ante estimate of their valuation. In this setting as well, the seller concentrates costly information on as few bidders as possible under mild conditions. The concentration results also carry over to auctions with reserve prices.

We compare the optimal overall level of advertising under the objective of maximizing the seller's revenues to the social optimum. If the bidders' valuations are drawn from a distribution with increasing failure rate (IFR), the seller overinvests in advertising. The reason behind is that under IFR, the difference between the first and second order statistic decreases in the number of participating bidders. Hence in this case, a larger number of participants is more desirable for a revenue-maximizing seller than for a welfare-maximizer: The selling price is more responsive to advertising than the valuation of the winning bidder. Conversely, the seller underinvests in advertising if the distribution function has a decreasing failure rate (DFR).

Furthermore, we demonstrate that the excess wealth order[4] serves as an appropriate tool to compare the levels of advertising given out under different distribution functions: If F is more dispersed than G in the excess wealth order, the seller gives out more advertising under F than under G. The excess wealth order focuses on the dispersion at the upper subintervals of the distribution supports. An in this sense more dispersed function offers a higher probability that adding another bidder in-

[3]See, e.g., Barlow and Proschan (1981). For previous applications of reliability theory to the auction literature, see Li (2005) and Moldovanu, Sela and Shi (2008) and the references in these.

[4]For an introduction to stochastic orders, see Shaked and Shantikumar (2007).

creases the winning bid. Thus advertising pays out more if the distribution function is more dispersed.

4.1.1 Related Literature

There is a considerable literature on auction models in which it is the bidders who decide to acquire more information on the object for sale before they make their bids. Compare the survey by Bergemann and Välimäki (2006).

Less attention has been spent on cases where it is the seller who controls the flow of information. Yet it is intuitive that it is often the seller who has much better access to information on the object for sale. Probably he cannot foresee how exactly further information will affect the bidders' valuations, but he can anticipate that in expectation releasing information has an impact.

Two related papers are Bergemann and Pesendorfer (2007) and Esö and Szentes (2007). Both assume that the seller releases information at no costs. In Esö and Szentes (2007), the seller can sell information via up-front fees and thus always releases all information. In Bergemann and Pesendorfer, the seller cannot sell the information. Hence, when giving out information, he faces a trade-off between potentially higher bids from better informed bidders and potential losses from giving up the insurance that less informed bidders never bid extremely low (as they stick to a rough estimate of their valuation). Considering a wide class of information structures, Bergemann and Pesendorfer show that the information structure in the revenue-maximizing mechanism treats bidders asymmetrically. Hagedorn (2009) further explores the optimality of asymmetric over symmetric information structures in the same setting.

The assumption that releasing information is costless for the seller fits well to situations where the seller possesses the information anyway and where he can costlessly pass it on to the bidders. If the seller has to acquire the information for the bidders, or if he has to explain to the bidders in detail what they want to know, costs come into play. This motivates us to assume that providing information is costly.

Ganuza and Penalva (2010) and Hoffmann and Inderst (2009) consider models of costly information release as well. Yet in their models, the seller cannot give different amounts of information to different bidders - which is the central feature of our model: Ganuza and Penalva (2010) explicitly rule out asymmetric allocations of information, while Hoffmann and Inderst (2009) focus on the one bidder case.

The idea of modeling the release of information via transmission probabilities goes back to Lewis and Sappington (1994) and Johnson and Myatt (2006). The same approach has also been taken by Hoffmann and Inderst (2009). In line with the literature, we assume that the seller is truthful: For example, one could assume that he does not know the bidders' specific needs well enough to be able to misrepresent information in a favorable way. Esö and Szentes (2007) and Hoffmann and Inderst (2009) assume that the bidders hold some preliminary information about their valuations. In Section 4.3.2 we show how to apply our results to a simple setting where preliminary information is included.

Compared to the rich mechanism-design problem studied in Bergemann and Pesendorfer (2007), our model of information transmission probabilities is rather simple. Obviously, one advantage of our approach is that we obtain an explicit characterization of the optimal information structure in a setting which is otherwise fairly general: In our setting we cannot only show that (as in the setting of Bergemann and Pesendorfer) symmetric information structures are non-optimal - we even obtain optimality of the most asymmetric structures.

Yet this is not the sole advantage of the simplicity of our model: Consider e.g. the seller of a house. It does not seem overly realistic that such a seller is able to let bidder 1 learn whether his valuation for the house is above or below $50,000$ \$ (and nothing else) and bidder 2 whether his valuation is above or below $200,000$ \$. But this is only a small fraction of the power the seller in the setting of Bergemann and Pesendorfer is assumed to possess. Considering such examples it becomes natural to think about the decision problem of a seller whose power does not go beyond an imperfect control of the amount of information transmitted to each bidder.

A natural question at this point is why we do not use a model of information transmission based on stochastic orders as is done e.g. in Ganuza and Penalva (2010). While an extension of our analysis along these lines would certainly be interesting, note that the theory of stochastic orders is not well-suited to the allocative questions we study: Our analysis requires comparisons such as the one between the costs of dividing a fixed amount of information among two bidders instead of concentrating the same amount on only one bidder. For this purpose an ordinal ranking on the space of information structures is not sufficient - we need a meaningful cardinal ranking of the information structures. Hence we believe the model of information transmission probabilities is a natural framework for addressing the questions we ask.

Our results on advertising are also related to the literature that studies auctions where bidders face entry costs. Seminal references are McAfee and McMillan (1987) and Levin and Smith (1994), see also the survey of Bergemann and Välimäki (2006). Typically, in these models it is assumed that bidders learn their valuations only after entry. In Levin and Smith (1994), bidders mix over their entry decision. This implies that bidders are indifferent between entering and not entering and that thus the seller's revenue coincides with social welfare. In other models such as McAfee and McMillan (1987) and Lu (2008), the seller can extract all surplus via individual entry fees. This also implies that - unlike in our models - welfare- and revenue-maximizing incentives coincide. The most closely related result in this literature is found in Lu (2008) who shows that welfare is maximal if the variance in bidders' entry is minimal. The same argument is made in Theorem 6.6 of Milgrom (2004) whose interpretation is closer to ours.

Our analysis of advertising extends the latter result in several directions. Most importantly, we prove a similar result for the case where the seller can only extract the second highest of the bidders' valuations. Technically, this is considerably more demanding, and as discussed above, some mild additional assumptions have to be made so that maximally asymmetric participation (or learning) probabilities are optimal. Additionally, we allow for cost functions that depend on the participation probabilities in a more general way than in Milgrom (2004) and Lu (2008) where linear cost functions are assumed.

The fact that the seller's objective does not coincide with welfare maximization raises the question of over- or underprovision of advertising. The microeconomic literature[5] distinguishes two types of advertising: Roughly, persuasive advertising changes the consumers' tastes about a product, whereas informative advertising

[5]See, e.g., Bagwell (2001) and Tirole (1994).

makes consumers know the product better. Clearly, we are concerned with informative advertising here. In a classical paper, Shapiro (1980) demonstrates that a standard monopolist who cannot price-discriminate and who can sell as many objects as he wants typically underprovides informative advertising. The reason is that he cannot extract the whole surplus from the consumers, hence he does not fully internalize the gains from advertising and then selling to more consumers. In our advertising setting, the product of the seller is scarce. Underprovision of advertising then occurs whenever the expected selling price in the auction reacts too little to advertising. Yet our analysis also shows that for many distribution functions, the seller actually overprovides advertising as the selling price reacts stronger to advertising than the winning bidder's valuation and thus welfare.

Our results also shed additional light on the well-known result of Bulow and Klemperer (1996) that for a revenue-maximizing seller adding one more bidder to an auction without reserve price dominates introducing any reserve price. Hence if it is equally costly to mobilize one more bidder or to learn the distribution of the bidders' valuations well enough to set an optimal reserve, the former should be preferred. At this point our results on advertising an auction come into play, providing mild conditions under which mobilizing one more bidder dominates mobilizing e.g. two bidders with probability 1/2. The result of Bulow and Klemperer suggests that negotiation with one or more bidders does not pay out if the full set of potentially interested bidders is not explored. Our study adds how this set of bidders should be explored optimally.

More generally, by studying auction models with endogenous numbers of bidders, we follow the advice of Klemperer (2004): According to him, with regard to practical auction design, endogenizing the participation of bidders is one of two key issues needing further exploration.[6]

The road-map is as follows: Section 4.2 introduces and analyzes advertising of auctions. Section 4.3 addresses auctions with information release. In Section 4.4, we provide a discussion of our modeling approach, notably the assumptions made on the cost functions. Section 4.5 concludes. The Appendix is structured as follows: Appendix D.1 provides all results on order statistics. Appendix D.2 develops the

[6] The other issue is collusion. Compare Klemperer (2004), Chapter 4.

concentration results in an abstract decision framework that covers both, revenue- and welfare-maximization, for all auction models considered. All results (including those of Appendices D.1 and D.2) are proved in Appendix D.3.

4.2 Advertising

There is one object for sale for which the seller has zero valuation. There are n potential, risk-neutral bidders i with independent, private valuations v_i drawn from a distribution F with finite expected value μ. Denote by $v_{j:k}$ the j^{th}-largest out of k independent F-distributed random variables. F is assumed to be common knowledge and admits a density f which is positive.

Initially all or most of the potential bidders are unaware of the second-price auction taking place. An unaware bidder i independently learns about the auction with probability γ_i. γ_i is set by the seller, and it is thought of as the intensity of advertising the seller concentrates on bidder i. Only bidders who got aware of the auction can submit a bid. We assume that the bidders taking part adhere to the dominant strategy equilibrium of bidding the best estimates they have of their valuations.

Setting a vector $\gamma = (\gamma_1, ..., \gamma_n)$ of advertising intensities costs the seller $c(\gamma)$. c is an increasing, symmetric, continuous, and quasi-concave function. The seller wants to maximize his revenues.

Denote by p_k the seller's expected payoff conditional on k of the bidders becoming aware of the auction. As in a standard second-price auction at least two bidders must be present to allow for a payoff above zero to the seller. Hence it holds $p_0 = p_1 = 0$ and for $k \geq 2$

$$p_k = E[v_{2:k}].$$

To sum up, the seller faces the following maximization problem:

$$\max_{\gamma \in [0,1]^n} \pi(\gamma) - c(\gamma)$$

where $\pi(\gamma)$ is the expected net surplus given a vector of advertising intensities γ

$$\pi(\gamma) = \sum_{j=0}^{n} \alpha_k(\gamma) p_k,$$

and $\alpha_k(\gamma)$ is the probability that exactly k of the bidders participate in the auction:

$$\alpha_k(\gamma) = \sum_{A \subseteq \{1,\dots,n\}, |A|=k} \prod_{j \in A} \gamma_j \prod_{l \in A^c} (1 - \gamma_l).$$

4.2.1 Optimality of Concentrated Advertising

We show in the following that under mild assumptions the seller chooses the γ_i as asymmetrically as possible, setting all but one γ_i either to zero or to one. Let us call such allocations *maximally asymmetric* allocations of advertising effort. Appendix D.2 proves in detail the intuitive result that maximally asymmetric allocations are favored provided that the payoff sequence p_k is essentially concave. The underlying principle that drives this result is Jensen's inequality. What requires more elaborated techniques is to ensure the concavity of the payoff sequence, hence the concavity of sequences of expected second order statistics and related quantities. Appendix D.1 is devoted to this.

Generally, sequences of second order statistics do not have to be concave. There is one crucial, mild but sharp condition that ensures their concavity, and hence leads to concentration of information by the seller. We call this the IFTR condition, where IFTR stands for Increasing Failure-out-of-Two Rate as is discussed below.

Definition 4.1. *The distribution F on $[0, \infty)$ with density f has an Increasing Failure-out-of-Two Rate (IFTR) iff*

$$\frac{f(x)}{(1 - F(x))^2}$$

is strictly increasing in $x > 0$.

In addition, we need the following slightly stronger assumption for part of our following results:

Definition 4.2. *The distribution F on $[0, \infty)$ with density f has an Unbounded Increasing Failure-out-of-Two Rate (UIFTR) iff F is IFTR and*

$$\lim_{x \to \infty} \frac{f(x)}{(1 - F(x))^2} = \infty.$$

Before we provide a discussion of these properties, let us state the proposition, which shows that the seller typically concentrates a given total amount of advertising efforts on as few bidders as possible:

Proposition 4.1. *(i) Assume F is IFTR. If γ^* is revenue-maximizing among all allocations γ for which $c(\gamma) = \kappa$ and if at least one bidder i participates in the auction for sure, γ^* must be maximally asymmetric.*

(ii) Assume F is UIFTR. Then there is an n^* such that for $n > n^*$ there is a κ^* with $c(0, \ldots, 0) < \kappa^* < c(1, \ldots, 1)$ such that, for all $\kappa \geq \kappa^*$, if γ^* is revenue-maximizing among all allocations γ for which $c(\gamma) = \kappa$, then γ^* must be maximally asymmetric.

The first part of the proposition shows that the seller allocates advertising maximally asymmetrically if at least one bidder is known to be present in the auction anyway. The second part states that the seller opts for maximally asymmetric allocations as well if all bidders are initially unaware of the auction and the total advertising he wants to allocate is sufficiently large.

For example, if F is the uniform distribution, the seller chooses to concentrate as soon as the total amount of advertising is enough to inform two bidders completely. Under the exponential distribution, he opts for maximally asymmetric allocations if the total advertising is enough to inform three bidders completely.

In the language of reliability theory, the expression $f/(1-F)^2$ from the IFTR condition can be interpreted as the instantaneous probability that a component A fails, given that components A and B have lasted until now. That there is a relation to second order statistics is hence not surprising. More rigorously, for the concavity of the payoff sequence it is necessary that the increments in the expected second order statistics, $E[v_{2:k+1} - v_{2:k}]$, are decreasing. In the Appendix, we show that

$$E[v_{2:k+1} - v_{2:k}] = E\left[\frac{(1 - F(v_{1:k}))^2}{f(v_{1:k})}\right].$$

Hence, if $f/(1-F)^2$ is increasing, the payoff increments are decreasing in the number of participating bidders.

The IFTR Condition reminds of the well-known increasing failure rate condition (IFR) that $f/(1-F)$ is increasing, but it is considerably weaker: Clearly, any distribution function that is IFR is also IFTR, as $1 - F$ is decreasing for sure. As many distribution functions are IFR, even more are IFTR, and hence lead to a strictly concave sequence of expected second order statistics. (A more detailed discussion of this and examples for distribution functions that are IFTR but not IFR can be found in the Appendix.)

To wrap up, IFTR guarantees the strict concavity of the payoffs $p_1, p_2, \ldots$ starting with one participating bidder. Thus, if at least one bidder is known to take part in the auction anyway, or if the seller wants to give out a sufficiently large total amount of information, he chooses a maximally asymmetric allocation.

The second part of Proposition 4.1 is based on the slightly stronger UIFTR condition, which requires additionally that

$$\frac{f(x)}{(1-F(x))^2}$$

goes to infinity in x. This ensures that $E[v_{2:k+1} - v_{2:k}]$ goes to zero as the number of participating bidders k gets large. This requirement is very mild,[7] yet necessary in order to prove that the non-concavity of the payoff sequence around zero (remember $p_0 = p_1 = 0$) plays a negligible role when the seller allocates larger amounts of information.

4.2.2 Comparisons to Welfare Maximization

This section explores the relation between revenue-maximizing and welfare-maximizing advertising. A first result shows that also under welfare maximization, advertising should be allocated maximally asymmetrically. We then compare the total amounts of advertising under the two objectives. It depends crucially on the distribution of valuations whether advertising is over- or underprovided: If F is IFR, the revenue-maximizing seller advertises more than in the social optimum. If F is DFR, the opposite is the case.

In our setting, social welfare is the expected valuation of the winning bidder minus advertising costs. Observe that the welfare maximizing allocation solves a similar maximization problem as in the subsection before with the sequence $q_0 = 0$ and

$$q_k = E[v_{1:k}]$$

[7] Yet it is an additional requirement which is stronger than IFTR and weaker than IFR, as we show in the Appendix.

in place of (p_k).

As Proposition D.1 in Appendix D.1 shows, such sequences of first order statistics are always strictly concave. Hence concentrating information is welfare-optimal:

Proposition 4.2. *If γ^w is socially optimal among all allocations γ for which $c(\gamma) = \kappa$, γ^w must be maximally asymmetric.*

We now turn to the comparison between the welfare- and the revenue-maximizing levels of advertising. For this purpose, we assume that the costs only depend on the sum of advertising efforts spent on each bidder:

Proposition 4.3. *Assume F is IFTR and assume that the costs are given by*

$$c(\gamma) = C\left(\sum_{i=1}^{n} \gamma_i\right)$$

for an increasing, convex and continuously differentiable function C. Assume that one bidder is known to be aware of the auction anyway. Then, if F is IFR, a revenue-maximizing seller implements a weakly higher advertising level than is socially optimal. Conversely, if F is DFR, a revenue-maximizing seller implements a weakly lower advertising level than is socially optimal.

The exponential distributions are the only distributions which are both IFR and DFR. Thus under exponentially distributed valuations the socially optimal and the revenue-maximizing level of advertising coincide. Recall that the level of advertising is linked to the expected number of bidders in the auction. Accordingly under IFR, the expected number of bidders entering the auction is higher under the revenue maximizing allocation than under the socially optimal allocation. The opposite is true under DFR.

The driving factor behind this result is that under IFR the difference between welfare and seller's payoff $(q_k - p_k)_k$ is decreasing in k while it is increasing under DFR. Thus under IFR choosing a high advertising level is more attractive for a revenue-maximizer than it is socially. Under DFR the opposite is the case. Observe that these results do not hinge on the fact that $E[v_{1:k}] > E[v_{2:k}]$: The choice of allocations would not change if either $(p_k)_k$ or $(q_k)_k$ was shifted by some constant.

4.2.3 Dispersion of Valuations

Next we show that both the socially optimal and the revenue maximizing levels of advertising are higher under a more dispersed distribution of valuations. For this purpose we make use of a result by Li and Shaked (2004) on the comparison of spacings of order statistics between probability distributions ordered in the excess wealth order which is defined as follows:[8]

Definition 4.3. *The distribution function F is more dispersed in the excess wealth order than the distribution function G, $F \geq_{EW} G$, if for all $p \in (0,1)$*

$$\int_{F^{-1}(p)}^{\infty} 1 - F(x)dx \geq \int_{G^{-1}(p)}^{\infty} 1 - G(x)dx.$$

The excess wealth order has a strong connection to the Lorenz curve in public economics. The Lorenz curve shows the distribution of wealth across a population and indicates the concentration of wealth among segments of the society. Likewise, the excess wealth order compares between two distributions how much the high valuations are concentrated.

Let us compare the optimal allocation of advertising when bidders have valuations $v_i \sim F$ to the optimal allocation under valuations $w_i \sim G$. Denote as usual by $F_{2:2}$ the distribution of the minimum of two F-distributed random variables.

Proposition 4.4. *Assume F and G are IFTR, and assume that the cost function is given by*

$$c(\gamma) = C\left(\sum_{i=1}^{n} \gamma_i\right)$$

for an increasing, convex and continuously differentiable function C. Assume that one bidder is known to be aware of the auction.

(i) If $F \geq_{EW} G$ the socially optimal advertising level under F is higher than the socially optimal advertising level under G.

(ii) If $F_{2:2} \geq_{EW} G_{2:2}$ the revenue-maximizing advertising level under F is higher than the revenue-maximizing advertising level under G.

[8] Throughout we slightly deviate from the notational convention in the theory of stochastic orders and write stochastic orders as relations between distribution functions, not between random variables.

Note that these results do not hinge on the number of bidders. The excess wealth order considers all possible restrictions of the underlying distribution function to all values above a threshold. If in this sense a distribution exhibits more dispersion than another one, it pays out more to include more bidders into the auction: The probability to end up with bidders with high valuations is more responsive to advertising if the dispersion is higher.

Both parts of the proposition hold also under the stronger assumption that $F \geq_{\text{disp}} G$ where $\geq_{\text{disp}}$ denotes the more frequently used dispersive order, see the discussion in Section D.1 of the Appendix. In there, we also show that $F_{2:2} \geq_{\text{EW}} G_{2:2}$ implies $F \geq_{\text{EW}} G$.

4.2.4 Reserve Prices

Let us now allow for an arbitrary but fixed reserve price $r > 0$. As usual, if no bid above the reserve is submitted, the object for sale remains with the seller. If only one bidder submits a bid above r then this bidder wins the auction and pays r. Hence obviously bidders who got informed about the auction are only payoff-relevant to the seller if their valuation is larger than r. Thus we obtain the modified payoff sequence $p_0 = 0$,

$$p_1 = E[r\, 1_{v_{\{1:1\}} \geq r}]$$

and for $k \geq 2$

$$p_k = E[r\, 1_{\{v_{1:k} \geq r \,\wedge\, v_{2:k} \leq r\}} + v_{2:k}\, 1_{\{v_{2:k} > r\}}].$$

This sequence is generally not concave. Yet by Proposition D.3 in the Appendix there exists a k^* such that under IFTR the sequence p_k is strictly concave for $k > k^*$. The intuition is straightforward: If many bidders take part in the auction, the reserve price has comparatively little influence on the seller's revenues. Hence for larger k, the payoff sequence (p_k) essentially behaves like in the setting without reserve. Likewise, under UIFTR, the increments of (p_k) again go to zero as k becomes large. This allows us to conclude that advertising is provided in a maximally asymmetric way if k^* bidders are known to be present or if a sufficiently large level of advertising is chosen.

Proposition 4.5. *(i) Assume F is IFTR. If γ^* is revenue-maximizing among all allocations γ for which $c(\gamma) = \kappa$ and if at least k^* bidders participate in the auction with certainty, then γ^* must be maximally asymmetric.*

(ii) Assume F is UIFTR. Then there is an $n^* > k^*$ such that for $n > n^*$ there is a κ^* with $c(0, \ldots, 0) < \kappa^* < c(1, \ldots, 1)$ such that, for all $\kappa \geq \kappa^*$, if γ^* is revenue-maximizing among all allocations γ for which $c(\gamma) = \kappa$, then γ^* must be maximally asymmetric.

With a reserve price, the welfare sequence q_k is given by $q_0 = 0$ and

$$q_k = E[v_{1:k} 1_{\{v_{1:k} \geq r\}}].$$

This is again a strictly concave sequence, and hence welfare-optimal allocations must be maximally asymmetric as well. This result is easily obtained by showing that (q_k) is a sequence of standard first order statistics from a different distribution function constructed out of F.[9]

Finally recall that as shown by Myerson (1981) under mild regularity conditions, the optimal reserve price in a second-price auction is independent of the number of bidders. Since bid-strategies are unaffected by the uncertainty in the number of bidders, this implies that the same choice of reserve price is optimal also in our auction with stochastic participation.

4.3 Information Release

This section analyzes a related model about the release of information. Assume now that all bidders are aware of the auction. The distribution function F is common knowledge. With probability γ_i bidder i learns his valuation. Informed bidders bid their valuation while uninformed bidders bid their best estimate μ. As a motivation, think of a complex product bidders need special information on to understand their personal value as they have to figure out how well the object suits to their specific needs. As before, the seller chooses the probabilities $\gamma = (\gamma_i)$ at costs $c(\gamma)$ where c is increasing, quasi-concave, continuous and symmetric. In this setting, we interpret γ_i as an amount of information transmitted to bidder i. We assume that the seller cannot sell information by demanding a fee from a bidder before delivering the information.[10]

[9] See Proposition D.1 in the Appendix. Note, however, that this trick is not possible for dealing with the payoff-sequence (p_k).

[10] Assuming such up-front fees only simplifies the analysis, see the discussion in Section 4.4.1.

4.3.1 Concentration of Information

In order to avoid technicalities associated with informing the final bidders we assume that there are infinitely many uninformed bidders. Thus there are always two uninformed bidders left who bid μ so that the selling price never falls below μ.[11]

Under information release, the payoff sequence (p_k) is given by $p_0 = p_1 = \mu$ and

$$p_k = E[\max(\mu, v_{2:k})]$$

for $k \geq 2$. Note that this sequence is not covered by the advertising with reserve price analysis. Yet again, for larger numbers of bidders, the payoff sequence behaves like $E[v_{2:k}]$. Hence we obtain similar results as in the case of advertising with a reserve:

Proposition 4.6. *(i) Assume F is IFTR. If γ^* is revenue-maximizing among all allocations γ for which $c(\gamma) = \kappa$ and if at least k^* bidders know their valuations with certainty, γ^* must be maximally asymmetric.*

(ii) Assume F is UIFTR. Then there is a κ^* with $c(0, \ldots, 0) < \kappa^* < c(1, \ldots, 1)$ such that, for all $\kappa > \kappa^*$, if γ^* is revenue-maximizing among all allocations γ for which $c(\gamma) = \kappa$, then γ^* must be maximally asymmetric.

The welfare sequence (q_k) in this model turns out to be globally strictly concave.[12] (q_k) is given by $q_0 = \mu$, and for $1 \leq k$,

$$q_k = E[\max(\mu, v_{1:k})].$$

Thus in the welfare-optimal allocations, information is always maximally concentrated.

It is straightforward to include reserve prices into the model: A reserve price $r \leq \mu$ is without bite while a reserve price $r > \mu$ drives all uninformed bidders out of the auction leading to a situation covered by the advertising model with reserve price from Section 4.2.4.

[11] For most natural distributions and sufficiently large but finite n this effect from informing the final two bidders becomes negligible anyway, since $E[\max(\mu, v_{2:n})]$ is arbitrarily close to $E[v_{2:n}]$.

[12] See Proposition D.1 in the Appendix.

4.3.2 Preliminary Information

We now briefly consider a setting where bidders possess preliminary private information about their respective valuations. We focus on a simple model and on showing maximal asymmetry of the welfare maximizing allocation. By this we want to demonstrate two things: That our results are robust and that similar techniques may be applied to numerous further problems.

Preliminary information is also considered in Esö and Szentes (2007) and in Hoffmann and Inderst (2009). In these papers, however, the focus is on extraction of surplus results which are non-trivial because the preliminary information of the bidders makes it comparatively complicated to construct the optimal mechanism. We abstract from this mechanism design problem and ask how a given amount of information should be allocated in order to maximize the (interim)-valuation of the winning bidder in a second-price auction. Since the model is comparatively complex we resort to a concrete distributional assumption.

Assume there are n bidders. Assume that the valuation of bidder i is given by $v_i + w_i$ where v_i and w_i are independent exponentially distributed random variables. Bidder i knows v_i but learns w_i only with probability γ_i. A bidder who does not learn w_i bids $v_i + E[w_i] = v_i + \mu$. In this model the welfare sequence is given by

$$q_k = E[\max(v_1 + w_1, \ldots, v_k + w_k, v_{k+1} + \mu, \ldots, v_n + \mu)].$$

In Lemma D.2 in the Appendix it is shown that the sequence (q_k) is strictly concave. Hence, to maximize welfare information should be maximally concentrated.

4.4 Discussion

In this section we discuss and relax some of the assumptions we made. We first consider the popular assumption that the seller can sell his information by raising ex-ante fees. Then we relax the assumptions on the cost function.

4.4.1 Selling Information

Consider the model of information release from Section 4.3. A popular assumption in the related literature is that the seller can sell his information by demanding an individual fee from each bidder for the delivery of information. This assumption is found e.g. in Esö and Szentes (2007) and in Hoffmann and Inderst (2009), and also in many papers from the literature on auctions with endogenous, costly entry (such as Levin and Smith (1994)). Such fees have the effect that the seller can extract the whole ex-ante expected surplus from the bidders. (Among the related papers which do not make this assumption are, e.g., Bergemann and Pesendorfer (2007) and Ganuza and Penalva (2010).) Thus with up-front fees the seller's maximization problem coincides with the problem of welfare maximization.

A similar up-front fee could be introduced into our advertising model as well. Yet this would be slightly inconsistent with the interpretation of advertising: Then bidders who are unaware of the auction would be willing to pay the seller to make them aware of the auction. The bidders would even know the auction well enough to pay exactly their expected payoff.

4.4.2 More General Cost Functions

In this section we discuss and relax the assumptions we made on the cost function: continuity, strict monotonicity, symmetry and quasi-concavity.

Continuity and strict monotonicity. As can be seen in the proofs, the assumptions of continuity and strict monotonicity are made for convenience only. These two assumptions ensure that in every iso-cost set there is a maximally asymmetric allocation. Without these assumptions, similar results could be derived at the cost of a more complex notation.

