Auctioning risk: The all-pay auction under mean-variance preferences*

Bettina Klose[♮] Paul Schweinzer[♯]

Abstract

We analyse the all-pay auction with incomplete information and variance-averse bidders. We characterise the symmetric equilibrium for general distributions of valuations and any number of bidders. Variance aversion is a sufficient assumption to predict that high-valuation bidders increase their bids relative to the risk-neutral case while low types decrease their bid. Considering an asymmetric two-player environment with uniform valuations, we show that a more variance-averse type always bids higher than her less variance-averse counterpart. Utilising our analytical bidding functions we discuss all-pay auctions with variance-averse bidders from a designer's perspective. We extend our basic model to include noisy signals and allow for the possibility of variance-seeking preferences and type-dependent attitudes towards risk.
(JEL *C7, D7, D81.* Keywords: *Auctions, Contests, Mean-Variance preferences.*)

1 Introduction

In economic contests players make irrecoverable investments in order to increase their chances of winning a prize. As such the nature of a bid in an all-pay auction is very similar to choosing a lottery. A high bid corresponds to lower payoffs in both the states of winning and losing, but increases the likelihood of the positive state. A low bid yields a lower probability of winning, but payoffs in both states are higher. Although the all-pay auction is therefore a model of risky choices, typical analyses focus on risk-neutral players, i.e., the maximisation of expected payoffs. Mean-variance preferences (Markowitz, 1952) have long been successfully applied to portfolio choice investment problems where asset managers evaluate alternative portfolios on the basis of the mean and variance of their return. It therefore may be surprising that the mechanism design literature and, specifically, the large literature on auctions has not yet addressed the decision making problem of players endowed with mean-variance preferences over their wealth. The present paper attempts to close this gap by

*Thanks to Olivier Bos, Thomas Giebe, Todd Kaplan, Dan Kovenock, Sergio Parreiras, Ron Siegel and anonymous referees for helpful comments and discussions. Both authors are grateful for the hospitality of their co-authors institutions and for financial support through the University of York Research and Impact Support Fund. Klose gratefully acknowledges the financial support from the European Research Council (ERC Advanced Investigator Grant, ESEI-249433) and the Swiss National Science Foundation (SNSF 100014 135257). Schweinzer appreciates the generous hospitality of CESifo, Munich. [♮]University of Zurich, Department of Economics, Blümlisalpstr. 10, CH-8006 Zurich, Switzerland, bettina.klose@econ.uzh.ch. [♯]Department of Economics, University of York, Heslington, York YO10 5DD, United Kingdom, paul.schweinzer@york.ac.uk. (23-Jun-2014)

fully characterising bidding and revenue-optimal sales behaviour in one of the standard auction types, the all-pay auction. This auction type may be viewed as a natural candidate because it exposes a bidder to the inherent risk of either winning the object (potentially at a bargain) or losing one's bid without gaining anything.

We offer a group of motivational examples underpinning the relevance of this approach. All of our examples have two random components: i) the 'endogenous' variance centred on the bidding process, i.e., the tension between winning a prize (consisting of a private valuation minus the own bid) and paying the own bid for sure, and ii) the 'exogenous' variance of the prize (or outside option) itself. As our leading example we consider the research and design competition on 'wearable tech,' widely argued to be the 'next big thing' in consumer electronics.[1] Technology firms Apple, Google, Samsung and others compete on the first breakthrough in this market. In the smartwatch war, these giants were beaten to the market by puny Pebble Technology Corp, a crowdfunded startup from California.[2] Similarly extreme asymmetries in budgets can be seen in the glasses market where the highly publicised Google Glass market entry was preempted by Vuzix, a comparably small technology firm from Rochester, New York. Both cases have in common that deep pockets firms compete in the same market with firms whose budgets they dwarf. This opens the door to predatory pricing strategies and thus induces very different risk preferences among competitors.[3]

A further and distinct set of examples lies in the practice of introductory offers ('burning money') with which firms try to ascertain uncertain future market shares through certain upfront losses. Our setup is as applicable to an entrepreneur's choice of team composition as it is to the research portfolio selection of heads of R&D or similar institutions. Similarly, our model seems to fit well with patent races under rivalry as analysed, for instance, by Dasgupta (1986) in a full information setup. Finally, the classical portfolio selection problem seems to be related as, clearly, portfolio choice is usually based on mean-variance considerations and anecdotal evidence is available which underlines all-pay aspects of fund managing practice.[4]