Symmetry. Clearly, asymmetries in costs give a second, independent reason for asymmetries in the γ_i. An extension to asymmetric quasi-concave cost functions is straightforward. It is easy to show that the seller still prefers maximally asymmetric allocations. Additionally he prefers to give information (or advertising) to those bidders who are comparatively cheap to inform.

Quasi-concavity. We have seen that the first three of our assumptions are mostly technical and easy to relax. Yet the assumption of quasi-concavity requires a more thorough discussion. From a technical point of view, the quasi-concave cost functions we considered so far have the following property: Among all allocations which lead to the same expected number of informed bidders, the maximally asymmetric allocations maximize gross revenue and minimize costs. This implies that maximal asymmetry is a necessary condition for revenue optimality. If we depart from quasi-concavity, the ranking of allocations induced by the gross revenues still plays a role - favoring maximally asymmetric allocations - but further effects come into play as well.

Quasi-concave cost functions describe e.g. situations where costs depend only on total informational effort or settings where informing a bidder more deeply becomes cheaper if the bidder already received some information (possibly due to learning effects). Quasi-concavity does not cover situations in which giving more information to a single bidder gets costlier and costlier, or situations where the seller can cheaply pass on information already given to one bidder to further bidders. Clearly, such departures from quasi-concavity induce a tendency to spread information instead of concentrating it. Generally, an analysis of such cases sensitively depends on how large the effects from the concavity of the payoff sequence p_k are. The interplay between the cost function and the underlying payoff sequence is then crucial. The following example shows that the effects of p_k and hence gross revenues π_k typically remain important and lead to some concentration:

Let us first consider the extreme case in which costs only depend on $\max_i \gamma_i$,

$$c(\gamma) = C(\max_i \gamma_i).$$

Such a cost function would e.g. fit to the situation where the seller collects all information and publishes it on the internet: Gathering or writing down information is costly, but there are no further costs from delivering information to different numbers of (potential) bidders. Clearly, in this very special case, there is no concentration of information: The seller chooses an amount $\gamma_{\max} \in [0, 1]$ of information and sets $\gamma_i = \gamma_{\max}$ for all bidders. Yet as soon as we move away from this extreme case and allow for costs of delivering information to different bidders, e.g.

$$c(\gamma) = C\left(\max_i \gamma_i, \sum_i \gamma_i\right),$$

our concentration results come into play: We conclude from Proposition D.5 in Appendix D.2 that for fixed choices of $\gamma_{\max} \in [0,1]$ and $s \in [\gamma_{\max}, n\gamma_{\max}]$ (where as before $s = \sum_i \gamma_i$) the seller allocates information as asymmetrically as possible. Hence he gives $\gamma_{\max}$ to as many bidders as possible and the remaining information to one further bidder: The seller optimally chooses some $\gamma_{\max} \in (0,1)$ and $s \in (\gamma_{\max}, n\gamma_{\max})$ and then sets $\gamma_i = \gamma_{\max}$ for $i = 1, \ldots, m$ where $m = \lfloor s/\gamma_{\max} \rfloor$, $\gamma_{m+1} = s - m\gamma_{\max}$, and $\gamma_i = 0$ for $i > m+1$.

We thus see that the relevance of our results extends to settings beyond those explicitly included in our previous analysis.

4.5 Concluding Remarks

Our analysis shows that under mild conditions, for revenue or welfare maximization it is optimal to concentrate information on as few bidders as possible in a variety of second-price auctions: Few bidders knowing their valuations for sure typically lead to higher outcomes than a larger number of bidders who learn their valuations only with some probability.

There are numerous possibilities for extensions of our model which we want to outline:

- More general cost functions could be studied in more detail (especially costs that are convex in the amount of information given to a bidder). We expect that in that case the effects we studied remain important as our analysis of the seller's gross payoff would not be affected. Yet different effects will interfere as demonstrated in the final example of Section 4.4.2.

- It would be interesting to replace the second-price auction by other formats like the first-price or the all-pay auction. We believe that the effects we have identified still play a major role. The analysis of such auction types would be complicated by the fact that these formats are not interim-efficient, unlike the second-price auction. Intuitively, bidders with a different probability of getting informed would shade bids differently. Hence in some cases an informed bidder could win even though another informed bidder had a higher valuation.

- Furthermore, it could be interesting (but challenging) to generalize the analysis of Section 4.3.2 where bidders hold preliminary information to more general distribution functions and to revenue-maximization.

- Another challenging generalization would be to study optimal information allocations in more general models of costly information transmission. Notably, it would be interesting to find a way to transform the concepts of stochastic orders considered in Ganuza and Penalva (2010) into meaningful cardinal rankings of information structures that suit the allocative questions we study.

- Finally, it would be interesting to analyze the optimal release of information in auction environments with interdependent valuations.

Chapter 5

Optimal Disclosure of Costly Information Packages in Auctions

In many countries, legislation requires takeover and public procurement auctions to equip bidders with equal amounts of information. We discuss such level playing field requirements in an independent, private values, second price auction with entry fees. Informing the bidders is costly for the seller. We find that marginal gross revenues do not generally behave monotonically in total information release. In the two bidder case, essentially, any asymmetric allocation of information dominates the symmetric information allocation. Even the bidder who gets less information is willing to pay a higher entry fee for asymmetric information allocations than for the symmetric one. His entry fee coincides with that of the better informed bidder. Losses from allocating an amount of information non-optimally can be substantial.

5.1 Introduction

Different firms compete in a takeover auction. They are all interested in the target, but not for the same reasons: Bidder A is interested in the client list of the target. Bidder B is more focused on possible synergies to reduce production costs. Firm C is in need of the target's know-how. Most likely, it is not easy to get information on any of these aspects. Should the target be willing to open its books, show the client list, give the bidders access to its production processes? It is reasonable to assume that informing the bidders causes the target firm some costs - at least, revealing information takes costly time. Furthermore it is likely that every bidder asks specific questions and has to be monitored while searching through the target. This leads to the main question: How much information should be given to the different bidders in order to maximize the auction's revenue if revealing information is costly?

Closely connected are the following questions: How much revenue loss is generated by the often established rule that bidders have to be informed symmetrically - as is, for instance, stated in the British Takeover Code?[1] Given that there is only some capacity left for informing the bidders, should it be devoted to some selected bidders - and to which? Or should it be shared equally among them?

To consider questions of this type we need a model that allows for giving out different amounts of information to different bidders. In most parts of the chapter, we will consider the following setting: Each aspect a bidder is interested in about the target is wrapped in an information package. The seller (e.g. the target itself) possesses all these information packages. Yet giving out packages is costly, and so the seller will usually not give out all packages. In case a bidder does not get one of his desired information packages, he sticks to his commonly known prior about this aspect of the target. The seller does not know what the information in the packages means to the bidders. He only takes into account that giving out more packages means that bidders will be better informed in the auction (so that they bid higher with some probability).

We find that with two bidders, allocating packages symmetrically is dominated by most asymmetric, "unfair" allocations. For example, concentrating all packages on one bidder generates a higher payoff than splitting up the same amount of information equally. For the n bidder case, such a general statement cannot be made. Yet we will also give examples where the restriction to symmetric information allocations leads to substantial reductions of the seller's revenue for the n bidder case.

With two bidders, both bidders are willing to pay the same maximum fee for the release of more packages - no matter how asymmetrically these packages will be allocated. Thus a bidder who is still uninformed is willing to pay a higher fee for

[1]Compare Rule 20.2 of the Takeover Code. For another example see the EU Directive 2004/18/EC which regulates public procurement, article 29. For further discussion, see Cramton and Schwartz (1991).

additional information given to the other bidder than for receiving this amount of information himself.

With n bidders, in the special case that informing is a zero-one decision (meaning that each bidder only needs one package to get fully informed), we find that there are decreasing returns to information. Hence intermediate cost levels imply that the seller will treat bidders very differently - some will be informed perfectly, others not at all. If bidders need several packages to get fully informed, the seller's returns to giving out information do not necessarily behave monotonically.

As the seller can charge entry fees, maximizing his expected revenue means also that he maximizes social welfare: All expected surplus on the side of the buyers is extracted by the seller via entry fees.[2] All in all, our results show that there are many situations where an "unfair" information policy is the optimal one.

This study is related to some recent works on information acquisition and mechanism design. In particular, there are some other papers dealing with independent private values auctions where the seller chooses the degree to which he informs the bidders: Bergemann and Pesendorfer (2007) consider the case of no entry fees and no information costs. Esö and Szentes (2007) allow for entry fees, but rule out information costs. Ganuza and Penalva (2010) rule out entry fees, but allow for information costs. Thus our study addresses the fourth and still open case where entry fees can be charged and information costs are present. To describe the related papers in more detail, Esö and Szentes show an extraction-of-surplus result which is more general than ours as it allows for preliminary information on the side of the bidders. As informing the bidders is not costly to their seller, he will give out all the information he has. This is in contrast to our framework, where the seller has to face the trade-off between additional rents due to information revelation and additional costs. Hoffmann and Inderst (2009) generalize the results of Esö and Szentes (2007) to a setting where giving out information is costly. Their analysis is restricted to the case where there is only one bidder. In Ganuza and Penalva, the seller is restricted to inform the bidders symmetrically. Thus, our main question of how to allocate information optimally cannot be addressed in their framework. Ganuza and Penalva concentrate on a second price auction - without entry fees this is generally not a

[2]Hence, if entry fees were ruled out, our analysis would still identify the allocations that maximize total surplus. This may be interesting for public procurements without entry fees.

revenue-maximizing mechanism. Bergemann and Pesendorfer (2007) find that the seller will inform agents asymmetrically. But their seller - due to the lack of entry fees - cannot extract all informational rents and thus solves a different optimization problem.

Another related paper is Crémer, Spiegel and Zheng (2007) where the seller faces costs for sequentially inviting bidders to an interdependent value auction. There are no entry fees and no bids from uninformed bidders. Hence, the optimization problem considered differs from ours as well.

There is also a literature on auction environments with costly information acquisition by the bidders. This includes Persico (2000), Bergemann and Välimäki (2002), Compte and Jehiel (2005), Bergemann, Shi and Välimäki (2008), Shi (2007) and Crémer, Spiegel and Zheng (2009). Crémer, Spiegel and Zheng (2009) is closest to our study: They also consider a revenue-maximizing seller who can charge entry fees in an independent private value model and prove a full-extraction of expected surplus result.

Vagstad (2007) is in the middle between the two groups of papers. In there, the seller can costlessly inform either all or no bidders before entry. Then bidders decide whether they (costly) enter the auction. If the seller does not inform the bidders initially, bidders learn their valuations after entering.

All papers which assume that better information can be acquired at a higher cost face the problem of how to model better, more costly information. The simplest approach, taken by Compte and Jehiel (2005), by Vagstad (2007), by Crémer, Spiegel and Zheng (2007, 2009), by Bergemann, Shi and Välimäki (2008) - and also here in Section 5.3 - is to assume that each bidder either stays completely uninformed or learns his valuation perfectly. The other papers utilize more sophisticated models to allow for partial releases of information: Persico (2000) and Bergemann and Välimäki (2002) order the cost of informative signals according to the signal's effectiveness (also known as accuracy), a concept which goes back to Lehmann (1988). Ganuza and Penalva (2010) and Shi (2007) introduce some interesting new classes of orders for ranking the informativeness of signals based on different stochastic orders.

The information package model of our Section 5.4 takes an intermediate approach, allowing for a partial release of information but staying comparably simple and thus tractable: Each bidder's valuation is assumed to be the sum of m iid random variables (m packages of information). The seller decides how many packages he wants to reveal to each bidder at costs which depend on the total number of packages released. As discussed at the end of Section 5.4, this approach allows us not only to have an ordinal ranking (as in stochastic order approaches), but also to have a cardinal measure of how much information the bidders get.

In the literature on takeover bidding, there are several papers which study information acquisition, mostly in the context of two bidder auctions with sequential bidding (see for instance Giammarino and Heinkel (1986), Fishman (1988), and Hirshleifer and Png (1989)). In these papers, however, it is not the seller but the bidders who decide about information acquisition and pay for it. Furthermore, these papers concentrate on the situation where information acquisition is a zero-one decision between full information and no information. Thus they do not address the types of asymmetries we discuss here.

Somewhat complementary to our study, Dasgupta and Tsui (2003) find that for takeovers with two bidders a selling procedure which treats bidders asymmetrically may dominate symmetric selling mechanisms. The model of Dasgupta and Tsui differs considerably from the one in this study as they consider a common value environment without information acquisition. Their asymmetry result holds only if bidders are sufficiently asymmetric ex-ante. In contrast our result that an asymmetric treatment of bidders may be optimal holds for ex-ante symmetric bidders.

The chapter is organized as follows: Section 5.2 introduces the model. We also analyze how much entry fees the bidders are maximally willing to pay depending on how much information is given out. In Section 5.3, we discuss how many bidders should be informed under the restriction that informing a bidder is a zero-one decision. Section 5.4 - the main part which is also technically the most interesting section of this chapter - focuses on the two bidder case where information is spread over several packages. Section 5.5 briefly discusses the case of more than two bidders. Section 5.6 concludes. All proofs, including the calculations behind the examples, are in the Appendix.

5.2 The Model

A seller sells one indivisible object for which his valuation is zero via a second price auction. There are n risk-neutral bidders with independent (but not necessarily identically distributed) valuations $X_1, \ldots, X_n$ with expected values $\mu_1, \ldots, \mu_n$.[3] The bidders do not know their valuations initially. The seller offers against entry fees to give to each bidder i a certain amount of information (represented by a σ-algebra $\mathcal{F}_i$) such that each bidder can calculate an estimate $\widetilde{X}_i = E[X_i|\mathcal{F}_i]$ of his valuation. Each bidder only receives information on his own type, but not on the other bidders' valuations. Let $\widetilde{X}^{(1)}$ and $\widetilde{X}^{(2)}$ be the two highest order statistics among these estimates.

The precise timing of the model is as follows: First, the seller announces individual entry fees to each bidder. He commits to giving out an information structure (describing which bidder will get how much information) and commits also to excluding all bidders that refuse to pay. Second, the bidders decide if they want to pay their fee. Third, the bidders who have paid get their information. Fourth, all bidders who have paid participate in the second price auction. Without loss of generality, ties in the auction are broken with equal probability.

Throughout the chapter, we will assume that the bidders stick to their weakly dominant strategy of bidding the best estimate they have of their valuations. We assume that giving information to the bidders is costly to the seller. We will, however, not specify this assumption before the next sections when we further restrict the $\mathcal{F}_i$. In Proposition 5.1 the entry fees are calculated that the seller can maximally charge such that bidders still participate.

Proposition 5.1. *If the seller offers to release the information sets $(\mathcal{F}_1, \ldots, \mathcal{F}_n)$, each bidder i is willing to pay an entry fee of*

$$e_i = E[(X_i - \widetilde{X}^{(2)})1_{\{i\ wins\}}] = E[(\widetilde{X}^{(1)} - \widetilde{X}^{(2)})1_{\{i\ wins\}}].$$

This leads to an gross expected revenue for the seller of $E[\widetilde{X}^{(1)}]$.

[3]For convenience, it is assumed throughout this chapter that all random variables are not almost surely constant. Without this assumption, all arguments will still go through but some strict inequalities will hold only weakly. We also assume that all random variables are L^1 integrable, i.e. $E[|\cdot|] < \infty$.

Thus the seller can extract all surplus from releasing more information, and this is an optimal selling mechanism in this framework.

With only two bidders, both bidders' entry fees coincide if they have the same prior estimates of their valuations $\mu_1 = \mu_2$:

Proposition 5.2. *Consider the setting of the previous proposition but with $n = 2$. Assume $\mu_1 \geq \mu_2$. Then bidder 1 pays*

$$e_1 = \frac{1}{2}E[|\widetilde{X}_1 - \widetilde{X}_2|] + \frac{1}{2}(\mu_1 - \mu_2)$$

and bidder 2 pays

$$e_2 = \frac{1}{2}E[|\widetilde{X}_1 - \widetilde{X}_2|] - \frac{1}{2}(\mu_1 - \mu_2).$$

Thus we see that the difference in entry fees between the two bidders always equals the difference in priors $\mu_1 - \mu_2$. This does not change, even if one bidder gets much more information than his competitor. The reason is that more information can raise a bidder's valuation - but it can lower it as well, thus reducing competition in favor of the other bidder. With two bidders, good news to the first bidder translate one to one into bad news to the second bidder. With more than two bidders, such a simple relation does not hold anymore.

5.3 Indivisible Information

In this section we consider the case where informing a bidder is a zero-one decision: At a cost of c the seller can inform a bidder completely (and thus make the bidder learn his valuation perfectly). If the seller decides not to do so, the bidder only knows his prior μ_i. We first consider the case where all n bidders' valuations are drawn independently from the same distribution function. We consider this special setting because it is rather tractable, also with more than two bidders. We find that there are decreasing returns to information: The higher the level of c, the lower the number of agents who get informed by the seller. In the following let $X^{l:k}$ denote the lth largest of the random variables X_1,...,X_k.

Proposition 5.3. *Assume all bidders' valuations X_i are iid, with distribution function F and mean μ. If the seller decides to inform $k < n$ bidders, his expected gross revenue is*

$$E[\max(X^{1:k}, \mu)].$$

If the seller decides to inform n bidders, his expected gross revenue is

$$E[X^{1:n}].$$

Additionally,

1) a bidder who will not be informed is only willing to pay a positive entry fee if he will be the only uninformed bidder in the auction.

2) A bidder who will not get informed pays a weakly lower entry fee than a bidder who will get informed.

3) The increase in the seller's expected revenue from informing one more bidder gets strictly smaller with every informed bidder.

Note that - because of $E[X^{1:n}] < E[\max(X^{1:n}, \mu)]$ - informing the last bidder comes along with a slightly smaller increase in total surplus than if there were additional uninformed bidders left. Informing the last bidder means that there is no bidder left who will bid μ for sure. Thus there is some loss in expected revenue compared to a situation with one or more additional, uninformed bidders.

Let us have a look at a simple example with two bidders, illustrating part 3 of Proposition 5.3:

Example 5.1. *Consider the case of two bidders whose valuations X_1 and X_2 are distributed uniformly on $[0, 1]$. Thus an uninformed bidder bids $E[X_1] = E[X_2] = \frac{1}{2}$.*

If the seller gives out no package, his expected net revenue is $\frac{1}{2}$. If the seller gives out one package, his expected net revenue is $E[\max(X_1, \frac{1}{2})] - c = \frac{5}{8} - c$ (which is is the expectation of the highest estimated valuation minus the costs). If the seller gives out two packages, his expected net revenue is $E[\max(X_1, X_2)] - 2c = \frac{2}{3} - 2c$.

Thus the seller will inform no bidder if $c > \frac{1}{8}$, one bidder if $\frac{1}{24} < c < \frac{1}{8}$ and both bidders if $c < \frac{1}{24}$.

We see that there is a wide range of cost levels where the seller decides to inform only one bidder. From Proposition 5.2, we know that both bidders pay the same entry fee, no matter whether they get different amounts of information. With more than two bidders, this does not hold: If bidder 1 gets a package, but bidders 2 and 3 do not, bidders 2 and 3 will not be willing to pay any entry fee. This is because, even if bidder 1 learns that his valuation is very low, bidders 2 and 3 will still bid their priors μ. Thus they do not expect to make any revenues in the auction and are not willing to pay any entry fees.

5.4 Divisible Information with Two Bidders

In Section 5.3, the seller was forced to treat the n bidders asymmetrically in case he wanted to release less than n packages. In this section we will consider a more general model where the seller may inform bidders asymmetrically although he could spread the same number of information packages evenly as well. We will focus on the two bidder case and illustrate what changes with more bidders.

We now assume two bidders that have valuations $X_1 + \ldots + X_m$ and $Y_1 + \ldots + Y_m$ where $X_1, \ldots, X_m$, $Y_1, \ldots, Y_m$, the packages of information, are iid random variables with distribution function F and mean μ. Each package represents an independent privately-valued aspect of the object for sale. Again, the revenue-maximizing seller decides how many packages each bidder should get. The seller has a cost of c per revealed package.

From Propositions 5.1 and 5.2 we know that if for some $k, j \leq m$ the seller reveals $X_1, \ldots, X_k$ and $Y_1, \ldots, Y_j$, his expected gross revenue is

$$E[\max(X_1 + \ldots + X_k + (m-k)\mu, Y_1 + \ldots + Y_j + (m-j)\mu)] \tag{5.1}$$

and each bidder pays an entry fee of

$$e_1 = e_2 = \frac{1}{2}E[|X_1 + \ldots + X_k - Y_1 - \ldots - Y_j + (j-k)\mu|]. \tag{5.2}$$

Then Proposition 5.4 shows that the seller prefers to give a fixed number of packages to one bidder, leaving the other bidder uninformed. Splitting up the packages evenly among the two bidders would give him lower revenues:

Proposition 5.4. *1) Assume $2k \leq m$ packages are to be allocated by the seller. If the distribution of the packages is asymmetric around the mean, it is strictly more* *profitable to concentrate the $2k$ packages of information at one bidder than to split them equally between the two bidders.*

2) Assume $j \leq 2m$ packages are to be allocated by the seller. If the distribution of packages is symmetric around the mean, the seller's revenue does not depend on how the packages are allocated.

Note that the proposition implies that a bidder may pay a higher fee for an additional package given to the other bidder than for receiving this information package himself.

The proof of Proposition 5.4 relies on the fact that for iid random variables X and Y

$$E[|X+Y|] \geq E[|X-Y|] \tag{5.3}$$

with equality if and only if the distribution of X and Y is symmetric around the mean. This result is found for instance in Jagers, Kallenberg and Kroese (1995). That such a formula is useful for the proof is already plausible from (5.2), the formula for the entry fees: Shifting packages from bidder 1 to bidder 2 means turning plus signs into minus signs in the formula. Note that in the two bidder case, maximizing the difference between the higher and lower order statistic (which is twice the entry fee) is equivalent to maximizing the higher order statistic.

Exploiting (5.3) a little more, we can generalize part 1 of Proposition 5.4 as follows:

Proposition 5.5. *Assume $2k < 2m$ packages are to be allocated and distributions are asymmetric. Then any split-up of the type $(2l, 2h)$ where $l > h$ and $l + h = k$ is better than (k, k).*

To illustrate the proposition, consider the following example: Assume the seller wants to give out six packages in total. Then he should not give three packages to each bidder, but rather give four or six packages to one bidder. The proposition does, however, not make a statement about giving five packages to one bidder and one package to the other bidder. Despite this obstacle (which does not seem to be easy to remove except in the case where each package can be rewritten as a sum of two iid random variables[4]) Proposition 5.5 says that splitting up information evenly is, essentially, the least profitable decision the seller can take.

[4] This is possible for instance for infinitely divisible probability distributions like the exponential distribution.

Two things drive the result: First, concentrating packages at one bidder may lead to a very high interim valuation of that bidder (e.g. if all packages he receives contain good news). Such a high interim valuation can never occur if the seller instead splits the total amount of packages evenly among the bidders. Second, giving all packages to one bidder makes sure that the other bidder, who receives no package, has an interim valuation of $m\mu$. Thus the first order statistic of the interim valuations, which our seller wants to maximize, cannot get smaller than $m\mu$. This insurance effect also has some power if a bidder receives not none, but few packages, and gets less important the more bidders take part in the auction (as it gets unlikely that the first order statistic becomes small anyway). Yet with few bidders, the insurance aspect of leaving bidders completely or nearly uninformed plays a crucial role.

We have seen that splitting up information equally is definitely not optimal. But which information policy is best? Is concentrating all information at one bidder best? The answer can depend sensitively on the distribution F as we can see in the following example:

Example 5.2. *Consider the question of how to allocate six packages of information optimally among two bidders whose valuations consist of six packages each. Assume that the packages take only two values,* 0 *and* $\frac{1}{6}+b$*, and have mean* $\mu = \frac{1}{6}$*. Via* b*, we vary the asymmetry of the probability distribution of the packages. An uninformed bidder's expected valuation is* $6 \cdot \frac{1}{6} = 1$*. Denote by* π_{ij} *the seller's expected gross revenue from giving* i *packages to one bidder and* j *packages to the other bidder. In Figure 5.1, we see* π_{60}*,* π_{51}*,* π_{42} *and* π_{33} *drawn as functions of* b*. Recalling Propositions 5.4 and 5.5 it is not surprising that* $\pi_{33}(b)$ *is strictly dominated by the other curves (except for the symmetric case* $b = \frac{1}{6}$ *where all four curves coincide). Beyond these facts, however, the seller's optimization problem is rather complex: The set of values of* b *for which concentrating all information at one bidder is optimal consists of five disjoint intervals. Every asymmetric allocation is strictly optimal for some values of* b*. Hence it depends sensitively on* b *whether allocations* $(4,2)$*,* $(5,1)$*, or* $(6,0)$ *are best.*

Figure 5.2 depicts $\max(\pi_{60}, \pi_{51}, \pi_{42})/\pi_{33}$*, the relation between revenues from allocating six packages optimally and from allocating them equally. As* b *increases, the loss soon reaches a substantial amount. (As* b *gets large allocating information optimally generates over 30 percent more revenue than allocating information equally.)*

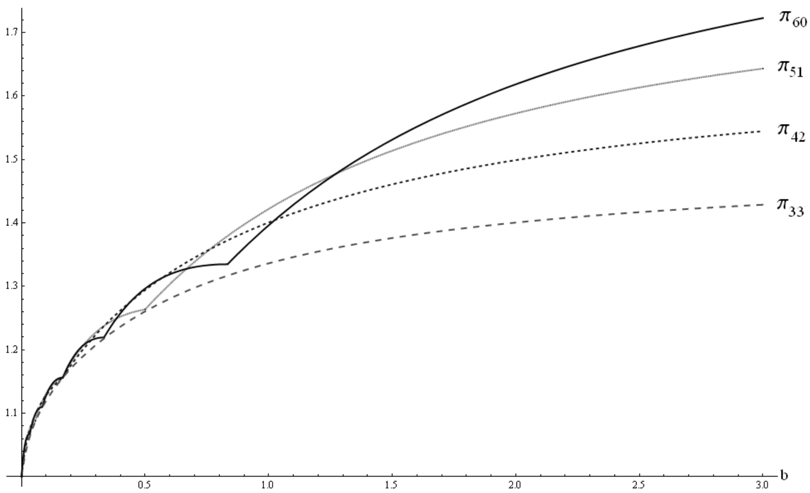

Figure 5.1: Payoffs from Different Allocations in Example 5.2

Thus we see that the optimal allocation depends quite sensitively on the probability distribution of the packages.

So far we have only discussed how to allocate a given number of packages optimally. Now we look at how many packages in total the seller should release. The following lemma shows that for some cost levels the seller will decide to release some, but not all information. (This justifies the approach we have taken in this section so far: The lemma assures that it is worthwhile to think about how an intermediate, fixed amount of packages should be split up among the bidders. Considerations of this type were pointless if the seller always decided to give out all the packages he has.) We see that the first package that is given out is the most profitable one:

Lemma 5.1. *The first package of information given out leads to a strictly higher increase in the seller's revenue than any additional package.*

The question arises of how much additional expected revenue can be made by giving out a second, a third, a fourth package. Is there a result like Proposition 5.3, such that the second package leads to a higher revenue increase than the third, the third package to a higher increase than the fourth, and so on? The following examples show that this depends on the probability distribution F:

Example 5.3. *Assume that each bidder's valuation consists of two packages, i.e.* $m = 2$.

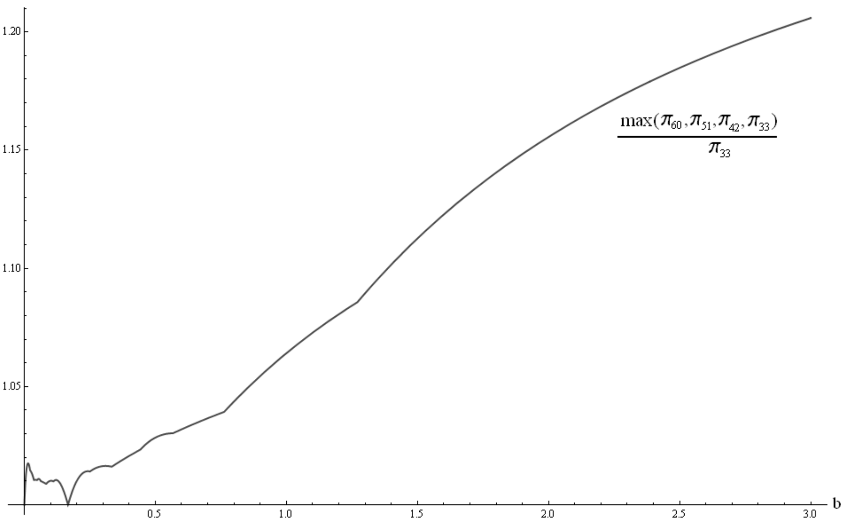

Figure 5.2: Maximal Loss in Example 5.2

1) Assume packages are distributed uniformly on $[0, 1]$*. Then the release of each further package leads to a smaller increase in the seller's expected revenue than the package released before. So we find concavity of expected gross revenue in the number of released packages.*

2) Assume each package takes the values 0 *and* 1 *with equal probability. Then the release of the second or the fourth package does not influence the seller's expected gross revenue. Thus, depending on the level of costs, the seller will give out no, one, or three packages of information. Notably, the seller will only inform the second bidder perfectly if* $c = 0$.