What is a general motivation to consider risk aversion in winner-pay auctions? A bidder in a (first-price) winner-pay auction controls, through her bid, both the probability of winning and the amount she wins. A risk averse bidder is willing to sacrifice some of this payoff (the individual value minus her bid) for a higher probability of winning (through a higher bid). Hence, in a first-price, winner-pay auction, risk aversion causes an increase in equilibrium bids relative to the risk neutral case.[5] In all-pay auctions, in addition, increasing one's bid has the direct negative effect of increasing the certain payment *independently* of both other effects. In consequence, a low valuation

[1] For sources, see Wired, *Why Wearable Tech Will Be as Big as the Smartphone*, 17-Dec-2013 and NBC news, *Samsung's Galaxy Gear marks official start to smartwatch war*, 4-Sep-2013.

[2] Pebble is a crowd funding success story, it "set out to raise $100,000. After 37 days, it had raised $10,266,845, which is the highest amount ever secured on the [Kickstarter] platform." Wired, *How Pebble kickstarted a trend*, 23-Jan-2014.

[3] For early accounts of predatory pricing and its effects on risk see, for instance, Tesler (1966) or Milgrom and Roberts (1982).

[4] "Five-star funds. Four-star funds. Those seem to be the only mutual funds people want to buy." Investors Business Daily, "Making Money in Mutuals: Don't Focus too Narrowly on Star Ratings," by A. Shell, 22-Jun-1998, cited in Bagnoli and Watts (2000). Hence, although all funds invest efforts, only the most highly ranked funds obtain significant investments.

[5] See Maskin and Riley (1984).

bidder under risk aversion wants to *decrease* her losses while a risk averse, high-type bidder wants to *increase* her probability of winning through more aggressive bidding in the all-pay auction.

Apart from intellectual curiosity, we field three main arguments in order to justify the attention we place on mean-variance preferences in this paper. First, the typically employed risk-neutral, expected payoff analysis of auctions simply ignores any risk considerations; compared with that, a mean-variance analysis certainly represents progress. Second, if all relevant probability distributions have the same elliptic form, then the mean and variance represent a sufficient statistic to identify the true distribution of returns. Then, the mean-variance approach does not differ from a *full* account of expected utility using a general representation of risk aversion. Third, financial practitioners make the vast majority of their day-to-day portfolio choice decisions on the basis of the mean and variance of portfolios. It would seem likely that this group could benefit from a similar representation of their choices for auctioning activities.

Literature

To the best of the authors' knowledge there are no existing papers which analyse auctions or incomplete information contests under mean-variance preferences. Most existing work on risk aversion in contests applies to full information Tullock contests.[6] An attempt to model mean-variance preferences in this full information case is Robson (2012) who derives an 'irrelevance result' in the sense that for two-player Tullock contests bidding behaviour is not affected by the introduction of an aversion to variance. A more general analysis in terms of risk aversion of the same setup is Cornes and Hartley (2012b) who focus on existence questions of both symmetric and asymmetric Nash equilibria (for the case of loss aversion see Cornes and Hartley, 2012a). The only existing works on risk aversion for the incomplete information all-pay auction of which we are aware are Fibich, Gavious, and Sela (2006), Cingottini and Menicucci (2006), and Parreiras and Rubinchik (2010).[7] Fibich, Gavious, and Sela (2006) show that an analytic characterisation of equilibrium strategies cannot be usually obtained for von Neumann-Morgenstern risk-averse players. Thus, contrasting our fully analytical approach, they turn to perturbation analysis to obtain revenue rankings for the case that players are 'almost' risk neutral. Esö and White (2004) show that under special conditions on valuations, decreasingly absolute risk averse players prefer the first-price auction to the all-pay auction. Fibich, Gavious, and Sela (2006) extend this ranking to the case of general risk aversion for independent valuations. Their results are limited, however, by the fact that they cannot generally obtain analytic forms of the equilibrium bidding strategies of risk averse players. We can overcome this limitation at the price of focusing attention to the class of linear mean-variance preferences.