3) Assume packages are distributed exponentially with parameter 1. Then the first package leads to a higher increase in revenue than the second which again leads to a higher increase than the fourth package. The third package, however, leads to a smaller increase in revenue than the fourth. Thus, depending on the costs, the seller will release no, one, two, or four information packages: If the second bidder gets informed at all, he gets fully informed.

The latter two examples have shown that the sequence of the seller's expected gross revenues if he releases a total of l packages is generally not concave in l.[5] But

[5] Such nonconcavities, in a different context, are the main focus of Radner and Stiglitz (1984), see also Chade and Schlee (2002).

the following lemma shows that the sequence cannot be too far from being concave either, namely, that it is increasing and bounded from above by a concave function:

Lemma 5.2. *Let the bidders' valuations consist of m independent, identically distributed packages, each package with mean μ and standard deviation σ.*[6] *The seller's expected gross revenue if he releases a total of $l \leq 2m$ packages is weakly increasing in l and bounded from above by $m\mu + \frac{\sigma}{2}\sqrt{l}$.*

Setting the upper bound on expected payoffs from the lemma in relation to the payoff $m\mu$ from giving out no information at all also gives us a rough upper bound on the maximal losses from allocating information suboptimally.

Finally in this section we want to discuss two assumptions we made - firstly, that packages are identically distributed, and secondly, that there is no preliminary information on the side of the bidders.

To discuss the first assumption let us consider the situation where packages are not identically distributed. The proofs of our results about how to allocate information do not go through in that setting as we cannot rely on inequality (5.3). It is also intuitive that the results themselves do not carry over in general: Assume X_1 and Y_1 had a much larger variance than the remaining packages. Then the seller should release X_1 and Y_1 first, even though this would be an equal split-up of information, in order to create as much variability in the interim valuations as possible.

The following proposition underlines in the case of only one package how the seller's allocation problem depends on the variability of the packages:

Proposition 5.6. *Consider the case $m = 1$, i.e., bidders have valuations X_1 and Y_1. Assume that X_1 and Y_1 are independent and have means μ_X and μ_Y. The condition for informing bidder 1 being strictly more profitable in expectation than informing bidder 2 is*

$$E[|X_1 - \mu_X|] > E[|Y_1 - \mu_Y|].$$

In the special case where $\mu_X = \mu_Y = \mu$ this condition becomes

$$E[|X_1 - \mu|] > E[|Y_1 - \mu|],$$

[6] As we need a finite standard deviation for this lemma, we have to assume here (and only here) that the random variables are L^2 and not just L^1.

i.e., the bidder with the higher mean absolute deviation should be informed.[7] This condition is also equivalent to

$$E[(X_1 - \mu)^+] > E[(Y_1 - \mu)^+]. \tag{5.4}$$

As the uninformed bidder's valuation is certain and guarantees μ, the seller compares in (5.4) a call-option on the valuation of bidder 1 to a call-option on the valuation of bidder 2 (both with strike μ).

Our model of identically distributed, equally costly packages may look rather limited at first sight. This may be especially true if one identifies our information packages with concrete properties of the object for sale. Our packages should be taken as an abstract division of a large amount of information into small pieces which are only loosely related to concrete aspects of the object.[8] The package units allow us to have an exact measure of how much information a bidder gets, and to compare the information amounts different bidders get not only in an ordinal, but also in a cardinal ranking. This is an advantage compared to approached based on stochastic orderings as in Persico (2000) and Ganuza and Penalva (2010): Stochastic orders can express that one bidder is better informed than another. Yet they cannot say what it means that, e.g., bidder 1 gets twice as much information as bidder 2. Our package model is a natural and relatively tractable way to achieve this goal.

Generally, the analysis of a model with non-identically distributed packages would be complicated by the same factor that made obtaining "clean" solutions difficult in the identically distributed model: Expected values of absolute values of sums and differences of random variables are much more difficult to handle than, for example, variances of sums and differences of random variables. Note however that the bounds of Lemma 5.2 immediately translate to any non-identically distributed, independent packages (with the sum of package variances instead of $m\sigma^2$). Furthermore, via continuity arguments it should be possible to extend our analysis to the case of packages which are almost identically distributed.

[7] Note that a random variable X having a larger absolute deviation than Y is not equivalent to X having a larger variance than Y. In many natural examples the two properties however go together.

[8] One could, e.g., assume that $X_1 + \ldots + X_l$ stands for one aspect of the object, and $X_{l+1} + \ldots + X_m$ for another one. Alternatively, one could interpret the packages as hours the seller spends on informing the different bidder.

To illustrate this point consider the case of two bidders with valuations with X_1+X_2 and Y_1+Y_2. The X_i and Y_i are independent, X_1 and Y_1 are exponentially distributed with parameter 1 and X_2 and Y_2 are distributed exponentially with parameter λ. Assume the seller wants to release two packages in total. An elementary calculation shows that for $0.97 < \lambda < 1.04$ it is optimal to inform one bidder fully. Hence we see that the optimality of concentrating information is robust around the value $\lambda = 1$ (at which all packages are iid).

A second assumption of our analysis which may seem rather strong is that bidders do not have preliminary private information. While the effects present in our analysis should still play a role in such a setting, the seller's allocation problem would typically be dominated by other concerns: He should sell information to bidders which value it highly while still keeping the auction competitive enough. The bidders' willingness to pay for information depends crucially on their private information. This leads to a quite complex mechanism design problem of which a solution is beyond the scope of our study. So far, two important special cases of the problem have been considered in the literature: Esö and Szentes (2007) study the problem in the case that releasing information is costless (so that the seller always gives out all information). Hoffmann and Inderst (2009) study the one-bidder-case with costs of information. In both cases, the question of how to split up intermediate amounts of information among the bidders does not arise at all. Still, even for these "simple" cases the optimal mechanisms are intricate. Thus finding the revenue maximizing mechanism in a model that allows for preliminary information and for unequal split-ups of information would be very interesting but also highly challenging.

Nevertheless, our analysis can solve the following non-trivial problem with preliminary information: An efficiency maximizing auctioneer decides about giving out costly information before the auction takes place. He cannot charge fees for information provision (as in Bergemann and Pesendorfer (2007) and in Ganuza and Penalva (2010)). There are two bidders with valuations $X_1 + X_2$ and $Y_1 + Y_2$ where the X_i and Y_i are independent and identically distributed with an asymmetric distribution. Initially, bidder 1 privately knows X_1 while bidder 2 is uninformed. Assume the seller wants to reveal one package in total. Then from Proposition 5.4 we can deduce the following: Revealing X_2, i.e., informing bidder 1 fully, strictly dominates revealing Y_1. Hence we can immediately see with the techniques we developed so far that the auctioneer will want to unlevel the playing field further in a situation where he could level it as well.

Another case which is covered by our analysis is, of course, the one where bidders' valuations have some common value component which is commonly known and where only the information which makes the difference between the bidders is in the hand of the seller.

5.5 Divisible Information with More than Two Bidders

So far we have focused on the two bidder case and found that a fixed amount of information will generally be allocated unequally among the bidders. To get a more complete picture we now have a brief look at examples with more than two bidders. The first example shows that we cannot hope for an equally general asymmetry result as in the two bidder case. It is a three bidder version of Example 5.2.

Example 5.4. *Consider the problem of how to allocate six packages of information among three bidders. Assume that the probability distribution of the packages is the same as in Example 5.2. Let $\pi_{ijk}(b)$ denote the expected gross revenue from allocating i, j and k packages to the three bidders. Figure 5.3 compares the allocations π_{ijk}.*

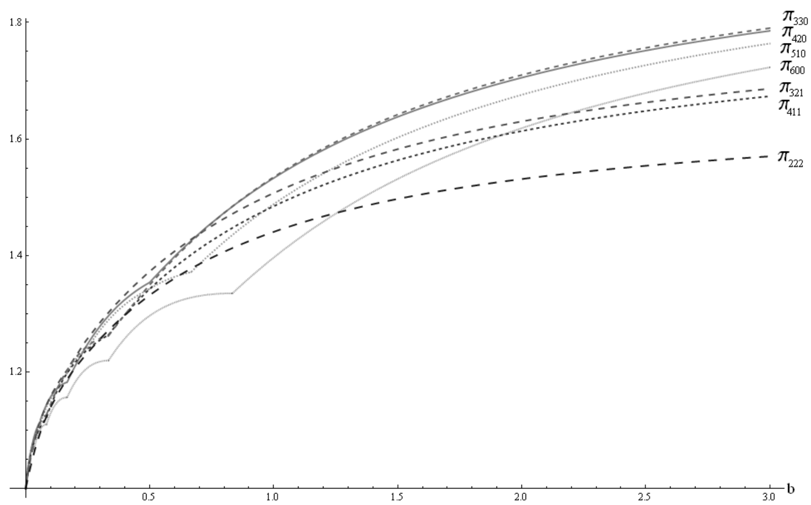

Figure 5.3: Payoffs from Different Allocations in Example 5.4

We see that $\pi_{222}(b)$ is no longer the universal worst choice for all b, but it is still far from optimal. The optimal information policy still depends sensitively on b.

In Figure 5.3 the curves no longer coincide for the symmetric case $b = \frac{1}{6}$. To see why, we compare the three bidder case with the two bidder case: With two bidders, for any symmetric probability distribution, revenues were independent of the package allocation, e.g. $\pi_{20}(\frac{1}{6}) = \pi_{11}(\frac{1}{6})$. From this we can conclude that introducing a third, uninformed bidder must change the picture: $\pi_{20}(b) = \pi_{200}(b)$ and $\pi_{11}(b) < \pi_{110}(b)$ as two uninformed bidders are not more useful to the seller than one uninformed bidder. This implies $\pi_{200}(\frac{1}{6}) < \pi_{110}(\frac{1}{6})$. Hence with a symmetric probability distribution F and more than two bidders a fixed amount of information should rather be spread equally among two bidders than concentrated at one bidder. Thus there is no reason to expect the curves in Figure 5.3 to coincide at $\frac{1}{6}$.

We close our look at the case of more than two bidders with the following observation: In contrast to the two bidder case and to our previous example, it can sometimes be strictly optimal to choose a symmetric allocation of packages when there are more bidders.

Example 5.5. *Assume the seller wants to allocate three packages of information among at least three bidders. If the packages are distributed uniformly on $[0, 1]$ or exponentially with parameter 1 it is strictly most profitable for the seller to give the three packages to three different bidders.*

5.6 Conclusion

We have studied an independent, private values, second price auction with entry fees in which the seller can split up total amounts of information differently among the bidders. We have found that if giving out information is costly, the seller will often decide not to provide all the information he has. Clearly, forcing the seller to choose a “fair” allocation in a setting like ours must be disadvantageous for the seller (or harmless at its best). We find that the disadvantage can be huge: In the two bidder case, choosing a “fair” allocation of information is essentially the worst decision the seller can make. Any other split-up of packages (at least into even numbers) would lead to higher revenues. Furthermore, we have seen that the restriction to “fair” allocations can lower the seller’s revenues (and overall welfare) substantially.

Several examples in the chapter have shown that the optimal information policy may depend sensitively on the probability distribution of the information packages. Allocating information the best way is thus a procedure that demands careful examination of the setting. To find a good information policy, there is no simple rule of thumb that can replace a thorough investigation of how bidders may incorporate additional information.

We want to point out that the difficulties we find in our simple model indicate that a more general analysis must be highly complex. For instance, we have seen that the optimal split-up of information depends very sensitively on how the bidders' valuations are distributed. We saw as well that returns to giving out information are not monotonically decreasing. It remains a challenge to find a more tractable model that still allows for an easy and quite natural cardinal ranking of informativeness. All classes of models with release of information that contain our model of independent information packages will suffer from non-monotonicity and sensitivity with regard to distributional assumptions.

Bibliography

[1] Allen, Beth and Martin Hellwig. Bertrand-Edgeworth Oligopoly in Large Markets. Review of Economic Studies, 8: 175-204, 1986.

[2] Allen, Beth and Martin Hellwig. The Approximation of Competitive Equilibria by Bertrand-Edgeworth Equilibria in Large Markets. Journal of Mathematical Economics, 18: 103-127, 1989.

[3] Allen, Beth and Martin Hellwig. Bertrand-Edgeworth Duopoly with Proportional Residual Demand. International Economic Review, 34: 39-60, 1993.

[4] Amann, Erwin and Wolfgang Leininger. Asymmetric All-Pay Auctions with Incomplete Information: The Two-Player Case. Games and Economic Behavior, 14, 1-18, 1996.

[5] Amegashie, J. Atsu and Ximing Wu. Self-selection in competing all-pay auctions. Working Paper, University of Guelph, 2006.

[6] Archak, Nikolay and Arun Sundararajan. Optimal Design of Crowdsourcing Contests. Proceedings of the Thirtieth International Conference on Information Systems, 2009.

[7] Asch, Steven M., Elizabeth A. McGlynn, Mary M. Hogan, Rodney A. Hayward, Paul Shekelle, Lisa Rubenstein, Joan Keesey, John Adams and Eve A. Kerr. Comparison of Quality of Care for Patients in the Veterans Health Administration and Patients in a National Sample. Annals of Internal Medicine, 141(12): 938-945, 2004.

[8] Bagwell, Kyle (ed.). The Economics of Advertising. Edward Elgar Publishing, Cheltenham, 2001.

[9] Banerjee, Abhijit and Lawrence H. Summers. On Frequent-Flyer Programs and other Loyalty-Inducing Economic Arrangements. Working Paper, Harvard University, 1987.

[10] Barlow, Richard and Frank Proschan. Statistical Theory of Reliability and Life Testing. Holt, Rinehart and Winston, New York, 1981.

[11] Barut, Yasar, and Dan Kovenock. The Symmetric Multiple Prize All-Pay Auction With Complete Information. European Journal of Political Economy, 14: 627-644, 1998.

[12] Baye, Michael, Dan Kovenock and Casper de Vries. It takes two to tango: Equilibria in a model of sales. Games and Economic Behavior, 4: 493-510, 1992.

[13] Baye, Michael, Dan Kovenock and Casper de Vries. The all-pay-auction with complete information. Economic Theory, 8: 291-305, 1996.

[14] Baye, Michael R. and John Morgan. A Folk Theorem for One-Shot Bertrand Games. Economics Letters, 65: 59-65, 1999.

[15] Beggs, Alan and Paul Klemperer. Multi-Period Competition with Switching Costs. Econometrica, 60: 651-666, 1992.

[16] Bergemann, Dirk and Martin Pesendorfer. Information structures in optimal auctions. Journal of Economic Theory, 137: 580-609, 2007.

[17] Bergemann, Dirk and Juuso Välimäki. Information in Mechanism Design. in: Richard Blundell, Whitney Newey and Torsten Persson, eds., Proceedings of the 9th World Congress of the Econometric Society. Cambridge University Press, Cambridge, 2006, 186-221.

[18] Bergemann, Dirk and Juuso Välimäki. Information Acquisition and Efficient Mechanism Design. Econometrica, 70: 1007-1034, 2002.

[19] Bergemann, Dirk, Xianwen Shi and Juuso Välimäki. Information Acquisition in Interdependent Value Auctions. Journal of the European Economic Association, 7(1): 61-89, 2009.

[20] Bester, Helmut and Emmanuel Petrakis. Coupons and Oligopolistic Price Discrimination. International Journal of Industrial Organization, 14: 227-242, 1996.

[21] Borgida, Eugene and Richard E. Nisbett. The differential impact of abstract versus concrete information on decisions. Journal of Applied Social Psychology, 7: 258- 71, 1977.

[22] Brooks, Janet R. Veterans' health system blazing trails. Canadian Medical Association Journal, 179: 518-519, 2008.

[23] Bulow, Jeremy I. and Paul Klemperer. Auctions versus negotiations. American Ecomonic Review 86: 180-194, 1996.

[24] Caminal, Ramon and Carmen Matutes. Endogenous Switching Costs in a Duopoly Model. International Journal of Industrial Organization, 8: 353-374, 1990.

[25] Chade, Hector and Edward Schlee. Another Look at the Radner-Stiglitz Nonconcavity in the Value of Information. Journal of Economic Theory, 107(2): 421-452, 2002.

[26] Che, Yeon-Koo and Ian Gale. Caps on Political Lobbying. American Economic Review, 88: 643-651, 1998.

[27] Chen, Yongmin. Paying Customers to Switch. Journal of Economics & Management Strategy, 6: 877-897, 1997.

[28] Chilton, Lydia and John J. Horton. The Labor Economics of Paid Crowdsourcing. Proceedings of the 11th ACM Conference on Electronic Commerce, to appear, 2010.

[29] Clark, Derek J. and Christian Riis. Competition Over More Than One Prize. American Economic Review, 88, 276-289, 1998.

[30] Compte, Olivier and Philippe Jehiel. Auctions and information acquisition: sealed bid or dynamic formats? RAND Journal of Economics, 38(2): 355-372, 2007.

[31] Cramton, Peter and Alan Schwartz. Using Auction Theory to Inform Takeover Regulation. Journal of Law, Economics, and Organization, 7: 27-53, 1991.

[32] Crémer, Jacques, Yossi Spiegel and Charles Zheng. Optimal search auctions. Journal of Economic Theory, 134: 226-248, 2007.

[33] Crémer, Jacques, Yossi Spiegel and Charles Zheng. Auctions with costly information acquisition. Economic Theory, 38(1): 41-72, 2009.

[34] Dasgupta, Sudipto and Kevin Tsui. A 'matching auction' for targets with heterogeneous bidders. Journal of Financial Intermediation, 12: 331-364, 2003.

[35] David, Herbert A. Order Statistics. Wiley, New York, 1970.

[36] Deneckere, Raymond, Dan Kovenock and Robert Lee. A Model of Price Leadership Based on Consumer Loyalty. Journal of Industrial Economics, 40: 147-156, 1992.

[37] DiPalantino, Dominic and Milan Vojnovic. Crowdsourcing and All-Pay Auctions. Proceedings of the 10th ACM International Conference on Electronic Commerce, 119-128, 2009.

[38] Ellison, Glenn. Bounded Rationality in Industrial Organization. In: R. Blundell, W. Newey and T. Persson (eds.), Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Cambridge University Press, 2006.

[39] Ellison, Glenn and Drew Fudenberg. Word-of-Mouth Communication and Social Learning. Quarterly Journal of Economics, 110: 93-125, 1995.

[40] Enkin, Murray W. and Alejandro R. Jadad. Using anecdotal information in evidence-based health care: Heresy or necessity? Annals of Oncology, 9: 963-966, 1998.

[41] Esö, Peter and Balasz Szentes. Optimal Information Disclosure in Auctions. Review of Economic Studies, 74: 705-731, 2007.

[42] Fagerlin, Angela, Catherine Wang and Peter A. Ubel. Reducing the influence of anecdotal reasoning on people's health care decisions: is a picture worth a thousand statistics? Medical Decision Making, 25(4): 398-405, 2005.

[43] Faigle, Ulrich, Walter Kern and Georg Still. Algorithmic Principles of Mathematical Programming. Kluwer, Dordrecht, 2002.

[44] Feller, William. An Introduction to Probability Theory and Its Applications. Volume II. 2nd edition, Wiley, New York, 1971.

[45] Fishman, Michael. A Theory of Preemptive Takeover Bidding. RAND Journal of Economics, 19(1): 88-101, 1988.

[46] Finucane, Melissa L., Ali Alhakami, Paul Slovic and Stephen M. Johnson. The affect heuristic in judgments of risks and benefits. Journal of Behavioral Decision Making, 13: 1-17, 2000.

[47] Fontanarosa, Phil B. and George D. Lundberg. Alternative Medicine Meets Science. Journal of the American Medical Association, 280: 1618-1619, 1998.

[48] Ganuza, Juan-José and José Penalva. Signal Orderings based on Dispersion and the Supply of Private Information in Auctions. Econometrica, 78(3): 1007-1030, 2010.

[49] Garfield, Joan and Andrew Ahlgren. Difficulties in Learning Basic Concepts in Probability and Statistics: Implications for Research. Journal for Research in Mathematics Education, 19(1): 44-63, 1988.

[50] Gattellari, Melina, Jeanette E. Ward and Michael J. Solomon. Randomized, controlled trials in surgery. Diseases of the Colon & Rectum, 44(10): 1413-1420, 2001.

[51] German Federal Ministry of Health, Gesundheitsfonds fördert Wettbewerb, News Release, 08 January 2009, http://www.bundesregierung.de/nn_774/Content/DE/Artikel/2009/01/2009-01-08-gesundheitsfonds.html

[52] Glenton, Claire, Elin S. Nilsen and Benedicte Carlsen. Lay perceptions of evidence-based information - a qualitative evaluation of a website for back pain sufferers. BMC Health Services Research, 6: 34, 2006.

[53] Giammarino, Ronald and Robert Heinkel. A Model of Dynamic Takeover Behavior. Journal of Finance, 41(2): 465-480, 1986.

[54] Grimmett, Geoffrey R. and David R. Stirzaker. Probability and Random Processes, 2nd Edition, Clarendon Press, Oxford, 1992.

[55] Hagedorn, Marcus. The Value of Information for Auctioneers. Journal of Economic Theory, 144(5), 2197-2208, 2009.

[56] Hamill, Ruth, Timothy D. Wilson and Richard E. Nisbett. Insensitivity to sample bias: generalizing from atypical cases. Journal of Personality and Social Psychology, 39: 578-89, 1980.

[57] Hermalin, Benjamin. Lecture Notes for Economics. UC Berkeley, 2009.

[58] Hillman, Arye L. and John G. Riley. Politically Contestable Rents and Transfers. Economics and Politics, 1: 17-39, 1989.

[59] Hillman, Arye L. and Dov Samet. Dissipation of Contestable Rents by Small Numbers of Contenders. Public Choice, 54: 63-82, 1987.

[60] Hirshleifer, David and Ivan Png. Facilitation of Competing Bids and the Price of a Takeover Target. Review of Financial Studies, 2(4): 587-606, 1989.

[61] Hoffmann, Florian and Roman Inderst. Price discrimination and the Provision of Information. Working Paper, University of Frankfurt, 2009.

[62] Howe, Jeff. The Rise of Crowdsourcing. Wired Magazine, 14(6), 2006.

[63] Ireland, Norman. The provision of consumer information in a Bertrand oligopoly. Journal of Industrial Economics, 41: 61-76, 1993.

[64] Jagers, Bert, Wilbert Kallenberg and Albert Kroese. Problem 314. Proposed by A.A. Jagers and W.C.M. Kallenberg, Solution by A.H. Kroese. Statistica Neerlandica, 49: 262-263, 1995.

[65] Jain, Shaili and David Parkes. The Role of Game Theory in Human Computation Systems. Proceedings of the Human Computation Workshop (KDD-HCOMP'09), 2009.

[66] Jenni, Karen E. and George F. Loewenstein. Explaining the "identifiable victim effect". Journal of Risk and Uncertainty, 14: 235-237, 1997.

[67] Johnson, Justin and David Myatt. On the Simple Economics of Advertising, Marketing, and Product Design. American Economic Review, 96(3): 756-784, 2006.

[68] Kahneman, Daniel and Amos Tversky. On the psychology of prediction. Psychological Review, 80: 237-51, 1973.

[69] Kallmes, David F., Bryan A. Comstock, Patrick J. Heagerty, Judith A. Turner, David J. Wilson, Terry H. Diamond, Richard Edwards, Leigh A. Gray, Lydia Stout, Sara Owen, William Hollingworth, Basavaraj Ghdoke, Deborah J. Annesley-Williams, Stuart H. Ralston and Jeffrey G. Jarvik. A Randomized Trial of Vertebroplasty for Osteoporotic Spinal Fractures. The New England Journal of Medicine, 361: 569-579, 2009.

[70] Kirkley, Alexandra, Trevor B. Birmingham, Robert B. Litchfield, J. Robert Giffin, Kevin R. Willits, Cindy J. Wong, Brian G. Feagan, Allan Donner, Sharon H. Griffin, Linda M. D'Ascanio, Janet E. Pope and Peter J. Fowler. A Randomized Trial of Arthroscopic Surgery for Osteoarthritis of the Knee. The New England Journal of Medicine, 359: 1097-1107, 2008.

[71] Klemperer, Paul. Markets with Consumer Switching Costs. Quarterly Journal of Economics, 102: 375-394, 1987a.

[72] Klemperer, Paul. The Competitiveness of Markets with Switching Costs. RAND Journal of Economics, 18: 138-150, 1987b.

[73] Klemperer, Paul. Entry Deterrence in Markets with Consumer Switching Costs. Economic Journal, 97: 99-117, 1987c.

[74] Klemperer, Paul. Competition when Consumers have Switching Costs. An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade. Review of Economic Studies, 62: 515-539, 1995.

[75] Klemperer, Paul. Auctions: Theory and Practice. Princeton University Press, Princeton, 2004.

[76] Konrad, Kai A. Altruism and envy in contests: an evolutionary stable symbiosis. Social Choice and Welfare, 22: 479-490, 2004.

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

[78] Konrad, Kai A. and Dan Kovenock. Contests with stochastic abilities. Economic Inquiry, 48: 89-103, 2010.

[79] Kovenock, Dan, Florian Morath and Johannes Münster. Information Sharing in contests. Working Paper, 2009, www.vwl1.ovgu.de/vwl1_media/downloads/paper/information_sharing_2009_06_25.pdf

[80] Krishna, Vijay. Auction Theory. Academic Press, San Diego, 2002.

[81] Lehmann, Erich. Comparing Location Experiments. Annals of Statistics, 16(2): 521-533, 1988.

[82] Levin, Dan and Janus L. Smith. Equilibrium in Auctions with Entry. American Economic Review, 84: 585-599, 1994.

[83] Lewis, Tracy and David Sappington. Supplying Information to Facilitate Price Discrimination. International Economic Review, 35: 309-327, 1994.

[84] Li, Xiaohu. A note on expected rent in auction theory. Operations Research Letters, 33(5): 531-534, 2005.

[85] Li, Xiaohu and Moshe Shaked. The observed total time on test and the observed excess wealth. Statistics & Probability Letters, 68(3): 247-258, 2004.

[86] Lim, Wooyoung and Alexander Matros. Contests with a Stochastic Number of Players. Games and Economic Behavior, 67, 584-597, 2009.

[87] Lipkus, Isaac M., Gregory Samsa and Barbara K. Rimer. General performance on a numeracy scale among highly educated samples. Medical Decision Making, 21(1): 37-44, 2001.

[88] Loewenstein, George F., Elke U. Weber, Christopher K. Hsee and Ned Welch. Risk as feelings. Psychological Bulletin, 127(2): 267-86, 2001.

[89] Longman, Phillip. The Best Care Anywhere: Why VA Health Care Is Better Than Yours. Second Edition, Polipoint Press, Sausalito, 2010.

[90] Lu, Jingfeng. Optimal Entry in Auctions with Valuation Discovery Costs. Applied Economics Research Bulletin, 2(1), 22-30, 2008.

[91] Maskin, Eric S. and John G. Riley. Auction Theory with Private Values. American Economic Review, 75: 150-156, 1985.

[92] McAfee, R. Preston. Endogenous Availability, Cartels and Merger in an Equilibrium Price Dispersion. Journal of Economic Theory, 62(1): 24-47, 1994.

[93] McAfee, R. Preston and John McMillan. Auctions with Entry. Economics Letters, 23: 343-347, 1987.

[94] McCulloch, Peter, Irving Taylor, Mitsuru Sasako, Bryony Lovett and Damian Griffin. Randomised trials in surgery: problems and possible solutions. BMJ, 324: 1448-1451, 2002.

[95] Milgrom, Paul. Putting Auction Theory to Work. Cambridge University Press, Cambridge, 2004.