Cingottini and Menicucci (2006) study an environment composed of ex-ante symmetric bidders sharing the same preferences exhibiting constant absolute risk aversion. They find that it is revenue-optimal for the seller to exclude all but two randomly chosen competitors. Their result, which is

[6] An analysis of the (repeated) full information case for more general success probabilities is Ireland (2004).

[7] It may be worth pointing out that the analysis of risk-aversion in Lazear and Rosen (1981) also boils down to mean-variance preference analysis as their output noise term is fully characterised by the mean and variance of the Normal distribution.

contrary to the monotonicity of revenue in the risk neutral case, is obtained provided that bidders are either highly risk averse or very likely to possess a particular, known valuation.

Parreiras and Rubinchik (2010) analyse bidding behaviour in contests where three or more players draw their valuations from asymmetric supports and may have asymmetric attitudes toward risk. They find that these ex-ante asymmetries may lead to player drop-out or, for sufficiently risk-averse players, the use of discontinuous 'all-or-nothing' strategies. Thus, both cases exhibit behaviour which is very different from the standard ex-ante symmetric equilibrium case. Although the authors cannot explicitly determine the equilibrium bidding functions in general, they construct a simple check for whether or not a particular bid can be part of a player's equilibrium strategy. This test is used to establish the above participation conclusions.

Papers relating to the analysis of risk aversion in general winner-pay auction environments are Maskin and Riley (1984) and Matthews (1987), both discussing risk-averse bidders' behaviour in auctions, Esö and White (2004), analysing precautionary bidding in auctions, and Esö and Futó (1999), who derive the revenue-optimal strategy for a risk-averse seller, and Hu, Matthews, and Zou (2010) who discuss reserve prices. The existing analyses of asymmetric auctions, for instance Amann and Leininger (1996), Lizzeri and Persico (2000), Maskin and Riley (2000), Fibich, Gavious, and Sela (2004), Parreiras and Rubinchik (2010), Kirkegaard (2012), or Kaplan and Zamir (2012), typically employ asymmetric distributions (or supports) while we use our idiosyncratic variance-aversion parameter. Next to Parreiras and Rubinchik (2010), to the best of the authors' knowledge, this is the only paper to analyse bidding in a contest when players are asymmetric in their degree of risk-aversion.

In terms of revenue and payoff analysis, Matthews (1987) compares payoffs for risk averse behaviour when bidders exhibit constant and increasing absolute risk aversion (CARA and IARA, respectively). For CARA, he finds that bidders are indifferent between first- and second-price auctions, while for IARA bidders prefer the first-price auction. Smith and Levin (1996) show that this ranking can be reversed under decreasing absolute risk aversion.

The present paper is dealing with variance aversion which, in general, is different from risk aversion.[8] Mean-variance preferences can be transformed into expected utility form under certain assumptions on the location, scale, and concordance parameters of the environment. For the precise relation of von Neumann-Morgenstern preferences to mean-variance preferences, see, for instance, Sinn (1983), Kroll, Levy, and Markowitz (1984), Mayer (1987), or, more recently, Eichner (2008), or Eichner and Wagener (2009). A recent survey and discussion of the differences between the approaches is presented by Markowitz (2012).

2 The model

There is a seller with one indivisible object for sale. The seller's valuation of the item is (normalised to) zero. There are $n \geq 2$ potential buyers with valuations θ_i, $i \in \mathcal{N} = \{1, 2, \ldots, n\}$, respectively. The own valuation is private information of each buyer and all players' valuations, θ_i, $i \in \mathcal{N}$, are

[8] For a discussion of the differences see, for instance, Rothschild and Stiglitz (1970).

assumed to be independent draws from the same increasing and absolutely continuous distribution F. Let $f(\cdot) = F'(\cdot)$ represent the associated probability density function and $[\underline{\theta}, \bar{\theta}] = [0, 1]$ its support.