[96] Moldovanu, Benny, and Aner Sela. Contest Architecture. Journal of Economic Theory, 126: 70-97, 2006.

[97] Moldovanu, Benny, Aner Sela and Xianwen Shi. Competing Auctions with Endogenous Quantities. Journal of Economic Theory, 141: 1-27, 2008.

[98] Münster, Johannes. Contests with an Unknown Number of Contestants. Public Choice, 129: 353-368, 2006.

[99] Münster, Johannes. Repeated contests with asymmetric information, Journal of Public Economic Theory. 11(1): 89-118, 2009.

[100] Myerson, Roger B. Optimal Auction Design. Mathematics of Operations Research, 6: 58-73, 1981.

[101] Myerson, Roger B. and Karl Wärneryd. Population Uncertainty in Contests. Economic Theory, 27: 469-474, 2006.

[102] Osborne, Martin J. and Ariel Rubinstein. A course in game theory. MIT Press, Cambridge, 1994.

[103] Osborne, Martin J. and Ariel Rubinstein. Games with Procedurally Rational Players. American Economic Review, 88: 834-849, 1998.

[104] Padilla, A. Jorge. Mixed Pricing in Oligopoly with Consumer Switching Costs. International Journal of Industrial Organization, 10: 393-411, 1992.

[105] Parreiras, Sergio O. and Anna Rubinchik. Contests with three or more heterogeneous agents. Games and Economic Behavior, 68(2): 703-715, 2010.

[106] Paul, Anand and Genaro Gutierrez. Mean sample spacings, sample size and variability in an auction-theoretic framework. Operations Research Letters, 32: 103-108, 2004.

[107] Persico, Nicola. Information Acquisition in Auctions. Econometrica, 68: 135-148, 2000.

[108] Radner, Roy and Joseph Stiglitz. A Nonconcavity in the Value of Information. In: "Bayesian Models in Economic Theory", Marcel Boyer and Richard Kihlstrom (eds.), Elsevier Science Publishers, New York, 1984.

[109] Rubinstein, Ariel and Ran Spiegler. Money Pumps in the Market. Journal of the European Economic Association, 6(1): 237-253, 2008.

[110] Schwartz, Lisa M., Steven Woloshin, William C. Black and H. Gilbert Welch. The role of numeracy in understanding the benefit of screening mammography. Annals of Internal Medicine, 127: 966-972, 1997.

[111] Shaffer, Greg and Z. John Zhang. Pay to Switch or Pay to Stay: Preference-Based Price discrimination in Markets with Switching Costs. Journal of Economics & Management Strategy, 9: 397-424, 2000.

[112] Shaked, Moshe and J. George Shantikumar. Stochastic Orders. Springer, New York, 2007.

[113] Shapiro, Carl. Advertising and Welfare: Comment. Bell Journal of Economics, 11 (2): 749-752, 1980.

[114] Shi, Xianwen. Optimal Auctions with Information Acquisition. Working Paper, 2007.

[115] Siegel, Ron. All-Pay Contests. Econometrica, 77, 71-92, 2009.

[116] Siegel, Ron. Asymmetric Contests with Conditional Investments. American Economic Review, 2010, forthcoming.

[117] Simon, Leo K. and William R. Zame. Discontinuous Games and Endogenous Sharing Rules. Econometrica, 58: 861-872, 1990.

[118] Spiegler, Ran. The Market for Quacks. Review of Economic Studies, 73(4): 1113-1131, 2006a.

[119] Spiegler, Ran. Competition over Agents with Boundedly Rational Expectations. Theoretical Economics, 1(2): 207-231, 2006b.

[120] Szech, Nora. Becoming a Bad Doctor. Working Paper, Bonn Graduate School of Economics, 2010.

[121] Tirole, Jean. The Theory of Industrial Organization. MIT Press, Cambridge, 1988.

[122] Tversky, Amos, and Daniel Kahneman. Belief in the law of small numbers. Psychological Bulletin. 76(2): 105-110, 1971.

[123] Vagstad, Steinar. Should auctioneers supply early information for prospective bidders? International Journal of Industrial Organization, 25: 597-614, 2007.

[124] Varian, Hal. A Model of Sales. American Economic Review, 70(4): 651-659, 1980.

[125] Weinstein, Neil D. What does it mean to understand a risk? Evaluating risk comprehension. Journal of the National Cancer Institute Monographs, 25: 15-20, 1999.

[126] Weinstein, Neil D., Kathy Atwood, Elaine Puleo, Robert Fletcher, Graham Colditz, and Karen M. Emmons. Colon cancer: risk perceptions and risk communication. Journal of Health Communication, 9(1): 53-65, 2004.

[127] Xu, Maochao and Xiaohu Li. Some Further Results on the Winner's Rent in the Second-Price Business Auction. Sankhya, Volume 70-A (1), 124-133, 2008.

[128] Yamagishi, Kimihiko. When a 12.86% mortality is more dangerous than 24.14%: implications for risk communication. Applied Cognitive Psychology, 11: 495-506, 1997.

Appendices

A Proofs of Chapter 1

Proof of Proposition 1.1

The proof is structured as follows: We collect observations about the shape of the equilibrium until we can calculate a unique equilibrium candidate. It is then easy to verify that this candidate is indeed an equilibrium.

The observations made in the main text together with the usual arguments for complete information all-pay auctions yield that

- the strategies' supports must go down to zero for some types,
- at most one bidder sets an atom on effort-level zero,
- no bidder mixes over intervals on which his opponent is inactive,
- bidders' supports do not have gaps, in particular, there is no gap between the strategies of strong i and weak i,
- the supports of strong i and weak i are distinct except that possibly both set an atom in zero, where strong i mixes over higher effort levels than weak i,
- if on an interval I strong i and weak j are active, strong i mixes with density C_j/p_i and weak j mixes with density $c_i/(1-p_j)$. Generally a type's density is always the quotient of the opponent type's costs and his own type probability.

Thus we know that the supports form an interval $[0, \overline{e}]$ for some $\overline{e} > 0$. We also know that there are points e_1 and e_2 in this interval such that weak i mixes over $[0, e_i]$ and strong i mixes over $[e_i, \overline{e}]$. Additionally, the weak type of one of the bidders may place an atom on zero. Finally, we explicitly know the densities chosen by the different types against different opponents.

We can thus sequentially calculate an equilibrium candidate: Fix some value of $\overline{e}$. Let the strong types of both bidders mix down from $\overline{e}$ until one of them, say i, has used up all his probability mass. We call this point (known only in reference to $\overline{e}$) e_i. At this point, the weak type of bidder i comes in. Repeat this procedure downwards to the point e_0 where both types of one bidder have used up all their probability mass. The opponent must put his remaining probability mass on an atom in e_0 since both bidders' have to mix over the same support. Note that for any $\overline{e}$, this procedure necessarily produces unique values e_0, e_1 and e_2. Now recall that e_0 must equal zero in equilibrium. This uniquely determines the values of e_1, e_2 and $\overline{e}$. We thus have found a unique equilibrium candidate. It is tedious but straightforward to verify that this sequential procedure leads to the equilibrium candidate stated in the proposition and that it is indeed an equilibrium. □

Proof of Proposition 1.2

It is easy to verify that the equilibrium is well-defined, i.e., that the e_i form an increasing sequence, that the F_i are continuous at the concatenation points and increasing from zero to one. Note that strong i's expected payoff $\pi_i(e)$ from playing e while his strong-type opponents utilize the strategies F_k is given by

$$\pi_i(e) = \prod_{k \neq i} (p_k F_k(e) + (1 - p_k)) - ce. \tag{A.1}$$

It is easy to verify by inserting the definition of F_k from the proposition into (A.1) that $\pi_i(e)$ is constant over $[e_i, e_{n+1}]$ and equals $\prod_{j=2}^{n}(1 - p_j)$.[1] It is also obvious that no bidder wishes to deviate to effort levels above e_{n+1}.

The following argument shows that no bidder i wishes to deviate to effort levels in the interval $[0, e_i]$: Consider $\pi_i(e)$ outside the support of bidder i, i.e. on the interval $[e_k, e_{k+1}]$ with $i \geq k + 1$. Using the fact that only the k bidders with the largest p_l

[1] Equation (A.1) is also the key ingredient for finding the F_k: We know (see the main text) that the equilibrium payoffs of the strong types equal $\prod_{j=2}^{n}(1 - p_j)$. Thus the right hand sides of (A.1) are known and must be identical for all i who are active at e. This allows us to identify an equilibrium candidate.

are active here, (A.1) becomes

$$\pi_i(e) = \prod_{j>k, j\neq i} (1 - p_j) \prod_{l=1}^{k} (p_l F_l(e) + (1 - p_l)) - ce. \tag{A.2}$$

Inserting the definitions of the F_l, this becomes

$$\pi_i(b) = \prod_{j>k, j\neq i} (1 - p_j) \left(\frac{\prod_{l=2}^{n} (1 - p_l) + ce}{\prod_{l=k+1}^{n} (1 - p_l)} \right)^{\frac{k}{k-1}} - ce. \tag{A.3}$$

A few algebraic manipulations turn (A.3) into

$$\pi_i(e) = (\prod_{l=2}^{n} (1 - p_l) + ce) \left(\frac{\prod_{l=2}^{n} (1 - p_l) + ce}{(1 - p_i)^{k-1} \prod_{l=k+1}^{n} (1 - p_l)} \right)^{\frac{1}{k-1}} - ce$$

and then into

$$\pi_i(e) = \prod_{l=2}^{n} (1 - p_l) + (\prod_{l=2}^{n} (1 - p_l) + ce) \left(\sqrt[k-1]{\frac{\prod_{l=2}^{n} (1 - p_l) + ce}{(1 - p_i)^{k-1} \prod_{l=k+1}^{n} (1 - p_l)}} - 1 \right).$$

Note that the first summand on the right hand side is bidder i's payoff on the support of his strategy F_i. Thus, in order to show that bidder i does not want to deviate, we have to consider the sign of the remainder of the expression. To see that the sign is negative, we have to show that the fraction

$$\rho(e) = \frac{\prod_{l=2}^{n} (1 - p_l) + ce}{(1 - p_i)^{k-1} \prod_{l=k+1}^{n} (1 - p_l)}$$

under the $(k+1)^{\text{th}}$-root is smaller than 1. Since $\rho(e)$ is increasing in e it is sufficient to verify that $\rho(e) \leq 1$ for the largest e in the interval, $e = e_{k+1}$. Inserting the definition of e_{k+1} gives us

$$\rho(e_{k+1}) = \frac{(1 - p_{k+1})^{k-1}}{(1 - p_i)^{k-1}} \leq 1.$$

Thus bidder i cannot gain from deviating to effort levels in $[0, e_i]$. □

Proof of Proposition 1.3

Recall from Lemma 1.1 that all strong types must earn $\prod_{j\neq 1}(1 - p_j)$ in equilibrium and define $\pi = \prod_{j\neq 1}(1-p_j)$. The main part of this proof consists of showing that the vector of strategies given in Proposition 1.2 is the only equilibrium candidate. Once this has been established, the results about equilibrium existence and tie-breaking follow easily: Tie-breaking only concerns bidder 1 who places an atom on zero.

Clearly, bidder 1 only earns π from playing zero if tie-breaking against low-types is in his favor at zero. If bidder 1 does not earn π in zero, the only equilibrium candidate is not an equilibrium so that no equilibrium exists.[2] The case where $p_1 = p_2$ is special because then no bidder places an atom on zero so that almost surely ties do not occur in equilibrium.

Let us first recall a number of observations which are similar to the complete information case: Any equilibrium must be in mixed strategies. Some strong types must mix down to zero. There are no atoms with the exception that at most one bidder places an atom on zero. As a consequence, the distribution functions of all bidders' equilibrium strategies must be continuous for positive efforts. The union of all bidders' supports cannot have gaps. Likewise, there are no intervals on which only one bidder is active. Now the key part of the proof lies in showing the following two steps:

Step 1: Consider any equilibrium and denote by H_i bidder i's equilibrium strategy. Denote by A the set of bidders who is actively mixing on an interval of effort-levels I. Assume that I and A are such that bidders in A are active on the whole interval I. Then there exists a constant D which is independent of i and e such that H_i is given by

$$H_i(e) = \frac{1}{p_i}\left[\sqrt[|A|-1]{\frac{\pi + ce}{D}} - (1 - p_i)\right]$$

for all $e \in I$.

Step 2: Consider $e^* > 0$, $\varepsilon > 0$ and i such that bidder i is active on $(e^* - \varepsilon, e^*]$ but inactive on $(e^*, e^* + \varepsilon)$. Then for sufficiently small ε none of the bidders is active on $(e^*, e^* + \varepsilon)$.

Note that the expression for H_i in Step 1 is well-defined because $|A| \geq 1$ implies $|A| \geq 2$. Before proving these steps, let us see how to conclude the proof from here on. Since there are no gaps, Step 2 implies that all bidders must mix up to the same highest effort level. Recalling payoff-uniqueness we know that this highest effort level must then be $\bar{e} = \frac{1}{c}(1 - \pi)$. Furthermore, Step 2 implies that each

[2]Note that our payoff-uniqueness result was independent of tie-breaking. In particular, bidder 1 must earn at least π in any equilibrium because from a bid marginally above zero he can secure a payoff arbitrarily close to π regardless of his opponents' behavior.

bidder's support must be a gapless interval. Armed with these insights we can start to explicitly calculate an equilibrium candidate: We know all bidders are active on an interval below $\overline{e}$. Step 1 specifies the distribution functions on this interval up to a constant D. We can calculate D for this interval from the identity $H_i(1) = 1$. Moving downwards from $\overline{e}$ we get to a point where some of the H_i become zero. All other bidders must remain active and their H_i must be continuous. Iterating this construction until all but one H_i becomes zero allows us to calculate a unique equilibrium candidate. It is straightforward to see that this unique candidate is indeed the vector of strategies given in Proposition 1.2.

Proof of Step 1

For $i \in A$ consider the payoff $\pi_(e)$ from making an effort $e \in I$,

$$\pi_i(e) = \prod_{j \in A \setminus \{i\}} (p_j H_j(e) + (1 - p_j)) \prod_{k \in A^c} (p_k H_k(e) + (1 - p_k)) - ce. \tag{A.4}$$

Note that by assumption the product over A^c is for $e \in I$ a constant which is independent of e and i. Denote this constant by D. Recall that by payoff-uniqueness we also have $\pi_i(e) = \pi$. Hence (A.4) becomes

$$\frac{ce + \pi}{D} = \prod_{j \in A \setminus \{i\}} (p_j H_j(e) + (1 - p_j))$$

for all $i \in A$. But since the left hand side of this identity is independent of i so must be the right hand side. This implies that all factors on the right hand side must be identical - the product is the same regardless of which bidder i is taken out. We can thus conclude that

$$\frac{ce + \pi}{D} = (p_i H_i(e) + (1 - p_i))^{|A|-1}$$

for all $i \in A$. Solving for H_i shows Step 1.

Proof of Step 2

Choose ε small enough such that there is a fixed set of active bidders on $I = (e^*, e^* + \varepsilon)$ and we can apply Step 1. Consider bidder i's payoff from deviating into the interval I,

$$\pi_i(e) = \prod_{j \in A} (p_j H_j(e) + (1 - p_j)) \prod_{k \in A^c \setminus \{i\}} (p_k H_j K(e) + (1 - p_k)) - ce. \tag{A.5}$$

Again, the second product is independent of e. Inserting the expression for H_i from Step 2 into the first product and collecting the constant D and the second product in a new constant E yields

$$\pi_i(e) = (ce + \pi)^{\frac{|A|}{|A|-1}} E - ce.$$

After some algebraic manipulations this becomes

$$\pi_i(e) = \pi + (ce + \pi)(E \sqrt[|A|-1]{ce + \pi} - 1).$$

Now since we assumed that $\pi_i(e^*) = \pi$ we know that E must be such that the second factor in the second summand equals zero for $e = e^*$. But this implies that $\pi_i(e)$ is increasing for $e > e^*$ contradicting our assumption of being in equilibrium. This proves Step 2. □

Proof of Lemma 1.2

In the following we consider the more general game where bidders choose probabilities p_1 and p_2 in order to maximize their payoffs given by

$$\pi_1(p_1, p_2) = (p_1(1 - \min(p_1, p_2))\alpha + (1 - p_1)(\max(0, p_1 - p_2))\beta$$

where $\alpha \geq \beta > 0$. The game in the lemma is obviously the special case $\alpha = \beta = \theta$ of this game. This more general formulation is more adaptable to generalizations.[3] One can easily verify that (independently of α and β) the first two types of equilibria stated in the lemma are indeed equilibria. We thus focus on the symmetric mixed strategy equilibrium. In order to identify an equilibrium candidate we take the following approach: Note that setting p_i to a value smaller than $\frac{1}{2}$ is strictly dominated by setting p_i to $\frac{1}{2}$. Assume thus that both bidders mix with a density f over $[\frac{1}{2}, 1]$.[4] The payoff of bidder 1 from choosing p_1 while his opponent mixes with density f can be written as

$$\pi_1(p_1) = \alpha p_1 \left[\int_{\frac{1}{2}}^{p_1} (1 - p_2) f(p_2) dp_2 + \int_{p_1}^{1} (1 - p_1) f(p_2) dp_2 \right] + \beta(1 - p_1) \int_{\frac{1}{2}}^{p_1} (p_1 - p_2) f(p_2) dp_2.$$

Denote by $F(p)$ the cumulative density function associated with f. Simple algebraic manipulations lead to

$$\pi_1(p) = \alpha p_1 \left[(1 - p_1) + p_1 F(p_1) - \int_{\frac{1}{2}}^{p_1} p_2 f(p_2) dp_2 \right] + \beta(1 - p_1) \left[p_1 F(p_1) - \int_{\frac{1}{2}}^{p_1} p_2 f(p_2) dp_2 \right]$$

[3] Notably, the case where the marginal costs of exerting effort are 1 and bidders have either valuation V or v where $V > v > 0$ is covered. In that case it can be deduced from Corollary 1.2 that $\alpha = V - v$ and $\beta = 1 - \frac{v}{V}$.

[4] In fact, this assumption does not turn out to be perfectly true: Except when $\alpha = \beta$, both bidders set an atom in 1. The point of this calculation is to demonstrate how to quickly identify an equilibrium candidate. With a more flexible notation, this slight inconsistency could be removed.

Defining

$$G(p_1) = p_1 F(p_1) - \int_{\frac{1}{2}}^{p_1} p_2 f(p_2) dp_2$$

and simplifying yields

$$\pi_1(p) = \alpha p_1(1 - p_1) + G(p)(\alpha p_1 + \beta(1 - p_1)).$$

By integration by parts we obtain

$$G(p_1) = \int_{\frac{1}{2}}^{p_1} F(p_2) dp_2$$

which implies that $G(\frac{1}{2}) = 0$. It also implies that $F(p) = G'(p)$. Now recall that in a symmetric mixed equilibrium we must have that $\pi_1(p)$ is constant in p, i.e. $\pi_1'(p) = 0$. This gives us the ordinary differential equation

$$0 = \alpha(1 - 2p_1) + G'(p_1)(\alpha p_1 + \beta(1 - p_1)) + G(p)(\alpha - \beta).$$

A solution to this differential equation under the boundary condition $G(\frac{1}{2}) = 0$ is given by

$$G(p) = \frac{\alpha(1 - 2p)^2}{4(\alpha p + (1 - \beta)p)}.$$

This leads to an equilibrium candidate of

$$F(p) = \frac{\alpha(2p - 1)(\alpha + 3\beta + 2(\alpha - \beta)p)}{4(\alpha p + (1 - \beta)p)^2}.$$

It follows that $F(\frac{1}{2}) = 0$ and $F(1) = \frac{3\alpha+\beta}{4\alpha} < 1$. Thus, in order to obtain a distribution function on $[\frac{1}{2}, 1]$ as desired, we have to assume an atom in 1 (except when $\alpha = \beta$). In the case $\alpha = \beta$ we obtain $F(p) = 2p - 1$, implying that F is the uniform distribution on $[\frac{1}{2}, 1]$. One can easily calculate that F is indeed an equilibrium. □

B Proofs of Chapter 2

Proof of Proposition 2.1

In McAfee (1994) it is shown that the vector of strategies in Proposition 2.1 is indeed an equilibrium. McAfee (1994) also shows *payoff uniqueness*, i.e., for all i, healer i's equilibrium payoff is given by

$$\pi_i = \alpha_i C \text{ where } C := \prod_{j=1}^{n-1} (1 - \alpha_i) > 0. \tag{B.1}$$

Observe that the first equality states that all healers' expected equilibrium payoffs *conditional on being recommended* must be identical.

We hence take these results as given and show how to obtain equilibrium uniqueness from this point on. We start with a number of preliminary observations:

- *In any equilibrium healers do not place atoms on prices except for possible atoms on* 0 *or* 1*:* If healer i sets an atom in $p \in (0,1)$, other healers playing prices right above p would want to shift their probability mass to prices marginally below p in order to substantially increase their winning probability (while only marginally decreasing prices). If no other healer played prices right above p, healer i could profitably shift the atom upwards.

- *In addition, at most one healer sets an atom on* 1 *in equilibrium:* If two or more healers played an atom in 1, this would result in a positive probability of ties. Thus at least one healer could profitably deviate by shifting his atom marginally downwards.

- *The union of the healers' strategy supports must go up to* 1*:* Playing higher prices is dominated. If the union of supports went only up to a lower price $p_H < 1$, any healer mixing up to p_H could profitably deviate to playing 1.

- Due to the positive equilibrium payoffs, the union of strategy supports must be bounded away from 0. Denote by $p_L > 0$ the infimum of the union of equilibrium supports.

- *The union of supports must be an interval* $[p_L, 1]$*, i.e. there cannot be any gaps in the union of supports:* If there was an interval $[\underline{p}, \overline{p}] \subset [p_L, 1]$ where no healer was active, a healer who would be playing prices right below $\underline{p}$ could deviate by shifting probability mass from a small interval below $\underline{p}$ to $\overline{p}$, yielding a substantially better price at a marginally lower probability of winning.

- *Furthermore, there cannot be a subset* $[\underline{p}, \overline{p}] \subset [p_L, 1]$ *where only one healer is active:* Such a healer could profitably deviate by concentrating all probability mass of the interval in an atom at $\overline{p}$. He would then receive a higher price at the same probability of winning.

Armed with these insights we turn to the first major step of the proof:

1) *In any equilibrium, the strategy support of each healer must go down to the same* $p_L > 0$. *Furthermore, in any equilibrium,* $p_L = C$.

Proof of 1): Consider two healers i and j with supports S_i and S_j. Assume $p_L^i < p_L^j$ where $p_L^k = \inf S_k$ for $k = i, j$. Then, with positive probability, healer i plays a price from $[p_L^i, p_L^j]$. Healer j's payoff from playing p_L^j must equal his equilibrium payoff $\alpha_j C > 0$. Yet this implies that healer i can earn more than his equilibrium payoff of $\alpha_i C$ by playing p_L^j: Since - unlike healer j - healer i does not compete against healer i (himself) as a possibly cheaper competitor when playing p_L^j, his expected payoff conditional on being recommended must be higher than that of j. This is a contradiction to (B.1). Hence the support of every healer must go down to the same lowest price p_L.

To see that $p_L = C$, note that for all j, healer j's payoff from playing p_L must be $\alpha_j p_L$: The other healers charge higher prices with probability 1. Thus healer j attracts all the patients to whom he is recommended and receives p_L from all of them. This leads to a payoff of $\alpha_j p_L$, which is only consistent with (B.1) if $p_L = C$.

The next step further characterizes the functional form of the healers' equilibrium distribution functions:

2) *Let* $\mathcal{D} \subset \{1, \ldots, n\}$ *denote the set of healers who are active on some interval* $I = (\underline{p}, \overline{p})$ *in some arbitrary but fixed equilibrium. Assume all healers* $j \in \mathcal{D}$ *are active at any* $p \in I$ *and let* $m = \#\mathcal{D}$. *(Note that from our preliminary observations it follows that* $m \neq 1$.*) Then for all* $j \in \mathcal{D}$ *any equilibrium distribution function* $H_j(p)$ *must satisfy for all* $p \in I$

$$H_j(p) = \frac{1}{\alpha_j}\left(1 - \sqrt[m-1]{\frac{L}{p}}\right) \tag{B.2}$$

where the constant $L > 0$ *is independent of* p *and* j. *Moreover,*

$$L = \frac{C}{\prod_{i \in \mathcal{D}^C}(1 - \alpha_i H_i(\underline{p}))}.$$

Proof of 2): Note that for all $j \in \mathcal{D}$ and all $p \in I$ the expected payoff of healer j from playing p must equal the equilibrium payoff of $\alpha_j C$. Using (2.2) and the fact that distribution functions of inactive healers are constant over I, this condition reads

$$\alpha_j C = p\alpha_j \left[\prod_{i\in\mathcal{D}^C}(1-\alpha_i H_i(\underline{p}))\right]\left[\prod_{k\in\mathcal{D}\setminus\{j\}}(1-\alpha_k H_k(p))\right].$$

Rearranging and using the definition of L yields for all $p \in I$ and $j \in \mathcal{D}$

$$\prod_{k\in\mathcal{D}\setminus\{j\}}(1-\alpha_k H_k(p)) = \frac{L}{p}. \tag{B.3}$$

Now consider (B.3) for two different healers $i, j \in \mathcal{D}$. Taking the quotient of (B.3) for i and (B.3) for j yields that for all $p \in I$

$$1 = \frac{1-\alpha_j H_j(p)}{1-\alpha_i H_i(p)}$$

which implies that there is a function $h(p)$ such that $h(p) = \alpha_k H_k(p)$ for all $k \in \mathcal{D}$. Substituting $h(p)$ for $\alpha_k H_k(p)$ on the left hand side of (B.3) and then solving for h yields

$$h(p) = 1 - \sqrt[m-1]{\frac{L}{p}}$$

and thus

$$H_j(p) = \frac{1}{\alpha_j}\left(1 - \sqrt[m-1]{\frac{L}{p}}\right)$$

as required.

The last main step shows that no healer has a gap inside his equilibrium price interval, i.e. no healer is inactive over some range of prices (above p_L) while putting positive probability mass on prices above that range:

3) *For all j the support of healer j's strategy is of the form $[p_L, p_H^j]$ for some $p_L < p_H^j \leq 1$.*

Proof of 3): Assume that some price $\overline{p} > p_L$ is in the support of the strategy of healer j but j is inactive on some interval directly below $\overline{p}$. Choose $\underline{p} < \overline{p}$ such that for all $p \in I = (\underline{p}, \overline{p})$ the set of healers who are active at p is identical. (This is possible since there are no atoms and thus the H_i are continuous.) Denote the set of healers active on I by $\mathcal{D}$. Using (2.2) as in Step 2) we can write the payoff of healer j from playing some $p \in I \cup \{\overline{p}\}$ as

$$\pi_j(p) = \alpha_j p \left[\prod_{i\in\mathcal{D}^C\setminus\{j\}}(1-\alpha_i H_i(\underline{p}))\right]\left[\prod_{k\in\mathcal{D}}(1-\alpha_k H_k(p))\right].$$

Defining the constant factor from the other inactive healers as

$$K := \left[\prod_{i \in \mathcal{D}^C \setminus \{j\}} (1 - \alpha_i H_i(\underline{p})) \right]$$

and making use of (B.2) from the last step, we can express $\pi_j(p)$ as

$$\pi_j(p) = \alpha_j p K \left(\sqrt[m-1]{\frac{L}{p}} \right)^m = \alpha_j K L^{\frac{m}{m-1}} \sqrt[m-1]{\frac{1}{p}}$$

where the constant L is defined as in Step 2. Note that this implies that $\pi_j(p)$ is strictly decreasing in p over $I \cup \{\overline{p}\}$. By assumption, healer j is active at $\overline{p}$ and thus must earn his equilibrium payoff there:

$$\pi_j(\overline{p}) = \alpha_j C.$$

Yet since $\pi_j(p)$ is decreasing, this implies that for $p \in I$

$$\pi_j(p) > \alpha_j C$$

such that healer j can profitably deviate - which is a contradiction.

To conclude the proof, we still have to show that the vector of strategies defined in the proposition is actually the only candidate for an equilibrium. We have seen that all supports start at $p_L = C$ and since healers do not set atoms or leave gaps in their supports, all healers remain active up to the price p_1 where the first healer(s) j have used up their probability mass, i.e. where $H_j(p_1) = 1$. Note that on any interval $[p_L, \overline{p}]$ where all healers are active, all distribution functions are uniquely determined by Step 2. Likewise, p_1 and the set of healers with $H_j(p_1) = 1$ are uniquely pinned down by this. Above p_1, all healers who still have probability mass to spend must remain active. By Step 2, distribution functions above p_1 are again uniquely determined, pinning down in turn the price $p_2 > p_1$ where the next supports end. Continuing this procedure sequentially until $p = 1$ or until all or all but one distribution functions equal 1 determines a unique candidate for an equilibrium. It is easy to calculate that this unique candidate is actually the vector of strategies stated in the proposition, and that this unique candidate is indeed an equilibrium.

□

Proof of Proposition 2.2

The payoff of healer i from playing p while the other healers play H_j is given by

$$\pi_i(p) = p\alpha_i E\left[\prod_{j\neq i}(1-\alpha_j H_j(p))\right] = p\alpha_i \prod_{j\neq i}(1-\overline{\alpha}_j H_j(p)) \qquad \text{(B.4)}$$

by the independence of the α_j. This differs from (2.2) only by a factor of $\overline{\alpha}_i/\alpha_i$ which is independent from p. Thus $H_i(p)$ must be a best response for healer i in the incomplete information game with qualities $\alpha_1, ..., \alpha_n$ as well. (Otherwise the H_i would not form a Nash equilibrium in the complete information game with $\overline{\alpha}_1, ..., \overline{\alpha}_n$.) □

Proof of Proposition 2.3

"⇐" follows almost immediately: If a strategy profile satisfies (2.3), the expected payoff of healer i from playing p while the other healers play $G_j^{\alpha_j}$ is the same as in the equilibrium studied in Proposition 2.2:[5]

$$\pi_i(p) = p\alpha_i \prod_{j\neq i}(1-E[\alpha_j G_j^{\alpha_j}(p)]) = p\alpha_i \prod_{j\neq i}(1-\overline{\alpha}_j H_j(p)). \qquad \text{(B.5)}$$

Hence it does not make any difference for healer i whether his competitors play H_j or $G_j^{\alpha_j}$. Healer i then does not have an incentive to deviate from $G_i^{\alpha_i}$ because this strategy has support in $[p_0, p_i]$, the support of H_i. Thus all healers playing $G_i^{\alpha_i}$ is an equilibrium.

For "⇒", we first verify that if the $G_j^{\alpha_j}$ are equilibrium strategies, we do not have to worry about atoms. This is needed to justify the expression (B.6) for the expected payoffs below. Note first that it is inconsistent with equilibrium behavior for a healer j to play a price $\widetilde{p} < 1$ with positive probability in expectation over α_j:[6] If other healers had probability mass on prices marginally above $\widetilde{p}$, they would shift this mass downwards. If no other healers had probability mass on prices marginally above $\widetilde{p}$, healer j could earn more by shifting his probability mass from $\widetilde{p}$ upwards. Additionally, at most one healer j plays a price of 1 with positive probability in expectation over α_j: If several healers did so, at least one of them would have an incentive to shift probability mass downwards to escape tie-breaking.

[5]Note that (2.3) implies that for all $p < 1$ in expectation over α_j the probability that healer j plays p is zero. We thus do not have to worry about atoms.

[6]Note that we do not rule out in the following that for fixed α_j the distribution function $G_j^{\alpha_j}$ contains atoms. We only show that there are no atoms the other healers can anticipate, i.e., atoms in $E[G_j^{\alpha_j}]$.

Since there are no atoms (except possibly one in 1) we can write the payoff of healer i from playing p while the other healers play $G_j^{\alpha_j}$ as

$$\pi_i(p) = p\alpha_i \prod_{j\neq i}(1 - E[\alpha_j G_j^{\alpha_j}(p)]). \tag{B.6}$$

Define for each healer i a distribution function

$$\widetilde{H}_i(p) = \frac{1}{\overline{\alpha}_i} E\left[\alpha_i G_i^{\alpha_i}(p)\right].$$

It must hold that $\widetilde{H}_i(p) = H_i(p)$: Clearly, with a similar reasoning as before, if all healers playing $G_i^{\alpha_i}$ is an equilibrium, all healers playing $\widetilde{H}_i(p)$ must be an equilibrium as well, as the expected payoffs from playing any price p are identical in both situations:

$$p\alpha_i \prod_{j\neq i}(1 - \overline{\alpha}_j \widetilde{H}_j(p)) = p\alpha_i \prod_{j\neq i}(1 - E[\alpha_j G_j^{\alpha_j}(p)]). \tag{B.7}$$

Note that $\widetilde{H}_i(p)$ does not depend on healer i's private information. From comparing the left hand side of (B.7) with (2.2) we see that all healers playing $\widetilde{H}_i(p)$ is also an equilibrium of the complete information game with qualities $\overline{\alpha}_1, ..., \overline{\alpha}_n$ since payoffs differ only by a constant factor between the two games. From Proposition 2.1 we know that $(H_1, \ldots, H_n)$ is the unique equilibrium of the complete information game. This yields $\widetilde{H}_i = H_i$ which by the definition of $\widetilde{H}_i$ implies (2.3).

Finally, we have to show that the support of $G_i^{\alpha_i}$ lies in $[p_0, p_i]$ for all values of α_i.[7] Note that it does not make any difference for healer i whether his competitors play $(G_j^{\alpha_j})_{\alpha_j}$ or H_j. But playing prices outside $[p_0, p_i]$ against competitors who play H_j is strictly dominated. (This is an easy calculation similar to Step 2 in the proof of Proposition 2.1). Thus, if $G_i^{\alpha_i}$ is a best response to $(G_j^{\alpha_j})_{\alpha_j}$, its support must be included in $[p_0, p_i]$.

That healers' expected payoffs are the same in all equilibria is a direct consequence of our result that (2.3) must hold in all equilibria. □

Proof of Proposition 2.4

We first start with a definition: Define the function $K_i(\alpha_i)$ as

$$K_i(\alpha_i) = \int_0^{\alpha_i} \beta f_i(\beta) d\beta.$$

[7] Note that (2.3) implies already that this holds for F_i-almost all values of α_i.

Note that K_i is strictly increasing since the integrand is positive for $\alpha_i > 0$. Furthermore, $K_i(0) = 0$ and $K_i(1) = \overline{\alpha}_i$.

Among others, we have to show that there exists a family $(G_i^{\alpha_i})_{\alpha_i}$ which satisfies the sufficient conditions for a Nash equilibrium from Proposition 2.3 and which consists only of distributions $G_i^{\alpha_i}(p)$ that put all mass on one price. For this purpose, note that a price setting function $\bar{P}_i(\alpha_i)$ translates into a family of distribution functions via

$$G_i^{\alpha_i}(p) = 1_{\{p \geq \bar{P}_i(\alpha_i)\}}.$$

Thus, for a price setting function, (2.3) becomes

$$\int_0^1 \alpha_i 1_{\{p \geq \bar{P}_i(\alpha_i)\}} f(\alpha_i) d\alpha_i = \overline{\alpha}_i H_i(p). \tag{B.8}$$

Now, to prove the proposition, we have to show that there exists a unique equilibrium in strictly increasing price setting functions. Define a strictly increasing price setting function via

$$\bar{P}_i(\alpha_i) = H_i^{-1}\left(\frac{K_i(\alpha_i)}{\overline{\alpha}_i}\right).$$

H_i^{-1} denotes the inverse of the restriction of H_i to $[p_0, p_i]$. Thus H_i^{-1} is a bijection from $[0, 1]$ to $[p_0, p_i]$. To see that $\bar{P}_i$ is a well-defined bijection from $[0, 1]$ to $[p_0, p_i]$ note also that $K_i(\alpha_i)/\overline{\alpha}_i$ is a bijection from $[0, 1]$ to $[0, 1]$. That $\bar{P}_i$ is strictly increasing follows because K_i and H_i are strictly increasing. Considering the inverse of $\bar{P}_i$, we see that $\bar{P}_i$ satisfies (2.3):

$$\begin{aligned}
& \bar{P}_i^{-1}(p) = K_i^{-1}(\overline{\alpha}_i H_i(p)) \\
\Leftrightarrow\ & K_i(\bar{P}_i^{-1}(p)) = \overline{\alpha}_i H_i(p) \\
\Leftrightarrow\ & \int_0^{\bar{P}_i^{-1}(p)} \alpha_i f_i(\alpha_i) d\alpha_i = \overline{\alpha}_i H_i(p).
\end{aligned} \tag{B.9}$$

As $\bar{P}_i$ is strictly increasing, the final equality is equivalent to (B.8) and thus to (2.3). Since $\bar{P}_i$ only takes values in $[p_0, p_i]$, we have hence shown (making use of Proposition 2.3) that the functions $\bar{P}_i$ form a Nash equilibrium. Furthermore, from (B.9) it is evident that $\bar{P}_i(\alpha_i)$ is the unique monotonically increasing equilibrium price setting function. □

Proof of Proposition 2.5

This proposition is an immediate corollary of results derived in Section 2.3 applied to the symmetric case $F_i = F$. In Proposition 2.3 we show that the healers' expected payoffs are identical in all equilibria. In Proposition 2.2 we prove that there is an equilibrium where the expected payoff of healer i is given by

$$\pi_i = \alpha_i(1-\overline{\alpha})^{n-1}.$$

Proposition 2.4 shows that there is a unique monotonically increasing strategy equilibrium, given by the price setting function

$$\bar{P}(\alpha_i) = H^{-1}\left(\frac{\int_0^{\alpha_i}\beta f(\beta)d\beta}{\overline{\alpha}}\right)$$

where H is defined in Proposition 2.1 as

$$H(p) = \frac{1}{\overline{\alpha}}\left(1 - \frac{1-\overline{\alpha}}{\sqrt[n-1]{p}}\right)$$

with support $[(1-\overline{\alpha})^{n-1}, 1]$. To verify that this is the price setting function stated in Proposition 2.5 we just have to calculate that

$$H^{-1}(k) = \left(\frac{1-\overline{\alpha}}{1-\overline{\alpha}k}\right)^{n-1}.$$

Inserting $k = \int_0^{\alpha}\beta f(\beta)d\beta/\overline{\alpha}$ yields the desired result. □

Proof of Proposition 2.6

Denote by $\alpha_{i:n}$ the i^{th} *lowest* of the values α_1,...,α_n. In the following, we make use of three well-known facts:

First, the density of $\alpha_{i:n}$ is given by[8]

$$f_{i:n}(\alpha) = n\binom{n-1}{i-1}F(\alpha)^{i-1}(1-F(\alpha))^{n-i}f(\alpha).$$

Second, recall the Binomial Theorem: For all $a, b > 0$ and $m \in \mathbb{N}$

$$\sum_{i=0}^{m}\binom{m}{i}a^i b^{m-i} = (a+b)^m. \tag{B.10}$$

Finally, we make use of the fact that

$$E[1-\widetilde{\alpha}|\widetilde{\alpha}<\alpha] \;=\; \int_0^{\alpha}(1-\beta)\frac{f(\beta)}{F(\alpha)}d\beta \;=\; 1 - \frac{1}{F(\alpha)}\int_0^{\alpha}\beta f(\beta)d\beta. \tag{B.11}$$

We now calculate γ_n in order to verify (2.5). Recall that each patient consults the worst healer who is recommended to him. Hence the probability that a patient

[8]Compare for instance David (1970).

consults healer i equals the probability that i is recommended and that all healers worse than i are not recommended. Thus

$$\gamma_n = E\left[\sum_{i=1}^{n} \alpha_{i:n}^2 \prod_{j=1}^{i-1}(1-\alpha_{j:n})\right].$$

Note that this calculation takes into account that if no healer is recommended, the patient receives a quality of zero. Observe that conditional on $\alpha_{i:n}$ taking some value α, $\prod_{j=1}^{i-1}(1-\alpha_{j:n})$ has the same distribution as a product of $i-1$ independent random variables:

$$\begin{aligned}\gamma_n &= \int_0^1 \sum_{i=1}^{n} \alpha^2 E\left[\prod_{j=1}^{i-1}(1-\widetilde{\alpha}_j)\middle|\, \widetilde{\alpha}_1,\ldots,\widetilde{\alpha}_{i-1}<\alpha\right] f_{i:n}(\alpha)d\alpha \\ &= \int_0^1 \sum_{i=1}^{n} \alpha^2 E[1-\widetilde{\alpha}|\widetilde{\alpha}<\alpha]^{i-1} f_{i:n}(\alpha)d\alpha \qquad \text{(B.12)}\end{aligned}$$

where the $\widetilde{\alpha}_j$ and $\widetilde{\alpha}$ are independent and distributed according to F. Plugging the definition of $f_{i:n}$ into (B.12) and rearranging yields

$$\gamma_n = n\int_0^1 \alpha^2 \left[\sum_{i=1}^{n} \binom{n-1}{i-1} \left(F(\alpha)E[1-\widetilde{\alpha}|\widetilde{\alpha}<\alpha]\right)^{i-1}(1-F(\alpha))^{n-i}\right] f(\alpha)d\alpha.$$

After shifting the summation index and inserting (B.11), this becomes

$$\gamma_n = n\int_0^1 \alpha^2 \left[\sum_{i=0}^{n-1} \binom{n-1}{i} \left(F(\alpha)-\int_0^\alpha \beta f(\beta)d\beta\right)^{i}(1-F(\alpha))^{(n-1)-i}\right] f(\alpha)d\alpha.$$

Applying (B.10), we obtain

$$\gamma_n = n\int_0^1 \alpha^2 \left[1-\int_0^\alpha \beta f(\beta)d\beta\right]^{n-1} f(\alpha)d\alpha$$

as we wanted to show.

We still have to prove that

$$\lim_{n\to\infty} \gamma_n = 0.$$

Consider the random variable Γ_n which is the quality of the treatment one fixed patient receives in equilibrium. Note that there are two levels of randomness in Γ_n, the randomness in the α_i and the randomness in the recommendations the patient gets. Clearly

$$E[\Gamma_n] = \gamma_n.$$

We first show that Γ_n converges to zero in probability, i.e., for every $\delta > 0$

$$\lim_{n\to\infty} \text{Prob}\,(\Gamma_n < \delta) = 1.$$

Fix $\delta > 0$. The idea is that for any k we can choose n large enough such that with high probability at least k healers have qualities in $(\frac{\delta}{2}, \delta)$. Each of these is recommended with a probability larger than $\frac{\delta}{2}$. Hence if k is large enough it is very likely that one of them is recommended. Precisely, we have to show that for every $\epsilon > 0$ there is some $n(\epsilon)$ such that

$$\text{Prob}\,(\Gamma_n < \delta) \geq 1 - \epsilon$$

for all $n > n(\epsilon)$. Denote by the random variable D_δ^n the number of healers whose qualities lie in the interval $(\frac{\delta}{2}, \delta)$. Denote by R_δ^n the number of healers with qualities in the interval $(\frac{\delta}{2}, \delta)$ who are recommended.

For any $k \leq n$

$$\begin{aligned} \text{Prob}\,(\Gamma_n < \delta) &\geq \text{Prob}\,(R_\delta^n \geq 1) \\ &\geq \text{Prob}\,(D_\delta^n \geq k)\, \text{Prob}\,(R_\delta^n \geq 1\, | D_\delta^n \geq k) \\ &\geq \text{Prob}\,(D_\delta^n \geq k) \left(1 - (1 - \frac{\delta}{2})^k\right). \end{aligned} \tag{B.13}$$

Choose $k(\epsilon)$ large enough such that

$$\left(1 - (1 - \frac{\delta}{2})^{k(\epsilon)}\right) > \sqrt{1 - \epsilon}.$$

Then choose $n(\epsilon)$ $(= \widetilde{n}(\epsilon, k(\epsilon)))$ large enough such that

$$\text{Prob}\left(D_\delta^{n(\epsilon)} \geq k(\epsilon)\right) > \sqrt{1 - \epsilon}.$$

This is possible because we have assumed $f > 0$ which implies that, independently, each healer has with positive probability a quality in $(\frac{\delta}{2}, \delta)$. Then by (B.13)

$$\text{Prob}\,(\Gamma_n < \delta) > \sqrt{1 - \epsilon}^{\,2} = 1 - \epsilon$$

for all $n > n(\epsilon)$. Thus we have shown that Γ_n converges to zero in probability. Since Γ_n is bounded, this implies that Γ_n converges to zero in mean (see, for instance, Grimmett and Stirzaker (1992)):

$$\lim_{n\to\infty} \gamma_n = \lim_{n\to\infty} E[\Gamma_n] = 0.$$

$\square$

Proof of Remark 2.1

Let $\sum_A$ denote $\sum_{A\subset\{1,\ldots,n\},\, A\neq\emptyset}$ and let $\#A$ denote the number of healers in A. Since price setting does not depend on the realizations of the α_i, all recommended healers have the same chance of offering the lowest price. Thus the expected quality a patient receives is the expectation of the average quality of the recommended healers:

$$\widetilde{\gamma}_n = E\left[\sum_A \frac{\sum_{i\in A}\alpha_i}{\#A} Prob[A \text{ is the set of recommended healers}]\right].$$

Obviously it holds that

$$Prob[A \text{ is the set of recommended healers}] = \prod_{j\in A}\alpha_j \prod_{k\in A^c}(1-\alpha_k).$$

Hence

$$\widetilde{\gamma}_n = E\left[\sum_A \frac{1}{\#A}\sum_{i\in A}\left(\alpha_i^2 \prod_{j\in A\setminus\{i\}}\alpha_j \prod_{k\in A^c}(1-\alpha_k)\right)\right].$$

By the independence of the α_i, this implies

$$\widetilde{\gamma}_n = \sum_A E[\alpha^2]E[\alpha]^{\#A-1}(1-E[\alpha])^{n-\#A}$$

where α is a random variable with distribution F. Since $\{1,\ldots,n\}$ has $\binom{n}{k}$ subsets with k elements we can rewrite this to

$$\widetilde{\gamma}_n = \frac{E[\alpha^2]}{E[\alpha]}\sum_{k=1}^{n}\binom{n}{k}E[\alpha]^k(1-E[\alpha])^{n-k}.$$

By the binomial theorem (B.10) this is the same as

$$\widetilde{\gamma}_n = \frac{E[\alpha^2]}{E[\alpha]}(1-(1-E[\alpha])^n).$$

Hence we are done. $\square$

C Proofs of Chapter 3

Proof of Proposition 3.1

Special case of Proposition 3.7. $\square$

Proof of Proposition 3.2

In order to identify an equilibrium candidate we first assume that the equilibrium is indeed given by a function F_j that can be decomposed into an aggressive and a defensive strategy as sketched in the main text, i.e., $F_j = q_j A_j + (1 - q_j) D_j$. Furthermore we postulate that the support of D_j corresponds to the support of A_i shifted upwards by r_j so that both distributions cover the same range of net prices for the customers in the home base of firm j. This is expressed by the system of equations

$$\overline{a}_i + r_j = \overline{d}_j \tag{C.1}$$

$$\overline{a}_j + r_i = \overline{d}_i \tag{C.2}$$

$$\underline{a}_i + r_j = \underline{d}_j \tag{C.3}$$

$$\underline{a}_j + r_i = \underline{d}_i. \tag{C.4}$$

Solving (3.2) and (3.3) for D_j and A_j and shifting the argument we get the following expressions for D_j and A_j

$$D_j(p) = 1 - \frac{\pi_i - m_i(p - r_i - r_j)}{m_j(p - r_j)(1 - q_j)} \tag{C.5}$$

and

$$A_j(p) = \frac{1}{q_j}\left(1 - \frac{\pi_i}{m_i p}\right). \tag{C.6}$$

From these functions it is easy to calculate $\underline{a}_j, \overline{a}_j, \underline{d}_j$ and $\overline{d}_j$ as the prices where A_j and D_j take the values 0 and 1. This yields

$$\underline{a}_j = \frac{\pi_i}{m_i}, \quad \overline{a}_j = \frac{\pi_i}{m_i(1 - q_j)}$$

and

$$\underline{d}_j = \frac{\pi_i + m_i(r_i + r_j) + r_j m_j(1 - q_j)}{m_i + m_j(1 - q_j)}, \quad \overline{d}_j = \frac{\pi_i}{m_i} + r_i + r_j.$$

What remains to be done in order to pin down our equilibrium candidate is eliminating π_i, π_j and q_i and q_j using the system of equations (C.1)-(C.4). Inserting the expressions for $\underline{a}_j, \overline{a}_j, \underline{d}_j$ and $\overline{d}_j$ the system becomes

$$\frac{\pi_j}{m_j(1 - q_i)} = \frac{\pi_i}{m_i} + r_i \tag{C.7}$$

$$\frac{\pi_i}{m_i(1 - q_j)} = \frac{\pi_j}{m_j} + r_j \tag{C.8}$$

$$\frac{\pi_j}{m_j} + r_j = \frac{\pi_i + m_i(r_i + r_j) + r_j m_j(1 - q_j)}{m_i + m_j(1 - q_j)} \tag{C.9}$$

$$\frac{\pi_i}{m_i} + r_i = \frac{\pi_j + m_j(r_i + r_j) + r_i m_i(1 - q_i)}{m_j + m_i(1 - q_i)}. \tag{C.10}$$

Solving this system for π_i, π_j and q_i and q_j yields the equilibrium candidate given in the statement of the proposition.[9] We still need to check that this is well-defined and that it is indeed an equilibrium. It is easy to see that A_i and D_i are indeed distribution functions, i.e., that they are monotonically increasing. (Then it follows by construction that they have the correct supports.) In order to check that q_i is indeed a probability, note that $q_i(m_i, m_j)$ only depends on the ratio m_i/m_j (and not on r_i and r_j). Thus it is sufficient to show that the univariate function $q_i(m_i, 1)$ only takes values in the interval $[0, 1]$. This is omitted here. To see that the supports of the defensive and aggressive strategies are adjacent, note that putting (C.8) and (C.9) together immediately yields

$$\overline{a}_j = \frac{\pi_i}{m_i(1-q_j)} = \frac{\pi_i + m_i(r_i + r_j) + r_j m_j(1-q_j)}{m_i + m_j(1-q_j)} = \underline{d}_j.$$

By construction, firm j earns an expected payoff of π_j from playing a price in $[\underline{a}_j, \overline{d}_j]$. Thus in order to show that we have indeed a Nash equilibrium it remains to be shown that prices below $\underline{a}_i$ or above $\overline{d}_i$ are weakly dominated. Clearly, we can restrict attention to prices which are close enough to the supports of the equilibrium strategies to keep the two firms in competition: if firm i sets a price above $\overline{d}_j + r_i$ it yields zero profits because all customers attend firm j and likewise there is a lower bound below which lowering the price even further will never lead to additional customers. Thus consider firm i playing a price p (not too far) below $\underline{a}_i$ while firm j plays its equilibrium strategy. It is easy to see that this leads to a payoff of

$$\pi_i(p) = m_i(p - r_i) + m_j p(1 - q_j A_j(p + r_j)) \tag{C.11}$$

for firm i. Now observe that by multiplying (3.3) with m_j/m_i and shifting the argument p we can conclude that for some constant C_1 (which does not depend on p)

$$m_j(p + r_j)(1 - q_j A_j(p + r_j)) = C_1.$$

This allows us to rewrite (C.11) to

$$\pi_i(p) = C_1 + m_i(p - r_i) - r_j(1 - q_j A_j(p + r_j)).$$

Thus $\pi_i(p)$ is an increasing function which implies that playing $\underline{a}_i$ dominates playing prices below it. We now turn to deviations to prices above $\overline{d}_i$. Playing such a price

[9]There are three more solutions which do not correspond to equilibria. The reader who is inclined to verify that this is indeed a solution is strongly advised to utilize a computer algebra system such as Wolfram Mathematica.

yields a payoff of

$$\pi_i(p) = m_i(p - r_i)(1 - q_j)D_j(p - r_i)). \tag{C.12}$$

From (3.2) we can conclude that for some constant C_2

$$\frac{m_i^2}{m_j}p + m_i(p - r_i - r_j)(1 - q_j)(1 - D_j(p - r_i)) = C_2.$$

This yields

$$\pi_i(p) = C_2 - \frac{m_i^2}{m_j}p + m_i r_j(1 - q_j)(1 - D_j(p - r_i)).$$

Thus $\pi_i(p)$ is a decreasing function. This implies that playing $\overline{d}_i$ dominates playing higher prices.

To conclude the proof of the proposition, we have to show that $A_i(p) = D_j(p + r_j)$. Note that by construction both $A_i(p)$ and $D_j(p + r_j)$ are probability distributions on $[\underline{a}_i, \overline{a}_i]$. Furthermore, by (C.5), $D_j(p + r_j)$ is given by

$$D_j(p + r_j) = 1 - \frac{\pi_i - m_i(p - r_i)}{m_j p(1 - q_j)} \text{ for } p \in [\underline{a}_i, \overline{a}_i].$$

Observe that both $D_j(p + r_j)$ and $A_i(p)$ are of the following form:

$$G(p) = \alpha - \frac{\beta}{p} \text{ for } p \in [\underline{a}_i, \overline{a}_i],$$

$G(\underline{a}_i) = 0$ and $G(\overline{a}_i) = 1$ where α and β are coefficients that do not depend on p. Then the boundary constraints $G(\underline{a}_i) = 0$ and $G(\overline{a}_i) = 1$ uniquely determine the values of the coefficients α and β. Thus $D_j(p + r_j)$ and $A_i(p)$ must be identical. □

Proof of Proposition 3.3

Case (i) is an immediate corollary of Proposition 3.2. The transition value r^* is calculated as the value of r for which $\overline{d} = 1 + r$. Likewise, it is easy to verify that the pure strategy equilibrium of Case (iii) is indeed an equilibrium. We can thus focus on Case (ii). An equilibrium candidate is constructed in a similar way as in the proof of Proposition 3.2: we still assume the existence of an aggressive and a defensive strategy whose respective supports differ by a shift by r. But in addition we make the restriction that $\overline{d} = 1 + r$ and allow for an atom of size $q^0 = 1 - q^A - q^D$ in $1 + r$. Here, q^A and q^D denote the probabilities of attacking and defending.[10]

[10]For convenience we drop the indices i and j throughout the proof. The analogous system of equations with possibly asymmetric payoffs and probabilities has the same symmetric equilibrium as its only solution which is a Nash equilibrium.

Analogously to (3.2) and (3.3) we now get

$$\pi = 1 - q^A \tag{C.13}$$

for $p = 1 + r$,

$$\pi = (p - r)(1 - q^A A(p - r)) \tag{C.14}$$

for $p \in [\underline{d}, \overline{d})$,

$$\pi = p - r + p(1 - q^A - q^d D(p + r)) \tag{C.15}$$

and for $p \in [\underline{a}, \overline{a}]$. Solving (C.14) and (C.15) for A and D and using (C.13) to eliminate π we get

$$D(p) = \frac{1}{q^D}\left(1 - q^A - \frac{1 - q^A - p + 2r}{p - r}\right)$$

and

$$A(p) = \frac{1}{q^A}\left(1 - \frac{1 - q^A}{p}\right).$$

Calculating the values where these functions become 0 or 1 yields the boundaries

$$\overline{a} = 1, \qquad \underline{a} = 1 - q^A,$$
$$\overline{d} = \frac{r(1 - q^A - q^D) + 1 - q^A + 2r}{2 - q^A - q^D}, \qquad \underline{d} = \frac{1 - q^A + 3r - rq^A}{2 - q^A}.$$

Solving the system of equations $\overline{a} + r = \overline{d}$ and $\underline{a} + r = \underline{d}$ yields the equilibrium values of q^A, q^D and (through (C.13)) of π. It is straightforward to verify that these strategies are well-defined and that they interpolate between the strategies of Cases (i) and (iii). By construction, all prices in the support of the equilibrium strategy lead to the same payoff (given that the opponent plays his equilibrium strategy). Thus, to complete the proof it remains to be shown that prices outside the supports of A and D are dominated. Clearly, deviating to prices above $\overline{d}$ leads to zero demand and is thus dominated. Playing prices between $\overline{a}$ and $\underline{d}$ attracts the same customers as playing a price of $\underline{d}$ and is thus dominated. Likewise, deviating to a price slightly (i.e., less than $\underline{d} - \overline{a}$) below $\underline{a}$ is dominated since it does not attract more customers than playing a price of $\underline{a}$. That deviating to even lower prices is dominated can be seen with an argument parallel to the one in the proof of Proposition 3.2. Likewise, the same argument as in the proof of Proposition 3.2 can be applied to show that $D(p + r) = A(p)$. □

Proof of Lemma 3.1

First we prove that no firm charges negative prices in equilibrium.

Step (i). When only one firm charges possibly negative prices, this firm yields a loss when it plays such a negative price since at least the customer which get no rebate from the other firm buy from this firm. This cannot be optimal since zero profits can always be guaranteed.

Step (ii). When two firms possibly charge negative prices, customers will buy for sure when at least one firm indeed charges a negative price. Hence, at least one firm will sell a positive amount with positive probability when it charges a negative price. This firm's expected profit from charging this price is therefore negative. This cannot be optimal since zero profits can always be guaranteed.

Next, we prove that a firm which offers a rebate charges prices well above zero. Suppose firm 1 offers a rebate $r_1 > 0$. From before we know that firm 2 will not charge negative prices. Hence, when firm 1 charges prices $\in [0, r_1)$ at least the customers in its home base buy from it. Firm 1's expected profit increases when it gets more likely that also other customers buy from it. Suppose that also the other customers buy with probability 1. Then

$$\pi_1(p_1) = (p_1 - r_1)m_1X(p_1 - r_1) + p_1m_2X(p_1).$$

We denote the total mass of customers by

$$m := m_0 + m_1 + m_2.$$

As X is non-increasing and $m_{\neg 1} = m - m_1$, we have

$$\begin{aligned}\pi_1(p_1) &\leq (p_1 - r_1)m_1X(p_1 - r_1) + p_1(m - m_1)X(p_1 - r_1) \\ &= ((p_1 - r_1)m_1 + p_1(m - m_1))\, X(p_1 - r_1)\end{aligned}$$

for all $p_1 \geq 0$. Hence, for all $p_1 \in (-\infty, r_1m_1/m)$ we must have $\pi_1(p_1) < 0$. These prices are clearly dominated. □

Proof of Proposition 3.5

Suppose firm 1 offers a rebate. From Lemma 3.1 we know that then $p_1 \geq r_1m_1/m$, where $m := m_0 + m_1 + m_2$.

Case 1: firm 2 offers no rebates. Firm 2 can set $p_2 \nearrow r_1m_1/m$. Then all customers who do not get a rebate from firm 1 will buy from firm 2, when they buy. When they buy, firm 2 yields a nontrivial profit. When they do not buy for this price, firm 2 can lower the price so that it sells a positive amount and yields a nontrivial positive profit.

Case 2: firm 2 offers a rebate. When firm 2 sets the price $p_2 \nearrow r_1 m_1/m + r_2$ it gets all customers in its home base, when they buy at all.

For both cases we have shown that there exists a lower bound on firm 2's expected profit which is well above zero. Call this lower bound $\underline{\pi_2}$. Next we have to prove that also firm 1 earns an expected profit well above zero. Since for p_2 near zero firm 2's expected profit is below $\underline{\pi_2}$, firm 2 must charge prices well above zero in equilibrium. This enables firm 1 to yield a nontrivial positive profit. The arguments correspond to the ones of Case 2 above. □

Proof of Lemma 3.2

First, note that when $\varepsilon_{x,p} > 1$ then the revenue $R(p) = pX(p)$ is decreasing in p. Hence, conditional on the customers from firm i's home base buying from firm j, firm j's profit from this customer segment is decreasing in the net price when the net price exceeds $\hat{p}$. Moreover, the probability that customers from firm i's home base buy from firm j is weakly decreasing in p_j for every price setting strategy of firm i. We next have to distinguish two cases. Suppose that firm j sets a price $\check{p} > \hat{p} + r_j$.

Case 1: the expected profit of firm j is positive for $\check{p}$. The price $\check{p}$ is dominated by the price $\hat{p} + r_j$ because then (i) the profit from selling to each customer segment is positive for $\check{p}$ and for $\hat{p} + r_j$, (ii) from the arguments before we know that setting $\hat{p} + r_j$ instead of $\check{p}$ leads to a weakly higher probability that customers buy and to higher revenues and profits, conditional that customers buy from firm j.

Case 2: the expected profit of firm j is non-positive for $\check{p}$. From Proposition 3.5 we know that there are prices so that that the firm yields a positive expected profit.

Hence, playing gross prices exceeding $\hat{p} + r_j$ is dominated. □

Proof of Proposition 3.6

Denote our game by G. Recall that we assumed monopoly payoffs and thus monopoly prices to be bounded. Denote by G' the modified game in which firms pricing strategies are restricted to lie in $[0, u_j]$ where $u_j = \hat{p} + \max\{r_i, r_j\}$. From Lemmas 3.1 and 3.2 we know that playing prices outside $[0, u_j]$ is strictly dominated in G. Thus any Nash equilibrium of G' is also a Nash equilibrium of G. Define the set S^* by

$$S^* = [0, u_1] \times [0, u_2] \setminus \{(s_1, s_2 | s_1 + r_1 = s_2 \text{ or } s_2 + r_2 = s_1)\}.$$

S^* lies dense in the set of actions $[0, u_1] \times [0, u_2]$. Furthermore, payoffs are bounded and continuous in S^*. Thus by Simon and Zame (1990, p. 864), there exists a tie-breaking rule in G' for which a Nash equilibrium exists. Now observe that tie-breaking occurs in any equilibrium with probability 0: suppose that tie-breaking occurs with positive probability. This can only be due to both firms setting atoms in a way that a tie occurs (i.e., at distance r_1 or r_2). By Proposition 3.5, the supports of both players' equilibrium strategies must be bounded away from 0. Hence at least one firm has an incentive to slightly shift its atom downwards. Thus we can conclude that G' has a Nash equilibrium for any tie-breaking rule. This Nash equilibrium is also a Nash equilibrium of G. □

Proof of Proposition 3.7

We first show that the prices $p_i = p^M + r_i$ and $p_j = p^M + r_j$ form a Nash equilibrium for sufficiently large r_i and r_j. Since these strategies imply that each firm earns monopoly profits from its market segment, a deviation can only be profitable if it attracts additional customers from the other firm's segment. Thus it is sufficient to consider deviations to prices $p \in [0, p^M]$. Suppose that r_i is sufficiently large so that $p^M - r_i < 0$. Then firm $i's$ profit from deviating to a price $p \in [0, p^M]$ can be bounded from above as follows:

$$m_i\ (p - r_i)X(p - r_i) + m_j\ pX(p) < m_i\ (p^M - r_i)X(p^M) + m_j\ p^M X(p^M),$$

since $X(p^M) \leq X(p - r_i)$ and since $pX(p) \leq p^M X(p^M)$. If r_i is sufficiently large the upper bound becomes negative so that deviations cannot be profitable.

So far we have shown that for sufficiently high rebates there exists an equilibrium where both firms earn monopoly profits in their market segment. Now we show that this has to be true in any equilibrium. From Lemma 3.1 we know that no firm will charge negative prices. From before we know that for sufficiently high rebates a firm yields a loss if it charges a price $p \in [0, p^M]$. Therefore, $p_1, p_2 > p^M$ in any equilibrium. Hence, by charging a price of $p^M + r_i$ firm i can guarantee a profit of at least $m_i p^M X(p^M)$. Therefore, in equilibrium the expected profit of firm i must be at least $m_i p^M X(p^M)$. This hold for both firms. Therefore, in an equilibrium the sum of both firms expected profits is at least $(m_1 + m_2)p^M X(p^M)$. By the definition of the monopoly profit the maximum sum of profits is $(m_1 + m_2)p^M X(p^M)$. All this is compatible only if firm 1 earns an expected profit of $m_1 p^M X(p^M)$ and firm 2 of $m_2 p^M X(p^M)$. That is, there can only be equilibria in which firms yield expected profits equal to the monopoly profits in their market segment.

Next, we prove the final part of the proposition which considers the case where $m_0, m_1, m_2 > 0$ and where a choke price exists. The proof is similar as the part before and is therefore only sketched. First, when rebates are sufficiently high a firm yields a loss when it charges a price below the choke price. Second, therefore in equilibrium the prices are above the choke price. Third, this implies that in equilibrium customers without rebate opportunities do not buy. Fourth, therefore customers without rebate opportunities can be ignored and the proof for the case $m_0 = 0$ applies. □

D Appendix to Chapter 4

D.1 Concavity Results for First and Second Order Statistics

In this appendix we develop a number of conditions that guarantee the strict concavity of sequences of expected first and second order statistics and related expressions.

Throughout denote by $X_{k:n}$ the k^{th} order statistic, i.e., the k^{th} largest of n non-negative, independent, identically distributed random variables X_i with distribution function F. Unless otherwise noted, we assume that $E[X_i] < \infty$. Due to

$$E[X_{k:n}] < nE[X_1]$$

the last condition immediately implies that $E[X_{k:n}] < \infty$ for all k. Assume throughout this section that F has a density f. For notational convenience, we assume additionally that F has full support on $[0, \infty)$. Yet the generalization to distributions with an interval support poses no difficulties.

First Order Statistics

Our first results deal with first order statistics (and hence with welfare maximization):

Proposition D.1. *(i) The sequence $E[X_{1:n}]$ is strictly concave in n.*

(ii) For any c with $F(c) \in (0, 1)$ the sequence

$$E[\max(X_{1:n}, c)]$$

is strictly concave in n.

(iii) For any c with $F(c) \in (0, 1)$ the sequence

$$E[X_{1:n} 1_{\{X_{1:n} > c\}}]$$

is strictly concave in n.

Part (i) is an elementary result which is proved mainly for expository reasons here. Note that Part (ii) of the proposition is an immediate consequence of Part (i) since by

$$E[\max(X_{1:n}, c)] = E[\max_i \max(X_i, c)]$$

the expressions $E[\max(X_{1:n}, c)]$ are first order statistics of the distribution G where G differs from F in concentrating the mass F puts on $[0, c]$ in an atom on c. Likewise, Part (iii) is an immediate consequence of Part (i) since the sequence

$$E[X_{1:n} 1_{\{X_{1:n} > c\}}]$$

is a sequence of first order statistics from a distribution which differs from F in concentrating all mass that F puts on $[0, c]$ on an atom in 0. Hence we found that sequences of expected first order statistics are easy to handle: Given that the underlying distribution has finite mean, the sequence is concave.

IFTR, UIFTR and Second Order Statistics

For second order statistics matters are more complex. We show that the IFTR condition we introduced and discussed in Section 4.2.1 is a mild but sharp condition for concavity of second order statistics. Recall that F is IFTR (Increasing Failure-out-of-Two Rate) iff the rate

$$\frac{f(x)}{(1-F(x))^2}$$

is increasing in x. If in addition this rate grows unboundedly we say that F is UIFTR. The UIFTR property is needed to guarantee that in addition to concavity the increments of the sequence of second order statistics go to zero.

To put the IFTR property into perspective, let us briefly recall a well-known result from reliability theory, see e.g. Barlow and Proschan (1981): Consider the expected difference between the first two order statistics $E[X_{1:n} - X_{2:n}]$. For a distribution F with positive density f, it holds that

$$E[X_{1:n} - X_{2:n}] = E\left[\frac{1-F(X_{1:n})}{f(X_{1:n})}\right]. \tag{D.1}$$

As the sequence $X_{1:n}$ is stochastically increasing, the monotonicity behavior of the failure rate $f/(1-F)$ is crucial: If the failure rate is increasing (i.e., if F is IFR), $E[X_{1:n} - X_{2:n}]$ is decreasing. Conversely, if F is DFR, $E[X_{1:n} - X_{2:n}]$ is increasing. Our main interest is in the monotonicity behavior of $E[X_{2:n+1} - X_{2:n}]$ and our central observation is that a similar approach is possible. As noted above, it holds that

$$E[X_{2:n+1} - X_{2:n}] = E\left[\frac{(1-F(X_{1:n}))^2}{f(X_{1:n})}\right].$$

From this it is easy to see that assuming IFTR guarantees the concavity of the sequence $E[X_{2:n}]$. Note finally that since $1/(1-F)$ is increasing, IFTR is considerably weaker than IFR.

In particular, all distributions which are IFR (such as the exponential and the uniform distribution) are also IFTR and UIFTR. Moreover, UIFTR is strictly stronger than IFTR, but strictly weaker than IFR: The distribution

$$F(x) = 1 - e^{2-2\sqrt{1+x}}$$

is DFR but satisfies UIFTR (and hence also IFTR). The distribution

$$F(x) = \frac{x - \log(1+x)}{1 + x - \log(1+x)}$$

is DFR and satisfies IFTR, but not UIFTR as

$$\frac{f(x)}{(1-F(x))^2} = 1 - \frac{1}{1+x}$$

is increasing, but bounded.

Proposition D.2. *If F is IFTR, the sequence $E[X_{2:n}]$ is strictly concave in $n \geq 2$. If furthermore F is UIFTR, then $E[X_{2:n+1}] - E[X_{2:n}]$ converges to zero as n becomes large.*

Conversely, if $f/(1-F)^2$ is (weakly) decreasing, the sequence $E[X_{2:n}]$ is (weakly) convex in n. For example, for the distribution $F(x) = 1 - 1/(1+x)$, for which $f/(1-F)^2 = 1$, it holds that $E[X_{2:n}] = n - 1$. For the distributions $F(x) = 1 - 1/(1+x)^\rho$ with $\rho \in (\frac{1}{2}, 1)$, the sequence of second order statistics is strictly convex (and $E[X_{2:n}] < \infty$ for all n). At first sight, this seems to contradict the result that the first order statistics form a strictly concave sequence, since $E[X_{1:n}]$ and $E[X_{2:n}]$ converge to the same limit as $n \to \infty$ and as $E[X_{2:n}] < E[X_{1:n}]$: Would the second order statistics not have to get ahead of the first order statistics from some point on? The solution to this apparent contradiction is that in these examples $E[X_{1:n}] = \infty$ for all n: If $f/(1-F)^2$ is weakly increasing, the second order statistics are weakly convex, implying that from some n on $E[X_{1:n}]$ must be infinite. From this we can conclude, however, that $E[X_1] = \infty$ (otherwise all order statistics would be finite). Thus $E[X_{1:n}] = \infty$ for all n since $E[X_{1:n}] > E[X_1]$ for all n. This demonstrates the well-known but puzzling fact that for some distributions $E[X_{2:n}] < \infty$ for all $n > 1$ while $E[X] = \infty$.

Next we give a result along the lines of the later parts of Proposition D.1 for second order statistics. It turns out that we generally do not get global concavity under IFTR. Yet the sequences are concave from a threshold on.

Proposition D.3. *(i) If F is IFTR, for all c with $c > 0$ there is a finite $n_0 \in \mathbb{N}$ such that the sequence*

$$q_n = E[\max(c, X_{2:n})]$$

is strictly convex in $n \geq 2$ until n_0 and strictly concave from then on.

(ii) If F is IFTR, for any $c > 0$ there is a finite $n_0 \in \mathbb{N}$ such that the sequence $q_n = E[X_{2:n} 1_{\{X_{2:n} > c\}}]$ is strictly concave in n for $n \geq n_0$.

(iii) If F is IFTR, for any $c > 0$ there is a finite $n_0 \in \mathbb{N}$ such that the sequence $q_n = E[c\, 1_{\{X_{1:k} \geq c \wedge X_{2:k} \leq c\}} + X_{2:k}\, 1_{\{X_{2:k} > c\}}]$ is strictly concave in n for $n \geq n_0$.

(iv) For the sequences q_k considered in (i), (ii) and (iii), the increments $q_{n+1} - q_n$ converge to zero as n goes to infinity provided that F is UIFTR.

Of course, the constant n_0 depends both on F and on c. Since the sequence $E[\max(c, X_{2:n})]$ is first convex and then concave, n_0 is the point where the discrete second derivative of the sequence becomes negative. From the proof one sees that the discrete second derivative is given by

$$\begin{aligned} & E[(\max(c, X_{2:n+1}) - \max(c, X_{2:n})) - (\max(c, X_{2:n}) - \max(c, X_{2:n-1}))] \\ = & \int_c^\infty n \left(F(x) - \frac{n-1}{n} \right) (1 - F(x))^2 F(x)^{n-2} dx \end{aligned} \tag{D.2}$$

which is straightforward to calculate for concrete choices of F, c and n. A similar criterion can be given for the latter part of the proposition.[11] For many natural distributions and moderate values of c, the sequences are concave from the start or from a moderate value n_0 on. Under $c = E[X]$, for the exponential distribution, the smallest n where (D.2) becomes negative is $n = 3$, for the uniform distribution, it is $n = 2$.

The proof of Proposition D.3 relies on the following lemma:

Lemma D.1. *Let $X_{n:n}$ be the lowest order statistic of the independent random variables X_i with distribution F, density f and support $[0, \infty)$. Then for any $c > 0$ the sequence*

$$E[X_{n:n} 1_{\{X_{n:n} < c\}}]$$

is either decreasing, or increasing until some n_0 and decreasing onwards.

For our auction where bidders hold preliminary information, we also need the following more specialized result:

Lemma D.2. *Let $X_1, \ldots, X_m$, $Y_1, \ldots, Y_m$ be independent exponentially distributed random variables. Then the sequence*

$$p_n = E[\max(X_1 + Y_1, \ldots, X_n + Y_n, X_{n+1} + 1, \ldots, X_m + 1)]$$

is strictly concave in n (for $1 \leq n \leq m$).

Dispersion

In order to study how allocations vary when moving from one distribution function to another, more dispersed one we rely on results from the theory of stochastic

[11] In the proof, the sequence is decomposed into a sum of two sequences which are first convex and then concave. If both sequences are concave for $n > n_0$, their sum is also concave for $n > n_0$.

orders. Recall the definition of the excess wealth order discussed already in Section 4.2.3: The distribution function F is more dispersed in the excess wealth order than the distribution function G if for all $p \in (0,1)$

$$\int_{F^{-1}(p)}^{\infty} 1 - F(x)dx \geq \int_{G^{-1}(p)}^{\infty} 1 - G(x)dx.$$

For this we write $F \geq_{\text{EW}} G$. The excess wealth order is also known as right spread order. It is weaker than the well-known dispersive order, i.e., $F \geq_{\text{disp}} G$ implies $F \geq_{\text{EW}} G$.[12]

The following result is very much based on Proposition 3.4 of Li and Shaked (2004). We only transfer it from spacings $E[X_{j+1:k} - X_{j:k}]$ of order statistics to the increments $E[X_{j:k+1} - X_{j:k}]$ that interest us. Furthermore, we add the observation that $F_{2:2} \geq_{\text{EW}} G_{2:2}$ implies $F \geq_{\text{EW}} G$ where as before $F_{2:2}$ denotes the distribution of the minimum of two F-distributed random variables.

Proposition D.4. *Consider two distribution functions F and G with densities f and g on $\mathbb{R}_+$. Denote by (X_i) and (Y_i) collections of independent random variables distributed, respectively, according to F and G.*

(i) Assume $F \geq_{EW} G$. Then for all $k \geq 1$

$$E[X_{1:k+1} - X_{1:k}] \geq E[Y_{1:k+1} - Y_{1:k}].$$

(ii) Assume $F_{2:2} \geq_{EW} G_{2:2}$. Then for all $k \geq 2$

$$E[X_{2:k+1} - X_{2:k}] \geq E[Y_{2:k+1} - Y_{2:k}].$$

(iii) $F_{2:2} \geq_{EW} G_{2:2}$ implies $F \geq_{EW} G$.

It is also easy to show that $F \geq_{\text{disp}} G$ implies $F_{2:2} \geq_{\text{EW}} G_{2:2}$. One word on the related literature seems in order. Paul and Gutierrez (2004) claim results based on the star-order that would be well-suited for our purposes if they were correct: As pointed out by Xu and Li (2008), the relevant Theorems 3 and 4 of Paul and Gutierrez (2004), reproduced in Theorem 4.B.19 of Shaked and Shantikumar (2007) are incorrect except for the case of first order statistics. Yet for first order statistics, the results are weaker than the result for the excess wealth order from Shaked and Li (2004) we apply here. A weaker alternative is Theorem 3.B.31 of Shaked and Shantikumar (2007) which relies on the dispersive order.

[12]See Shaked and Shantikumar (2007), p.166.

D.2 The Decision Problem

In this section we analyze the general decision problem which covers all the situations of welfare and revenue maximization in the models we discussed. Denote by the random variable K the number of bidders who are active in the respective auction and know their valuation realization. Denote by p_k the decision-maker's revenue conditional on K taking the value k. Note that this revenue either stands for welfare or for the seller's profits in the auction models. We call $(p_k)_k$ the payoff sequence. The decision-maker controls the probabilities γ_i. Setting a vector γ leads to costs of $c(\gamma)$ for the decision-maker where c is an increasing cost function.

The general optimization problem is hence given by

$$\max_{\gamma \in [0,1]^n} \pi(\gamma) - c(\gamma)$$

where $\pi(\gamma)$ is the expected value of p_K given probabilities γ,

$$\pi(\gamma) = E[p_K] = \sum_{j=0}^{n} \alpha_k(\gamma) p_k,$$

and $\alpha_k(\gamma)$ is the probability that K takes the value k,

$$\alpha_k(\gamma) = Prob[K = k] = \sum_{A \subseteq \{1,\dots,n\}, |A|=k} \prod_{j \in A} \gamma_j \prod_{l \in A^c} (1 - \gamma_l).$$

Observe that, with an increasing payoff sequence p_k, π is increasing in all γ_i.

Globally Concave Payoff Sequences

Let us for now assume that c is a quasi-concave, increasing, symmetric, continuous function and that the payoff sequence $(p_k)_k$ is strictly concave. The latter assumption is relaxed in Section D.2.

Let us first deal with the expected revenue π and focus on the optimal choice of γ_i and γ_j, while keeping the other γ_l fixed. We show that under the constraint $\gamma_i + \gamma_j = \delta$ the expected gross revenue is maximized by the most asymmetric allocation of probabilities:

Proposition D.5. *Assume that the sequence* $(p_k)_k$ *is strictly concave. Fix* γ_3,..., $\gamma_n \in [0, 1]$ *and* $\delta \in (0, 2)$. *Then the maximization problem*

$$\max_{\gamma_1,\gamma_2\in[0,1]} \pi(\gamma_1, \dots, \gamma_n) \text{ such that } \gamma_1 + \gamma_2 = \delta$$

is strictly solved by

$$\{\gamma_1^*, \gamma_2^*\} = \{\lfloor\delta\rfloor, \delta - \lfloor\delta\rfloor\}.$$

$\lfloor\cdot\rfloor$ denotes the Gaussian floor function, which maps δ to the largest integer weakly smaller than δ, hence either to 0 or to 1 in our case. Thus the decision-maker splits $\delta > 1$ such that one γ_i is set to 1 while he splits $\delta < 1$ such that one γ_i is set to 0. The intuition behind is as follows: Due to the concavity of the sequence $(p_k)_k$ the decision-maker prefers earning p_k for sure over earning p_{k-1} or p_{k+1} with equal probability. More accurately, keeping the sum of the γ_i fixed is equivalent to keeping $E[K]$ fixed. By Jensen's inequality,

$$E[p_K] \leq p_{E[K]}.$$

The inequality gets tighter when the randomness in K is reduced. Thus the decision-maker optimally chooses the γ_i such that the randomness is minimized. This same intuition also underlies our following results. We show that the decision-maker optimally sets at most one γ_i to an intermediate value inside $(0, 1)$.

Obviously, we can separate the decision-maker's problem of maximizing net revenue into first choosing the optimal allocation for each iso-cost set and then choosing the optimal iso-cost set. As we are mostly interested in the structure of the optimal allocation we will most of the time focus on the first of these two problems: We analyze how the decision-maker optimally sets the γ_i given that he wants to spend a fixed level of total costs.

Let us first introduce some more definitions: Define as H_s the feasible allocations for which the γ_i sum up to s, i.e.,

$$H_s = \left\{ \gamma \in [0, 1]^n \,\middle|\, \sum_{i=1}^{n} \gamma_i = s \right\}.$$

Note that H_s can also be characterized as the set of allocations γ which lead to $E[K] = s$. Define as K_s those allocations in H_s where all but one γ_i are either zero or one,

$$K_s = \{\gamma \in H_s \,|\gamma_i \in (0, 1) \text{ and } \gamma_j \in (0, 1) \Rightarrow\ i = j\}.$$

We call K_s the **maximally asymmetric allocations in** H_s. Observe that for $\gamma \in K_s$ the gross revenue π takes the following simple form:

$$\pi(\gamma) = (s - \lfloor s \rfloor)\, p_{\lfloor s \rfloor + 1} + (1 - (s - \lfloor s \rfloor))p_{\lfloor s \rfloor} \qquad \text{for} \qquad \gamma \in K_s. \tag{D.3}$$

From Proposition D.5 we obtain the following corollary:

Corollary D.1. *Assume that the sequence $(p_k)_k$ is strictly concave. The maximally asymmetric allocations K_s maximize gross revenue within H_s, i.e., for $s \in (0, n)$*

$$\gamma^* \in \arg\max_{\gamma \in H_s} \pi(\gamma) \quad \Leftrightarrow \quad \gamma^* \in K_s.$$

Our next result shows that a similar characterization of maximizers is possible for the net revenue under quasi-concave costs. Define, for an increasing cost function c, $\underline{c}$ as the minimal and $\overline{c}$ as the maximal costs arising from choosing a feasible allocation, i.e.,

$$\underline{c} = c(0, \ldots, 0) \qquad \text{and} \qquad \overline{c} = c(1, \ldots, 1).$$

Define also the iso-cost curve c_κ corresponding to a cost level of κ:

$$c_\kappa = \{\gamma \in [0, 1]^n | c(\gamma) = \kappa\}.$$

We obtain the following result:

Proposition D.6. *Assume that the payoff sequence $(p_k)_k$ is strictly concave. Assume that the cost function $c \,:\, [0, 1]^n \to \mathbb{R}^+$ is quasi-concave, strictly increasing, continuous and symmetric. Then for any $\kappa \in [\underline{c}, \overline{c}]$ there exists an $s \in [0, n]$ such that*

$$\gamma^* \in \arg\max_{\gamma \in c_\kappa} \pi(\gamma) - c(\gamma) \quad \Leftrightarrow \quad \gamma^* \in K_s.$$

The assumption of quasi-concavity is fulfilled, e.g., by cost functions of the form

$$c(\gamma) = C\left(\sum_{i=1}^{n} \gamma_i\right)$$

where $C \,:\, \mathbb{R}^+ \to \mathbb{R}^+$ is an increasing, continuous function.

Next we provide a simple criterion for comparing optimal allocations between two payoff sequences. For this we impose a more rigid assumption on the cost function, i.e. that it is a strictly convex transformation of a linear cost function. We presume the existence of an interior solution, i.e. that the decision-maker optimally gives out $\gamma \in K_s$ for some $s \in (0, n)$. An interior solution comes out naturally when costs are neither prohibitively high nor so low that the decision-maker optimally gives out all information. From (D.3) it is easy to see that a sufficient condition for an interior solution is given by $C'(0) < p_1 - p_0$ and $C'(n) > p_n - p_{n-1}$.

Proposition D.7. *Consider two payoff sequences $(p_k)_k$ and $(q_k)_k$. Assume that the costs are given by*

$$c(\gamma) = C\left(\sum_{i=1}^{n} \gamma_i\right)$$

for a strictly increasing, strictly convex and continuously differentiable function C with

$$C'(0) < \min(p_1 - p_0, q_1 - q_0) \text{ and } C'(n) > \max(p_n - p_{n-1}, q_n - q_{n-1}).$$

Denote by π_p and π_q the decision-maker's respective gross revenue under the sequences $(p_k)_k$ and $(q_k)_k$. Then the decision-maker's problem of maximizing, respectively, $\pi_p(\gamma) - c(\gamma)$ and $\pi_q - c(\gamma)$ is solved by $\gamma_p \in K_{s_p}$ and $\gamma_q \in K_{s_q}$ which are unique up to a relabeling of bidders with $s_p, s_q \in (0, n)$. Moreover, if $(p_k - q_k)_k$ is strictly increasing in k, then $s_p > s_q$.

Thus the proposition provides a clear condition under which one payoff sequence leads to more release of advertising or information than another. Note that $(p_k - q_k)_k$ being increasing in k is equivalent to $(p_k)_k$ having larger increments than $(q_k)_k$ since

$$p_k - q_k \geq p_{k-1} - q_{k-1}$$

is equivalent to

$$p_k - p_{k-1} \geq q_k - q_{k-1}.$$

Concavity from a Threshold on

So far we assumed that the payoff sequence (p_k) is strictly concave. Yet we also need results for cases where the sequence (p_k) is concave only from some threshold k^* on. We hence introduce the following assumption:

Assumption D.1. *Assume (p_k) is weakly increasing and there is a k^* such that the shifted sequence $\hat{p}_j = p_{k^*+j}$ is strictly concave for $j \geq 0$.*

Our first observation is useful but mathematically rather trivial: Assume that $\gamma_1, ..., \gamma_l$ are fixed to take the value 1. The decision-maker only decides how to optimally set $\gamma_{l+1}, ..., \gamma_n$. This corresponds to a shift in the sequence (p_k) which may remove the non-concavity:

Corollary D.2. *Under Assumption D.1, assume that $\gamma_1 = \cdots = \gamma_l = 1$ for some $l \geq k^*$ and that the seller sets $\hat{\gamma} = (\gamma_{l+1}, \ldots, \gamma_n)$ according to an increasing, continuous, symmetric, and quasi-concave cost-function $c(\hat{\gamma})$. Then for any $\kappa \in [\underline{c}, \overline{c}]$ there exists an $s \in [0, n-l]$ such that*

$$\hat{\gamma^*} \in \arg\max_{\hat{\gamma} \in c_\kappa} \pi(\hat{\gamma}) - c(\hat{\gamma}) \quad \Leftrightarrow \quad \hat{\gamma^*} \in K_s$$

where the definitions of K_s and c_κ are modified accordingly to account only for bidders $l+1$ to n.

Our next results show that, quite intuitively, under concavity from a threshold on, sufficiently large sums of efforts s should be allocated maximally asymmetrically. Corollary D.3 provides an explicit criterion for when an amount s is sufficiently large.

Corollary D.3. *Under Assumption D.1, assume there exists a weakly concave sequence $(q_k)_k$ which weakly dominates $(p_k)_k$ and which coincides with $(p_k)_k$ for $k \geq l^*$ for some $n > l^* > 0$. Denote by π_p and π_q the expected gross payoff with respect to the sequences $(p_k)_k$ and $(q_k)_k$. Denote by $\kappa^* \in (\underline{c}, \overline{c})$ the smallest level of costs for which the allocation involves setting a sum of efforts greater than l^*, i.e.,*

$$\kappa^* = \inf \left\{ \kappa \in [\underline{c}, \overline{c}] \,\middle|\, \left[\arg \max_{\gamma \in c_\kappa} \pi_q(\gamma) - c(\gamma)\right] \in K_s \text{ for some } s \geq l^* \right\}.$$

Then for any $\kappa \in [\kappa^*, \overline{c}]$, there exists an $s \in [l^*, n]$ such that

$$\gamma^* \in \arg\max_{\gamma \in c_\kappa} \pi(\gamma) - c(\gamma) \quad \Leftrightarrow \quad \gamma^* \in K_s.$$

Note that from Proposition D.6 and from the monotonicity of c it follows that κ^* is a well-defined cost-level inside the interval $(\underline{c}, \overline{c})$. Lemma D.3 shows that the criterion from Corollary D.3 for an effort sum s to be sufficiently large is widely applicable: A sequence $(q_k)_k$ as postulated in the corollary is constructible provided a mild additional condition holds and provided n is sufficiently large.

Lemma D.3. *Let $(p_k)_k$ be an (infinite) sequence in $\mathbb{R}$ which is weakly increasing and for which there is a k^* such that the shifted sequence $\hat{p}_j = p_{k^*+j}$ is strictly concave for $j \geq 0$. Assume in addition that the increments of $(p_k)_k$ go to zero as k becomes large, i.e.,*

$$\lim_{k\to\infty} p_{k+1} - p_k = 0.$$

Then there exists a number l^* and a weakly concave sequence $(q_k)_k$ for which $q_k \geq p_k$ for all $k \geq 0$ and for which $q_j = p_j$ for all $j \geq l^*$.

D.3 Proofs

Proof of Proposition 4.1

By Proposition D.2 in Appendix D.1, the sequence $p_1, p_2, \ldots$ is strictly concave under IFTR. Thus by Corollary D.2 in Appendix D.2 maximally asymmetric allocations are optimal if at least one bidder is present anyway. This proves (i). Under UIFTR, by Proposition D.2 the increments of (p_k) converge to zero. Thus by Lemma D.3 it is possible to construct a concave sequence which is weakly greater than (p_k) and coincides with (p_k) from some point on. Thus we can apply Corollary D.3 and conclude that sufficiently large amounts of advertising should be allocated maximally asymmetrically. This proves (ii). □

Proof of Proposition 4.2

By Proposition D.1, the sequence $q_0, q_1, \ldots$ is strictly concave. Thus by Proposition D.6 maximally asymmetric allocations are optimal. □

Proof of Proposition 4.3

Recall from (D.1) that $E[v_{1:k} - v_{2:k}]$ is monotonically decreasing under IFR and monotonically increasing under DFR. Thus one can apply Proposition D.7 and conclude that under IFR the welfare-optimal allocation under $q_k = E[v_{1:k}]$ involves less informational effort than the revenue-optimal allocation under $p_k = E[v_{2:k}]$. Under DFR, the opposite is true. □

Proof of Proposition 4.4

By Proposition D.4 the increments of both (p_k) and (q_k) are larger for F than for G under the given assumptions on F and G. Thus Proposition D.7 (applied separately to revenue- and to welfare-maximization) shows that a higher advertising level is chosen under F than under G. □

Proof of Proposition 4.5

By Proposition D.3 (iii) in Appendix D.1, the sequence (p_k) is strictly concave from some k^* on under IFTR. Thus by Corollary D.2 in Appendix D.2 maximally asymmetric allocations are optimal if at least k^* bidders are present. Under UIFTR, by Proposition D.3 (iv) the increments of (p_k) converge to zero. Thus by Lemma D.3 it is possible to construct a concave sequence which is weakly greater than (p_k) and which coincides with (p_k) from some point on. Hence we can apply Corollary D.3 and conclude that sufficiently large amounts of advertising should be allocated maximally asymmetrically. □

Proof of Proposition 4.6

The proof is the same as that of Proposition 4.5, except that in the first step we apply Proposition D.3 (i) instead of D.3 (iii). □

Proof of Proposition D.1

Recall[13] that $X_{1:n}$ has distribution function $F_{1:n}(x) = F(x)^n$ and that we can thus write

$$E[X_{1:n}] = \int_0^\infty 1 - F(x)^n dx. \tag{D.4}$$

In order to prove Part (i), we have to show that the sequence $d(n)$ defined by

$$d(n) = E[X_{1:n+1}] - E[X_{1:n}]$$

is strictly decreasing in n. By (D.4) we can conclude that

$$d(n) = \int_0^\infty F(x)^n(1 - F(x))dx.$$

Since the integrand is strictly decreasing in n over the support of F (and constant outside the support) and since we assumed that F is non-degenerate, $d(n)$ is strictly decreasing in n. This proves Part (i). How to conclude Parts (ii) and (iii) from Part (i) is indicated in the text below the statement of the proposition. □

Proof of Proposition D.2

Note that $X_{2:n}$ has distribution function

$$F_{2:n}(x) = F(x)^{n-1}(F(x) + n(1 - F(x))) = nF(x)^{n-1} - (n-1)F(x)^n.$$

[13]The formulae for the distribution functions and densities of order statistics used in this and the following proofs can be found, e.g., in David (1970).

Analogously to the proof of Proposition D.1, we write

$$E[X_{2:n}] = \int_0^\infty 1 - F_{2:n}(x)dx = \int_0^\infty 1 - nF(x)^{n-1} + (n-1)F(x)^n dx$$

and define

$$\begin{aligned} d(n) &= E[X_{2:n+1} - X_{2:n}] = \int_0^\infty nF(x)^{n-1} - 2nF(x)^n + nF(x)^{n+1}dx \\ &= \int_0^\infty nF(x)^{n-1}(1-F(x))^2 dx. \end{aligned}$$

To show that $E[X_{2:n}]$ is strictly concave, we have to prove that $d(n)$ is strictly decreasing or, equivalently, that $dd(n)$ given by

$$dd(n) = d(n) - d(n-1) = \int_0^\infty n(1-F(x))^2 F(x)^{n-2}\left(F(x) - \frac{n-1}{n}\right)dx$$

is negative.

To make use of our assumption that $f/(1-F)^2$ is strictly increasing we first need to make a number of observations: Recall that $X_{1:n}$ has density $nF(x)^{n-1}f(x)$ and that since this is a probability density it holds that

$$1 = \int_0^\infty nF(x)^{n-1}f(x)dx.$$

Subtracting from this identity the same identity with $n-1$ instead of n one obtains

$$\begin{aligned} 0 &= \int_0^\infty nF(x)^{n-1}f(x) - (n-1)F(x)^{n-2}f(x)dx \\ &= \int_0^\infty nF(x)^{n-2}\left(F(x) - \frac{n-1}{n}\right)f(x)dx. \end{aligned} \tag{D.5}$$

Now note that the integrand in our expression for $dd(n)$ is negative for $x < F^{-1}(\frac{n-1}{n})$ and positive for $x > F^{-1}(\frac{n-1}{n})$. Thus, if for some strictly increasing, positive function $h(x)$ it holds that

$$\int_0^\infty nF(x)^{n-2}\left(F(x) - \frac{n-1}{n}\right)(1-F(x))^2 h(x)dx \leq 0, \tag{D.6}$$

this implies $dd(n) < 0$. Now set $h(x) = f(x)/(1-F(x))^2$ and observe that with this choice of h the left hand side of (D.6) equals the right hand side of (D.5) and thus equals 0. Hence we have shown that, under our assumptions, $dd(n) < 0$. That $d(n)$ converges to zero in n provided that $f(x)/(1-F(x))^2$ goes to infinity as x gets large follows as the special case $c = 0$ from Part (iv) of Proposition D.3. □

Proof of Proposition D.3

(i) Note first that

$$\begin{aligned} E[\max(c, X_{2:n})] &= c + E[\max(0, X_{2:n} - c)] \\ &= c + \int_c^\infty (x-c) f_{2:n}(x)dx = c + \int_c^\infty 1 - F_{2:n}(x)dx, \end{aligned}$$

where $f_{2:n}$ and $F_{2:n}$ denote the density and the distribution function of $X_{2:n}$. The last equality holds as for any distribution function G with density g on $[0, \infty)$, for all $c \geq 0$:

$$\int_c^\infty (x-c)g(x)dx = \int_c^\infty 1 - G(x)dx. \tag{D.7}$$

For $c = 0$ this identity is well-known. In order to verify it for $c > 0$, note that both sides of the equation have the same (bounded) derivative $G(c) - 1$ with respect to c.

Thus it holds that (completely analogous to the corresponding quantities in the proof of Proposition D.2)

$$d(n) = E[\max(c, X_{2:n+1})] - E[\max(c, X_{2:n})] = \int_c^\infty nF(x)^{n-1}(1 - F(x))^2 dx$$

and

$$dd(n) = d(n) - d(n-1) = \int_c^\infty n(1-F(x))^2 F(x)^{n-2} \left(F(x) - \frac{n-1}{n} \right) dx.$$

To complete the proof we have to show that, under the assumptions made, $d(n)$ is either decreasing for all n or that there is an n_0 such that $d(n)$ is increasing until n_0 and decreasing from there on.

Define $h(x) = \frac{(1-F(x))^2}{f(x)}$. Note that by the definition of $f_{1:n}$ one can rewrite $d(n)$ to

$$d(n) = \int_c^\infty h(x) f_{1:n}(x)dx = E\left[h(X_{1:n})1_{\{X_{1:n} > c\}}\right].$$

Now define the random variables $Y_i = h(X_i)$ and the constant $a = h(c)$. Since h is strictly decreasing by assumption, one can conclude that $Y_{n:n} = h(X_{1:n})$. Thus

$$d(n) = E\left[Y_{n:n}1_{\{Y_{n:n} < a\}}\right].$$

Now apply Lemma D.1 to $d(n)$ and conclude that $d(n)$ is either decreasing, or it switches once from being increasing to being decreasing.

(ii) We show that the increments $\widetilde{d}$ of the sequence $E\left[X_{2:n}1_{\{X_{2:n}>c\}}\right]$ can be written as the sum of the sequence d from (i) and an easy to analyze sequence as follows:

$$\begin{aligned} E\left[X_{2:n}1_{\{X_{2:n}>c\}}\right] &= E\left[(X_{2:n}-c)1_{\{X_{2:n}>c\}}\right] + cE\left[1_{\{X_{2:n}>c\}}\right] \\ &= E\left[(X_{2:n}-c)1_{\{X_{2:n}>c\}}\right] + c\,(1-F_{2:n}(c)). \end{aligned} \tag{D.8}$$

Hence by the proof of (i) and by the definition of $F_{2:n}$,

$$\widetilde{d}(n) = E\left[X_{2:n+1}1_{\{X_{2:n+1}>c\}}\right] - E\left[X_{2:n}1_{\{X_{2:n}>c\}}\right] = d(n) + c\,(F_{2:n}(c) - F_{2:n+1}(c)).$$

Since $d(n)$ is decreasing in n for large enough n, it is sufficient to show that $(F_{2:n}(c) - F_{2:n+1}(c))$ is decreasing for sufficiently large n as well. Indeed,

$$F_{2:n}(c) - F_{2:n+1}(c) = nF(c)^{n-1}(1-F(c))^2$$

is decreasing in n for sufficiently large n. Hence we are done.

(iii) Note that with (D.8) one can rewrite our sequence of interest as follows:

$$\begin{aligned} & E[X_{2:n}\,1_{\{X_{2:n}>c\}} + c\,1_{\{X_{1:n}\geq c\,\wedge\, X_{2:n}\leq c\}}] \\ = \; & E\left[(X_{2:n}-c)1_{\{X_{2:n}>c\}}\right] + cE\left[1_{\{X_{2:n}>c\}}\right] + cE\left[1_{\{X_{1:n}\geq c\,\wedge\, X_{2:n}\leq c\}}\right] \\ = \; & E\left[(X_{2:n}-c)1_{\{X_{2:n}>c\}}\right] + cE\left[1_{\{X_{1:n}\geq c\}}\right]. \end{aligned} \tag{D.9}$$

That the first summand becomes concave for sufficiently large n was shown in (i). Thus we only have to consider the second summand

$$cE\left[1_{\{X_{1:n}\geq c\}}\right] = c(1-F(c)^n).$$

As this is a strictly concave sequence, we are done.

(iv) Consider the sequence $d(n)$ from Part (i) of the proof:

$$d(n) = E\left[Y_{n:n}1_{\{Y_{n:n}<a\}}\right] \leq E\left[Y_{n:n}\right].$$

By assumption $Y_i = h(X_i)$ where $h = (1-F)^2/f$. Under our assumption that $f/(1-F)^2$ goes to infinity, the Y_i are continuously distributed random variables whose support goes down to zero. Thus the expected lowest order statistic $E\left[Y_{n:n}\right]$ must converge to zero as n gets large. This proves our claim for the sequence from Part (i). Looking at the proofs from (ii) and (iii) the results for those sequences follow easily (because the increments can be decomposed into $d(n)$ from (i) and a term which always converges to zero). □

Proof of Lemma D.1

We have to show that the sequence

$$E[X_{n:n}1_{\{X_{n:n}<c\}}]$$

is decreasing for sufficiently large n and that it switches from increasing to decreasing at most once. Recall that the distribution function of $X_{n:n}$ is given by

$$F_{n:n}(x) = 1 - (1 - F(x))^n.$$

Note furthermore that

$$\begin{aligned} E[X_{n:n}1_{\{X_{n:n}<c\}}] &= \int_0^c x f_{n:n}(x)dx \\ &= \int_0^c F_{n:n}(c) - F_{n:n}(x)dx \\ &= \int_0^c (1 - F(x))^n - (1 - F(c))^n dx \end{aligned}$$

where the middle step can easily be checked using integration by parts. Define the increment $d(n)$ by

$$\begin{aligned} d(n) &= E[X_{n+1:n+1}1_{\{X_{n+1:n+1}<c\}}] - E[X_{n:n}1_{\{X_{n:n}<c\}}] \\ &= \int_0^c F(c)(1 - F(c))^n - F(x)(1 - F(x))^n dx. \end{aligned}$$

We first show that $d(n)$ switches signs at most once: Assume that

$$0 \geq d(n) = \int_0^c F(c)(1 - F(c))^n - F(x)(1 - F(x))^n dx. \quad \text{(D.10)}$$

Note that (D.10) implies

$$\int_0^c F(c)(1 - F(c))^{n+1} - F(x)(1 - F(x))^n(1 - F(c))\, dx \leq 0. \quad \text{(D.11)}$$

Since $x \leq c$, $1 - F(c) \leq 1 - F(x)$. Thus the integral in (D.11) is made only more negative by substituting the factor $1 - F(c)$ in the second term of the integrand by $1 - F(x)$. Yet this is equivalent to

$$d(n+1) = \int_0^c F(c)(1 - F(c))^{n+1} - F(x)(1 - F(x))^{n+1}dx \leq 0. \quad \text{(D.12)}$$

Hence we have seen than $d(n) \leq 0$ implies $d(n+1) \leq 0$. Completely analogously one can show that $d(n+1) \geq 0$ implies $d(n) \geq 0$. Hence $d(n)$ switches its sign at most once and if it does, then from positive to negative.

To show that the sequence $E[X_{n:n}1_{\{X_{n:n}<c\}}]$ is decreasing from some n_0 on one has to prove that $d(n)$ becomes negative for sufficiently large n. As a preliminary observation, note that for any positive real numbers x_1, x_2, y_1 and y_2 with $0 < y_1 < y_2$ there exists an n_0 such that

$$x_1 y_1^n - x_2 y_2^n < 0 \text{ for } n > n_0. \tag{D.13}$$

This holds because (D.13) is equivalent to

$$\frac{x_1}{x_2}\left(\frac{y_1}{y_2}\right)^n < 1,$$

which holds for sufficiently large n since the left hand side converges to 0.

Now note that one can bound $d(n)$ by

$$\begin{aligned} d(n) &= cF(c)(1-F(c))^n - \int_0^c F(x)(1-F(x))^n dx \\ &< cF(c)(1-F(c))^n - \int_{\frac{c}{3}}^{\frac{c}{2}} F(x)(1-F(x))^n dx \\ &< cF(c)(1-F(c))^n - \frac{c}{6}F\left(\frac{c}{3}\right)\left(1-F\left(\frac{c}{2}\right)\right)^n. \end{aligned} \tag{D.14}$$

Applying (D.13) to the final expression in (D.14) one can conclude that $d(n) < 0$ for n sufficiently large. This completes the proof. □

Proof of Lemma D.2

We have to show that the sequence

$$p_n = E[\max(X_1+Y_1,\ldots,X_n+Y_n,X_{n+1}+1,\ldots,X_m+1)]$$

is strictly concave. Denote by $G(x)$ the distribution function of X_i+1 and by $H(x)$ the distribution function of X_i+Y_i. Note that by the properties of the exponential distribution (see, e.g., Feller (1971)) it holds that

$$G(x) = (1-e^{-x+1})1_{\{x\geq 1\}} \qquad \text{and} \qquad H(x) = 1-(1+x)e^{-x}.$$

Note furthermore that $\max(X_1+Y_1,\ldots,X_n+Y_n,X_{n+1}+1,\ldots,X_m+1)$ has distribution function H^nG^{m-n} and thus

$$p_n = \int_0^\infty 1 - H(x)^n G(x)^{m-n} dx.$$

To prove the lemma we have to show that the second increment of p_n

$$dd(n) = p_{n+1} - 2p_n + p_{n-1}$$

is negative for $1 \leq n \leq m-1$. Note that, for $0 \leq n \leq m-1$,

$$p_n = 1 + \int_1^\infty 1 - H(x)^n G(x)^{m-n} dx.$$

Thus it is straightforward to calculate, using the definitions of G and H, that for $1 \leq n \leq m-2$

$$dd(n) = \int_1^\infty -\frac{\left(1-e^{1-x}\right)^{n-k}(1-e+x)^2\left(1-e^{-x}(1+x)\right)^k}{(e^x - e)(e^x - 1 - x)} dx.$$

Since the integrand is negative for any $x > 0$, it follows that $dd(n) < 0$ for $1 \leq n \leq m-2$. We now turn to $d(m-1)$. Plugging in the definitions of H and G, it is straightforward to calculate (taking into account that $G \equiv 0$ on $[0,1]$) that

$$dd(m-1) = \int_0^1 -e^{-nx}(e^x - 1 - x)^n dx + \int_1^\infty -e^{-nx}(e^x - 1 - x)^{n-2}(1-e+x)^2 dx.$$

Since both integrands are negative, we conclude that $dd(m-1) < 0$. Thus the sequence $p_1, \ldots, p_m$ is strictly concave. □

Proof of Proposition D.4

Proposition 3.4 of Li and Shaked (2004) yields that under the assumptions of (i), for $k \geq 2$,

$$E[X_{1:k} - X_{2:k}] \geq E[Y_{1:k} - Y_{2:k}]$$

and that under the assumptions of (ii), for $k \geq 3$,

$$E[X_{2:k} - X_{3:k}] \geq E[Y_{2:k} - Y_{3:k}].$$

Applying the recurrence relations

$$kE[X_{1:k} - X_{1:k-1}] = E[X_{1:k} - X_{2:k}]$$

and

$$\frac{k}{2}E[X_{2:k} - X_{2:k-1}] = E[X_{2:k} - X_{3:k}]$$

(which follow easily from the relations given on p. 45 of David (1970), or by direct calculation) immediately allows to conclude (i) and (ii). We thus turn to (iii). As shown in the proof of Proposition 3.4 of Li and Shaked (2004),[14] $F_{2:2} \geq_{\text{EW}} G_{2:2}$ is equivalent to

$$\int_p^1 (1-u)^2 d[F^{-1}(u) - G^{-1}(u)] \geq 0$$

for all $p \in (0,1)$. By Lemma 7.1(a) of Chapter 4 of Barlow and Proschan (1981), this implies

$$\int_p^1 (1-u) d[F^{-1}(u) - G^{-1}(u)] \geq 0$$

for all $p \in (0,1)$. This is however equivalent to $F \geq_{\text{EW}} G$ as shown by Li and Shaked in the proof of their Proposition 3.4. □

Proof of Proposition D.5

For $j = 0, \ldots, n-2$ denote by β_j the probability that exactly j bidders out of bidders $3, ..., n$ know their valuation and take part in the auction. Let $(\widehat{p}_m)_{m=0,1,2}$ be the expected gross revenue given that exactly m of bidders 1 and 2 know their valuation and take part in the auction. Since p_k is the expected gross revenue conditional on that k bidders out of bidders $1, ..., n$ know their valuation and take part, we can express $\widehat{p}_m$ as

$$\widehat{p}_m = \sum_{j=0}^{n-2} \beta_j p_{j+m}.$$

Since the sequence (p_k) is strictly concave, the sequence $(\widehat{p}_m)$ of weighted averages is also strictly concave. Via $(\widehat{p}_m)$ we can express the expected gross revenue as an expectation over the number of bidders among 1 and 2 who know their valuation and take part:

$$\begin{aligned} \max_{\gamma_1} \quad & \gamma_1(\delta-\gamma_1)\widehat{p}_2 + [\gamma_1(1-(\delta-\gamma_1)) + (1-\gamma_1)(\delta-\gamma_1)]\widehat{p}_1 + (1-\gamma_1)(1-(\delta-\gamma_1))\widehat{p}_0 \\ & \text{s.t. } \gamma_1 \in [0,1], \quad \delta - \gamma_1 \in [0,1]. \end{aligned}$$

For notational convenience we do not reiterate the constraints on γ_1 in the following. The optimization problem is not changed if we subtract $\widehat{p}_0$ and divide by $\widehat{p}_2 - \widehat{p}_0$. This (and simplifying) yields

$$\max_{\gamma_1} \ \gamma_1\delta - \gamma_1^2 + [\delta - 2\delta\gamma_1 + 2\gamma_1^2]\frac{\widehat{p}_1 - \widehat{p}_0}{\widehat{p}_2 - \widehat{p}_0}$$

The strict concavity guarantees that $2(\widehat{p}_1 - \widehat{p}_0) > \widehat{p}_2 - \widehat{p}_0$. Hence we can define a *positive* constant ϵ via

$$\epsilon = \frac{\widehat{p}_1 - \widehat{p}_0}{\widehat{p}_2 - \widehat{p}_0} - \frac{1}{2} > 0.$$

[14]Note that Li and Shaked denote our "$X_{k:n}$" by "$X_{(n-k)}$".

After inserting $\frac{1}{2}+\epsilon$ for $(\widehat{p}_1-\widehat{p}_0)/(\widehat{p}_2-\widehat{p}_0)$ and simplifying, the optimization problem becomes

$$\max_{\gamma_1} \; -2\epsilon\,\gamma_1(\delta-\gamma_1)+\frac{\delta}{2}+\epsilon\delta.$$

Hence the objective function has a unique minimum at the even split-up $\gamma_1=\frac{\delta}{2}$ and is symmetric around it. It is thus strictly optimal to choose γ_1 as far away from $\frac{\delta}{2}$ as possible. Since we have to ensure $\gamma_1\in[0,1]$ and $\delta-\gamma_1\in[0,1]$ this yields

$$\{\gamma_1^*,\gamma_2^*\}=\{\lfloor\delta\rfloor,\delta-\lfloor\delta\rfloor\}.$$

□

Proof of Corollary D.1

Note that by Proposition D.5, any γ with, for $i\neq j$, $0<\gamma_i,\gamma_j<1$ can be strictly improved either by setting $\gamma_i=0$ or $\gamma_i=1$ and adjusting γ_j accordingly. Thus only allocations where at most one intermediate $\gamma_i\in(0,1)$ is chosen are candidates for optimal allocations. Since all such allocations are identical up to relabeling of bidders, and since all other allocations are strictly dominated, the corollary follows.

□

Proof of Proposition D.6

Define the function $s:[0,1]^n\to\mathbb{R}$ by

$$s(\gamma)=\sum_{i=1}^{n}\gamma_i$$

Note that trivially for any $\gamma\in[0,1]^n$ it holds that $\gamma\in H_{s(\gamma)}$.

We have to show the following: Assume that for some fixed κ there is a γ_0 with $c(\gamma_0)=\kappa$. Then γ_0 maximizes revenues among all γ with $c(\gamma)=\kappa$ if and only if $\gamma\in K_{s(\gamma)}$.

As a preliminary observation, note that, by the assumption that c is continuous and strictly increasing, it holds that for any $\kappa\in[\underline{c},\overline{c}]$ there is as an $s\in[0,n]$ and a $\gamma\in K_s$ such that $c(\gamma)=\kappa$. Note furthermore that for $\gamma\in K_s$ and $\gamma'\in K_{s'}$ with $s<s'$ it holds that $c(\gamma)<c(\gamma')$ and $\pi(\gamma)<\pi(\gamma')$.

The main step of the proof is to show that if $\gamma_0\in H_{s(\gamma_0)}\setminus K_{s(\gamma_0)}$ then for any $\gamma_1\in K_{s(\gamma_0)}$ it holds that $c(\gamma_1)<c(\gamma_0)$ and $\pi(\gamma_1)>\pi(\gamma_0)$. Note first that by symmetry all configurations in $K_{s(\gamma_0)}$ lead to the same costs and gross payoffs for the seller. Furthermore $\pi(\gamma_1) > \pi(\gamma_0)$ follows from Corollary D.1. In order to prove that the configurations in $K_{s(\gamma_0)}$ lead to lower costs than the configurations in $H_{s(\gamma_0)}$ we show that any element of $H_{s(\gamma_0)}$ can be written as a convex combination of the elements of $K_{s(\gamma_0)}$: Note that $K_{s(\gamma_0)}$ is finite and denote its cardinality by m. Denote by $\gamma^1, \ldots, \gamma^m$ the elements of $K_{s(\gamma_0)}$. Then there exist coefficients $\rho_1, \ldots, \rho_m$ with $\rho_i \in [0,1]$ and $\sum_i \rho_i = 1$ such that

$$\gamma_0 = \sum_{i=1}^{m} \rho_i \gamma^i. \tag{D.15}$$

We first show how to complete the proof given that $(D.15)$ holds. At the end we prove $(D.15)$. By the definition of quasi-concavity and by symmetry $(D.15)$ implies

$$c(\gamma_0) \geq \min_i c(\gamma^i)$$

and thus $c(\gamma_0) \geq c(\gamma_1)$ for any $\gamma_1 \in K_{s(\gamma_0)}$.

Hence, if $\gamma_0 \in H_{s(\gamma_0)} \setminus K_{s(\gamma_0)}$, for any $\gamma_1 \in K_{s(\gamma_0)}$ it holds that $c(\gamma_1) \leq c(\gamma_0)$ and $\pi(\gamma_1) > \pi(\gamma_0)$. By our preliminary observations, there exists a γ_2 with $\gamma_2 \in K_{s(\gamma_2)}$ and $c(\gamma_2) = c(\gamma_0)$. Furthermore, since $c(\gamma_2) = c(\gamma_0) \geq c(\gamma_1)$ it holds that $s(\gamma_2) \geq s(\gamma_1)$ and thus $\pi(\gamma_2) \geq \pi(\gamma_1) > \pi(\gamma_0)$. Hence an allocation γ can only be optimal within its cost-level if $\gamma \in K_{s(\gamma)}$. The rest of the desired equivalence follows from symmetry and strict monotonicity.

To conclude the proof we have to show that representation (D.15) is valid, i.e., that for all $s \in [0, n]$ any $\gamma \in H_s$ can be expressed as a convex combination of the elements of K_s. To prove this we make use of the following result[15] from the theory of convex sets: Any element of a convex polytope P can be written as a convex combination of the extremal points of P. All we have to show is that H_s is a convex polytope and that the set of extremal points of H_s is contained in K_s.[16] Obviously, H_s is a convex polytope, i.e., a bounded subset of $\mathbb{R}^n$ which is defined through linear, weak inequality constraints: Note that the constraints $\gamma_i \in [0,1]$ can be rewritten as $\gamma_i \geq 0$ and $\gamma_i \leq 1$, and $\sum_i \gamma_i = s$ as $\sum_i \gamma_i \geq s$ and $\sum_i \gamma_i \leq s$. All these constraints are linear. For the boundedness, note that $H_s \subset [0,1]^n$.

[15]See, e.g., section 3.5. of Faigle, Kern and Still (2002).

[16]While this is not necessary for our purposes, it can be shown that actually K_s is the set of extremal points of H_s.

The set of extremal points of H_s is defined by

$$Ext(H_s) = \{\gamma \in H_s | \gamma = \rho\gamma^1 + (1-\rho)\gamma^2 \text{ for } \gamma^1, \gamma^2 \in H_s, \rho \in (0,1) \Rightarrow \gamma^1 = \gamma^2 = \gamma\},$$

i.e., $Ext(H_\gamma)$ is the set of points in H_γ which cannot be written as a non-trivial convex combination of two distinct points in H_γ. For our proof, it is sufficient to show that $Ext(H_\gamma) \subseteq K_\gamma$.

We prove this by showing that any $\gamma_0 \in H_s \setminus K_s$ does not lie in $Ext(H_s)$. Consider $\gamma_0 \in H_s \setminus K_s$. Note that γ_0 has at least two entries γ_0^j and γ_0^k, $j \neq k$ which lie in $(0,1)$. We have to construct $\gamma_1, \gamma_2 \in H_s$ such that there is a $\rho \in (0,1)$ with $\gamma_0 = \rho\gamma_1 + (1-\rho)\gamma_2$. First, set all components except j and k equal in the three vectors: $\gamma_1^l = \gamma_2^l = \gamma_0^l$ for $l \notin \{j,k\}$. Furthermore set

$$(\gamma_1^j, \gamma_1^k) = (\gamma_0^j + \epsilon, \gamma_1^k - \epsilon) \text{ and } (\gamma_2^j, \gamma_2^k) = (\gamma_0^j - \epsilon, \gamma_1^k + \epsilon).$$

Note that we can choose $\epsilon > 0$ small enough such that γ_1^j, γ_1^k, γ_2^j, γ_2^k all lie in $(0,1)$. Observe furthermore that, since the entries of γ_0 sum up to s, so do the entries of γ_1 and of γ_2. Finally note that

$$\gamma_0 = \frac{1}{2}\gamma_1 + \frac{1}{2}\gamma_2$$

which proves that $\gamma_0 \notin Ext(H_s)$ and thus $Ext(H_s) \subseteq K_s$.

□

Proof of Proposition D.7

Since we know about the optimality of maximally asymmetric allocations from Proposition D.6, we can treat the decision-maker's problem as one of maximizing, respectively, $\pi_p(s) - C(s)$ and $\pi_q(s) - C(s)$ where s denotes the sum of the γ_i. By the representation (D.3) of $\pi_p(s)$ and $\pi_q(s)$, the slope of $\pi_p(s)$ is larger than the slope of $\pi_q(s)$ if $p_k - q_k$ is increasing in k. Thus if $p_k - q_k$ is increasing, the optimal choice of s under $(p_k)_k$ is larger than the optimal choice of s under $(q_k)_k$. □

Proof of Corollary D.3

Consider first the modified setting where the decision-maker's payoff sequence is given by $(q_k)_k$. For the moment assume in addition that $(q_k)_k$ is *strictly* concave. In that case we can apply Proposition D.6 and conclude that an amount s of information with $s > l^*$ (or equivalently $C(s) > \kappa^*$) is split as asymmetrically as possible. Note that thus the decision-maker's optimal payoff does not depend on the numerical values $q_0, \ldots, q_{l^*-1}$. Notably, if we substitute $q_0, \ldots, q_{l^*-1}$ by the smaller values $p_0, \ldots, p_{l^*-1}$, the revenues out of some previously suboptimal allocations may get even lower, but the payoffs from allocating s as asymmetrically as possible do not change. Thus, maximally concentrating remains optimal.

Finally, note that *weak* concavity of $q_0, \ldots, q_{l^*-1}$ is sufficient for this argument: Even if the sequence $(q_k)_k$ is linear for small k, due to the strict concavity for $k \geq l^*$ the decision-maker splits $s > l^*$ maximally asymmetrically. □

Proof of Lemma D.3

For a weakly increasing sequence $(p_k)_k$ which is strictly concave from some value k^* on we construct a weakly concave sequence $(q_k)_k$ which dominates $(p_k)_k$ and which coincides with $(p_k)_k$ from some $l^* \geq k^*$ on. Our construction which is described now is also depicted in Figure D.1. The big red dots stand for the original sequence $(p_k)_k$. The small blue dots show the values $q_0, ..., q_{l^*-1}$ which extend $p_{l^*}, p_{l^*+1}, ...$ to a concave sequence.

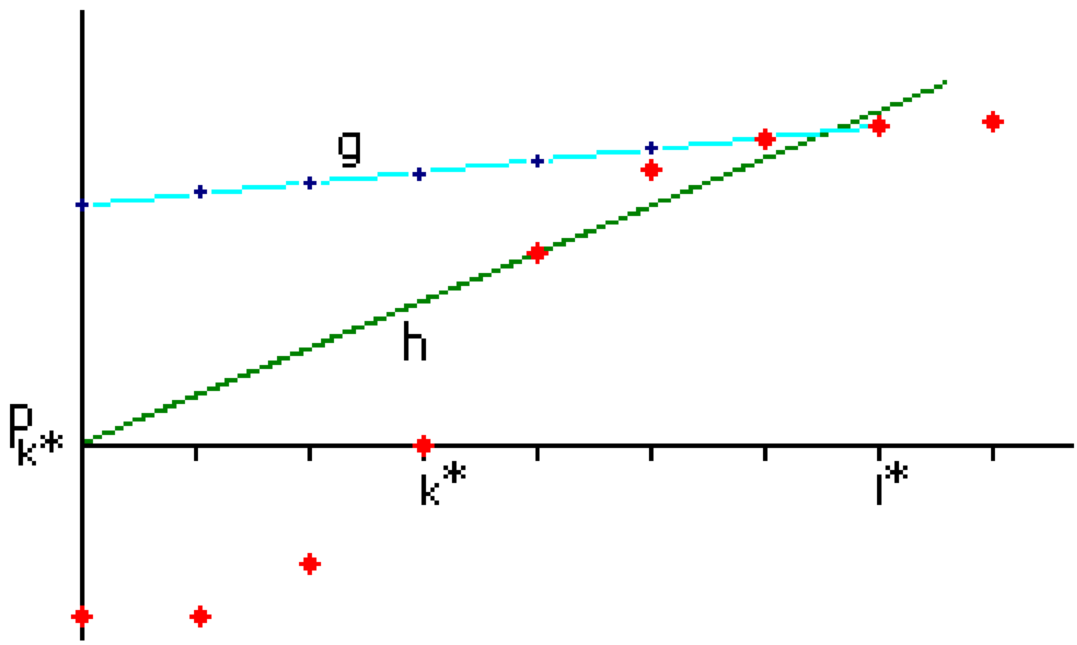

Figure D.1: Constructing a Concave Dominating Sequence

Denote by $h : \mathbb{R} \to \mathbb{R}$ the linear function interpolating the points $(0, p_k^*)$ and $(k^* + 1, p_{k^*+1})$, i.e.,

$$h(x) = p_k^* + (p_{k^*+1} - p_{k^*})\frac{x}{k^* + 1}.$$

Note that by strict concavity and weak monotonicity, $p_{k^*+1} - p_{k^*} > 0$. Now define l^* as the first integer after k^* where h dominates the sequence p_k:

$$l^* = \min\{l > k^* | h(l) > p_l\}.$$

Observe that the fact that the slope of h is constant while the increments of $(p_k)_k$ go to zero as k gets large guarantee that $l^* < \infty$. Now define a second linear function g which interpolates the points $(l^* - 1, p_{l^*-1})$ and (l^*, p_{l^*}):

$$g(x) = p_{l^*-1} + (p_{l^*} - p_{l^*-1})(x - (l^* - 1)).$$

By the definition of l^*, $h(l^*) > g(l^*) = p_{l^*}$ and $p_{l^*-1} = g(l^* - 1) \geq h(l^* - 1)$. Since g and h are linear, this implies that $h'(x) > g'(x)$ for all x. Furthermore, g and h intersect somewhere in $[l^* - 1, l^*]$. These observations imply that for all $i < k^*$

$$g(i) > h(0) = p_{k^*} \geq p_i.$$

Furthermore by the local concavity of (p_k) it follows that, for $k^* \leq i < l^*$, $g(i) \geq p_i$. We can now define the sequence (q_k) by

$$q_k = \begin{cases} g(k) & \text{for } k \leq l^* - 1 \\ p_k & \text{for } k \geq l^*. \end{cases}$$

By construction, this sequence coincides with (p_k) for large k and dominates (weakly) (p_k) for all k. Furthermore, (q_k) is weakly concave (as it follows a tangent to the concave sequence p_k on one side of l^* and the sequence (p_k) itself on the other side of l^*). □

E Proofs of Chapter 5

Proof of Proposition 5.1

Define $1_{\{i\}} = 1_{\{i \text{ wins}\}}$. To prove the formula for the entry fees, we only have to show that

$$E[(\widetilde{X}^{(1)} - \widetilde{X}^{(2)})1_{\{i\}}]$$

is the expected revenue of bidder i from the informed auction. Define $\overline{\mathcal{F}} = \sigma(\bigcup_i \mathcal{F}_i)$, i.e., $\overline{\mathcal{F}}$ is the smallest σ-algebra containing all the $\mathcal{F}_i$. Then we have, using independence of the $\mathcal{F}_i$, $\overline{\mathcal{F}}$-measurability of $1_{\{i\}}$ and the law of iterated expectations:

$$\begin{aligned} E[\widetilde{X}^{(1)} 1_{\{i\}}] = & \; E[E[X_i|\mathcal{F}_i] 1_{\{i\}}] & = E[E[X_i|\overline{\mathcal{F}}] 1_{\{i\}}] \\ = & \; E[E[X_i 1_{\{i\}}|\overline{\mathcal{F}}]] & = E[X_i 1_{\{i\}}]. \end{aligned}$$

We thus have

$$E[(\widetilde{X}^{(1)} - \widetilde{X}^{(2)}) 1_{\{i\}}] = E[(X_i - \widetilde{X}^{(2)}) 1_{\{i\}}] = e_i$$

which is the expected revenue of bidder i that the seller can extract as an entry fee. We further have to show that the seller's expected revenue equals

$$E[\widetilde{X}^{(1)}].$$

As $\sum_i 1_{\{i\}} = 1$ the agents' entry fees add up to

$$E[\widetilde{X}^{(1)} - \widetilde{X}^{(2)}].$$

Adding to this the expected selling price in the second price auction, which is $E[\widetilde{X}^{(2)}]$, we are done. □

Proof of Proposition 5.2

The formula for the entry fees is proved by showing that the difference between the fees equals $|E[X_1] - E[X_2]|$ and their sum equals $E[|\widetilde{X}_1 - \widetilde{X}_2|]$. The two agents' entry fees as calculated in the previous proposition are

$$e_1 = E[(\widetilde{X}_1 - \widetilde{X}_2) 1_{\{1 \text{ wins}\}}] \text{ and } e_2 = E[(\widetilde{X}_2 - \widetilde{X}_1) 1_{\{2 \text{ wins}\}}].$$

The difference in entry fees, $e_1 - e_2$, equals

$$\begin{aligned} & E[(\widetilde{X}_1 - \widetilde{X}_2) 1_{\{1 \text{ wins}\}}] - E[(\widetilde{X}_2 - \widetilde{X}_1) 1_{\{2 \text{ wins}\}}] \\ = & \; E[(\widetilde{X}_1 - \widetilde{X}_2) 1_{\{1 \text{ wins}\}}] + E[(\widetilde{X}_1 - \widetilde{X}_2) 1_{\{2 \text{ wins}\}}] \\ = & \; E[X_1] - E[X_2]. \end{aligned}$$

Their sum, $e_1 + e_2$, equals

$$\begin{aligned} & E[(\widetilde{X}_1 - \widetilde{X}_2) 1_{\{1 \text{ wins}\}}] + E[(\widetilde{X}_2 - \widetilde{X}_1) 1_{\{2 \text{ wins}\}}] \\ = & \; E[(|\widetilde{X}_1 - \widetilde{X}_2|) 1_{\{1 \text{ wins}\}}] + E[(|\widetilde{X}_1 - \widetilde{X}_2|) 1_{\{2 \text{ wins}\}}] \\ = & \; E[|\widetilde{X}_1 - \widetilde{X}_2|]. \end{aligned}$$

□

Proof of Proposition 5.3

The expressions for the expected revenue, and the fact that an uninformed bidder only pays a positive entry fee if he is the only uninformed bidder follow immediately from Proposition 5.1. To see that an uninformed bidder's entry fee will never exceed the entry fee of an informed bidder note that for the difference between the fees we have (in the only non-trivial case where all but one bidders get informed):

$$\begin{aligned} & E[\max(X_1 - \max(\mu, X_2, \dots, X_{n-1}), 0)] - E[\max(\mu - \max(X_1, \dots, X_{n-1}), 0)] \\ = \; & E[\max(X_1, X_2, \dots, X_{n-1})] - E[\max(\mu, X_2, \dots, X_{n-1})] \geq 0, \end{aligned}$$

where the final inequality follows from Jensen's inequality applied to the expectation with respect to X_1. For the concavity result we need to show that for all $0 \leq k \leq n-3$:

$$E[\max(X^{1:k+1}, \mu)] - E[\max(X^{1:k}, \mu)] > E[\max(X^{1:k+2}, \mu)] - E[\max(X^{1:k+1}, \mu)], \tag{E.1}$$

where we define $X^{1:0} = 0$, and

$$E[\max(X^{1:n-1}, \mu)] - E[\max(X^{1:n-2}, \mu)] > E[X^{1:n}] - E[\max(X^{1:n-1}, \mu)]. \tag{E.2}$$

We will prove (E.1) and (E.2) by showing that (E.1) holds for $k \geq 0$. (E.2) follows then immediately as the right hand side is even smaller than that of (E.1). Define $\bar{X}_i = \max(X_i, \mu)$. Let $\bar{F}$ denote the distribution function of $\bar{X}_i$. Now note that $E[\max(X^{1:n}, \mu)] = E[\bar{X}^{1:n}]$. Thus we can rewrite (E.1) to

$$E[\bar{X}^{1:k+1}] - E[\bar{X}^{1:k}] > E[\bar{X}^{1:k+2}] - E[\bar{X}^{1:k+1}]. \tag{E.3}$$

To show (E.3) note that, as $\bar{X}^{1:k}$ has distribution function $\bar{F}^k$, we have

$$E[\bar{X}^{1:k}] = \int_0^\infty 1 - \bar{F}(x)^k - \bar{F}(-x)^k dx$$

which implies

$$E[\bar{X}^{1:k+1}] - E[\bar{X}^{1:k}] = \int_{-\infty}^\infty \bar{F}(x)^k (1 - \bar{F}(x)) dx.$$

Remembering that we assumed the random variables X_i and thus also $\bar{X}_i$ to be not a.s. constant, we get that $0 < \bar{F} < 1$ on a set of positive measure. Thus the integral is strictly decreasing in k and (E.3) is proved. □

Proof of Proposition 5.4

In this and the following proofs we will make use of the fact that for $k, j \leq m$ the seller's expected revenue from revealing $X_1, \ldots, X_k$ and $Y_1, \ldots, Y_j$ can be rewritten as

$$\begin{aligned} & E[\max(X_1 + \ldots + X_k + (m-k)\mu, Y_1 + \ldots + Y_j + (m-j)\mu)] \\ = \; & m\mu + E[\max((X_1-\mu) + \ldots + (X_k - \mu), (Y_1 - \mu) + \ldots + (Y_j - \mu))] \\ = \; & m\mu + \frac{1}{2}E[|(X_1 - \mu) + \ldots + (X_k - \mu) - (Y_1 - \mu) - \ldots - (Y_j - \mu)|]. \quad \text{(E.4)} \end{aligned}$$

In order to compare the seller's expected revenue from different choices of k and j we can concentrate on the second summand in the last expression. As this expression only depends on the random variables $X_i - \mu$ and $Y_i - \mu$, we can set $\mu = 0$ without loss of generality. To prove that giving all packages to one bidder weakly dominates the equal split-up we just have to show that for independent, identically distributed, mean zero random variables X_i and Y_i

$$E[|X_1 + \ldots + X_k + X_{k+1} + \ldots + X_{2k}|] \geq E[|X_1 + \ldots + X_k - Y_1 - \ldots - Y_k|] \quad \text{(E.5)}$$

with equality exactly in the symmetric distribution case. Recall the inequality from Jagers, Kallenberg and Kroese (1995): For X and Y iid

$$E[|X+Y|] \geq E[|X-Y|]$$

(with equality exactly in the symmetric case). Setting $X = X_1 + \ldots + X_k$ and $Y = Y_1 + \ldots + Y_k$ this becomes

$$E[|X_1 + \ldots + X_k + Y_1 + \ldots + Y_k|] \geq E[|X_1 + \ldots + X_k - Y_1 - \ldots - Y_k|].$$

By the iid assumption we can substitute $Y_1, \ldots, Y_k$ on the left hand side by $X_{k+1}, \ldots, X_{2k}$ and get (E.5). (Note that a sum of iid random variables is symmetric around its mean if and only if the summands are symmetric around their means.)

□

Proof of Proposition 5.5

In order to circumvent an unnecessarily complicated notation for a simple variation of the proof of Proposition 5.4 we will just show that six packages should rather be split up into four and two than into three and three. All the other inequalities covered by Proposition 5.5 follow from analogous arguments. By (E.4) we only have to show that for iid mean-zero random variables X_i and Y_i

$$E[|X_1 + X_2 + X_3 + X_4 - Y_1 - Y_2|] \geq E[|X_1 + X_2 + X_3 - Y_1 - Y_2 - Y_2|]. \quad \text{(E.6)}$$

We start again with

$$E[|X + Y|] \geq E[|X - Y|].$$

Setting $X = X_1 + X_2 - X_3$ and $Y = X_4 + X_5 - X_6$ this becomes

$$E[|X_1 + X_2 + X_4 + X_5 - X_3 - X_6|] \geq E[|X_1 + X_2 + X_6 - X_3 - X_4 - X_5|].$$

Using the iid assumption we can rename summands on both sides of the inequality and obtain (E.6). To see that the distribution of $X_1 + X_2 - X_3$ is asymmetric if and only if the distribution of the X_i is asymmetric, note that $X_1 + X_2 - X_3$ is the sum of the (possibly asymmetric) random variable X_1 and the always symmetric random variable $X_2 - X_3$.

The other inequalities covered by the proposition follow with parallel arguments, setting $X = X_1 + X_2 + X_3 - X_4$, then $X = X_1 + X_2 + X_3 - X_4 - X_5$, etc. $\square$

Proof of Lemma 5.1

Consider $0 \leq k, j \leq m$ with $k + j \geq 2$ and $k \geq 1$. As adding another package always costs c we only have to consider by how much a package raises the seller's expected gross revenue. It is sufficient to compare the increase in revenue from revealing the k^{th} package to bidder 1 given that bidder 2 gets j packages with the revenue increase from revealing the first package to bidder 1 given that bidder 2 gets no package. So by (E.4) we have to show that for iid mean-zero random variables X_1,...,X_k and Y_1,...,Y_j

$$E[|X_1 + \ldots + X_k - Y_1 - \ldots - Y_j| - |X_1 + \ldots + X_{k-1} - Y_1 - \ldots - Y_j|] < E[|X_1|].$$

That this inequality holds weakly is an immediate consequence of the triangle inequality where we use that $E[|X_1|] = E[|X_k|]$. Noting that (since we have assumed the X_i and Y_i to be not a.s. constant) equality would contradict the independence of the X_i and Y_i it follows that we have a strict inequality and are done. $\square$

Proof of Example 5.3

Again we only need to consider gross revenues as every additional package costs c. By (E.4) we know that for the first two (symmetric) examples it is sufficient to compare the increments of the sequence 0, $E[|X_1|]$, $E[|X_1+X_2|]$, $E[|X_1+X_2+X_3|]$, $E[|X_1+X_2+X_3+X_4|]$ for independent random variables X_i distributed according to the distributions from the examples but shifted to have mean zero. For the uniform distribution on $[-\frac{1}{2}, \frac{1}{2}]$ we get

$$E[|X_1|] = \frac{1}{4}, \qquad E[|X_1 + X_2|] = \frac{1}{3},$$
$$E[|X_1 + X_2 + X_3|] = \frac{13}{32} \qquad \text{and} \qquad E[|X_1 + X_2 + X_3 + X_4|] = \frac{7}{15}.$$

For the distribution that takes $-\frac{1}{2}$ and $\frac{1}{2}$ with equal probability we get

$$E[|X_1|] = \frac{1}{2}, \qquad E[|X_1 + X_2|] = \frac{1}{2},$$
$$E[|X_1 + X_2 + X_3|] = \frac{3}{4} \qquad \text{and} \qquad E[|X_1 + X_2 + X_3 + X_4|] = \frac{3}{4}.$$

For the third example where - because of the asymmetry - the order in which packages are allocated matters, we have to compare the increments of the sequence 0, $E[|X_1|]$, $E[|X_1 + X_2|]$, $E[|X_1 + X_2 - X_3|]$, $E[|X_1 + X_2 - X_3 - X_4|]$ to see how the different packages affect the seller's revenue. Here the X_i are independent and distributed according to the exponential distribution shifted by its mean 1 to the left. We find

$$E[|X_1|] = 2e^{-1}, \qquad E[|X_1 + X_2|] = 8e^{-2},$$
$$E[|X_1 + X_2 - X_3|] = \frac{7}{2}e^{-1} \qquad \text{and} \qquad E[|X_1 + X_2 - X_3 - X_4|] = \frac{3}{2}.$$

Calculating these expectations is tedious but straightforward. Besides the formulas for the distribution functions of sums of uniformly and exponentially distributed random variables from Feller (1971), the following result from Jagers, Kallenberg and Kroese (1995) proved to be useful in the third example: Let X and Y be independent random variables with distribution functions F and G then

$$E[|X - Y|] = \int_{-\infty}^{\infty} F(x)(1 - G(x))dx + \int_{-\infty}^{\infty} G(x)(1 - F(x))dx.$$

□

Proof of Lemma 5.2

Denote by X_i and Y_i the packages of information of the two agents normalized so that they have mean zero. By (E.4) we have to prove that the sequence of the seller's gross revenues

$$P_l := m\mu + \max_{0\le k\le l} \frac{1}{2} E[|(X_1-\mu)+\ldots+(X_k-\mu)-(Y_1-\mu)-\ldots-(Y_{l-k}-\mu)|]$$

is bounded by $m\mu + \frac{\sigma}{2}\sqrt{l}$ and weakly increasing in l. Note that it is sufficient to prove this for $\mu = 0$. The upper bound on the sequence $(P_l)_l$ follows with Jensen's inequality:

$$\begin{aligned} E[|X_1+\ldots+X_k-Y_1-\ldots-Y_{l-k}|] &= E[\sqrt{(X_1+\ldots+X_k-Y_1-\ldots-Y_{l-k})^2}] \\ &\le \sqrt{\operatorname{Var}(X_1+\ldots+X_k-Y_1-\ldots-Y_{l-k})} \\ &= \sqrt{l\operatorname{Var}(X_1)} = \sqrt{l}\sigma. \end{aligned}$$

To see that the sequence $(P_l)_l$ is weakly increasing we show that the optimal split-up of $l+1$ packages is at least as good as the optimal split-up of l packages. Choose a k with $0 \le k \le l$ which maximizes

$$E[|X_1+\ldots+X_k-Y_1-\ldots-Y_{l-k}|],$$

and set $A = X_1+\ldots+X_k-Y_1-\ldots-Y_{l-k}$. Note that the function $g:\mathbb{R}\to\mathbb{R}$ given by

$$g(z) \equiv E[|A+z|]$$

is convex and thus by Jensen's inequality

$$g(0) = g(E[X_{k+1}]) \le E[g(X_{k+1})].$$

But this is the same as

$$E[|X_1+\ldots+X_k-Y_1-\ldots-Y_{l-k}|] \le E[|X_1+\ldots+X_{k+1}-Y_1-\ldots-Y_{l-k}|].$$

So we have found a split-up of $l+1$ packages that leads to a weakly higher gross revenue than the optimal split-up of l packages. Thus also the optimal split-up of $l+1$ packages leads to a weakly higher gross revenue than the optimal split-up of l packages. □

Proof of Proposition 5.6

The condition for informing bidder 1 being strictly more attractive than informing bidder 2 is

$$E[\max(X_1,\mu_Y)] > E[\max(\mu_X,Y_1)]$$

which (using the identity $\max(u,v) = \frac{1}{2}(u+v+|u-v|)$) is seen to be equivalent to

$$E[|X_1 - \mu_Y|] > E[|Y_1 - \mu_X|].$$

□

Proof of Example 5.5

Denote the packages of the first three bidders by X_i, Y_i and Z_i, respectively. We have to show that for both distributions all three packages should be given to three different bidders no matter whether the number of bidders is exactly 3 or greater. Thus we have to show that

$$\begin{aligned} E[\max(X_1 + X_2 + X_3, 3\mu)] &< E[\max(X_1 + 2\mu, Y_1 + 2\mu, Z_1 + 2\mu)] \\ E[\max(X_1 + X_2 + \mu, Y_1 + 2\mu, 3\mu)] &< E[\max(X_1 + 2\mu, Y_1 + 2\mu, Z_1 + 2\mu)] \end{aligned}$$

for exactly three bidders and

$$\begin{aligned} E[\max(X_1 + X_2 + X_3, 3\mu)] &< E[\max(X_1 + 2\mu, Y_1 + 2\mu, Z_1 + 2\mu, 3\mu)] \\ E[\max(X_1 + X_2 + \mu, Y_1 + 2\mu, 3\mu)] &< E[\max(X_1 + 2\mu, Y_1 + 2\mu, Z_1 + 2\mu, 3\mu)] \end{aligned}$$

for more than three bidders. Obviously, the first two inequalities imply the second two, and the first two are true because for the exponential distribution (where $\mu = 1$) we have

$$\begin{aligned} E[\max(X_1 + X_2 + X_3, 3\mu)] &= 3 + \frac{27}{2}e^{-3}, \\ E[\max(X_1 + X_2 + \mu, Y_1 + 2\mu, 3\mu)] &= 3 + e^{-1} + 4e^{-2} - \frac{7}{4}e^{-3}, \\ E[\max(X_1 + 2\mu, Y_1 + 2\mu, Z_1 + 2\mu)] &= \frac{23}{6}, \end{aligned}$$

and for the uniform distribution (where $\mu = \frac{1}{2}$) we have

$$\begin{aligned} E[\max(X_1 + X_2 + X_3, 3\mu)] &= \frac{109}{64}, \\ E[\max(X_1 + X_2 + \mu, Y_1 + 2\mu, 3\mu)] &= \frac{671}{384}, \\ E[\max(X_1 + 2\mu, Y_1 + 2\mu, Z_1 + 2\mu)] &= \frac{7}{4}. \end{aligned}$$

□