After realising their own valuations of the object, θ_i, all players simultaneously submit their bids, b_i, $i \in \mathcal{N}$. The player with the highest bid receives the object and all players forgo their bids. Player i's payoff is given by

$$\pi_i(b_i, b_{-i}; \theta_i) = \begin{cases} \theta_i - b_i & \text{if } b_i > b_j \forall j \neq i, \\ \frac{1}{m}\theta_i - b_i & \text{if } i \in Q = \{j \in \mathcal{N} | b_j = \max_{k \in \mathcal{N}} b_k\}, m = |Q|, \\ -b_i & \text{if } \exists j : b_i < b_j. \end{cases}$$

When buyers have mean-variance preferences, they maximise an objective function $u_i(\mu_i, \sigma_i^2)$, which is increasing in the expected payoff, μ_i, and decreasing in the variance of their payoff, σ_i^2. For our analytical investigation we use the following simple linear representation of mean-variance preferences[9]

$$u_i(\mu(b, \theta_i), \sigma^2(b, \theta_i)) = \mu(b, \theta_i) - \nu_i \sigma^2(b, \theta_i), \tag{1}$$

where the parameter $\nu_i \in [0, {}^1\!/_2]$ accounts for player i's variance-aversion.[10] The case of $\nu_i = 0$ represents the standard case of risk-neutral expected payoff maximisation. Bounding the degree of variance aversion from above guarantees the existence of a pure strategy equilibrium.

Provided a player's knowledge of her own type, her bid can be interpreted as choosing a lottery with two possible outcomes. If the player submits the highest bid, she will receive her valuation of the prize (represented by her type) minus the cost of her bid. In all other events, she will lose her bid. Note that the payoff difference between these two outcomes remains fixed for any bid and is just equal to the player's type. However, selecting a higher bid does not only decrease the respective payoffs for both outcomes, but also moves probability mass from the loosing to the winning outcome.

In section 5 we discuss various ways to generalise this model. We introduce exogenous noise to the model through independent shocks on the winning and losing outcomes. In this case, a player's type is merely a noisy signal of her true valuation for the prize which will only be realised after the contest. We further consider the effects of variance seeking behaviour in contests and type-dependent risk preferences, i.e. ex-ante symmetric players have private information about their type (θ_i, ν_i).

[9] A large body of empirical and theoretical work employs variants of this simple form on the basis of both tractability and testability. For discussions see, for instance, Tsiang (1972), Coyle (1992), Saha (1997) or the textbook treatment in Sargent and Heller (1987, p154–5). Recently, Chiu (2010) discusses the applicability of mean-variance preferences of this form to a large set of problems in finance and economics in choice theoretic terms. For an easy generalisation to the case of mean plus any monomial function of the variance see the discussion in the appendix.

[10] In section 5.2 we extend the range of admissible values of ν to $[-1/2, +1/2]$.

3 Bidding behaviour

3.1 The symmetric case: n identical bidders

Under the first-price, all-pay auction, a type-θ_i bidder's expected payoff when issuing a bid of b is given by

$$\pi(b,\beta;\theta_i) = \int_0^{\beta^{-1}(b)} \cdots \int_0^{\beta^{-1}(b)} \theta_i f(\theta_1)\cdots f(\theta_{n-1}) d\theta_1 \cdots d\theta_{n-1} - b \tag{2}$$

where $\beta(\theta)$ is the tentative symmetric equilibrium bid issued by a type-θ player. We conjecture that the function $\beta(\theta)$ is non-decreasing and denote the highest type who submits a bid no higher than b by $\theta = \beta^{-1}(b)$. It is well known, for instance from Milgrom (2004, p119), that the strategies

$$\beta_{\mathsf{rn}}(\theta) = \theta (F(\theta))^{n-1} - \int_0^{\theta} (F(\vartheta))^{n-1} d\vartheta, \tag{3}$$

maximise (2), hence constituting a symmetric equilibrium if players simply maximise their expected payoffs (i.e., $\nu_i = 0$ for all $i \in \mathcal{N}$).

With mean-variance preferences, symmetric players with $\nu \equiv \nu_1 = \cdots = \nu_n$ choose a bidding function which maximises (1), taking into account their payoff variance in addition to their expected payoff. These are given for the first-price, all-pay auction as

$$\begin{aligned} \mu(b,\theta_i) &= \theta_i (F(\beta^{-1}(b)))^{n-1} - b; \\ \sigma^2(b,\theta_i) &= (F(\beta^{-1}(b)))^{n-1} (\theta_i - b - \mu)^2 + (1 - (F(\beta^{-1}(b)))^{n-1}) (-b-\mu)^2 \\ &= (F(\beta^{-1}(b)))^{n-1} (1 - (F(\beta^{-1}(b)))^{n-1}) \theta_i^2. \end{aligned} \tag{4}$$

Inserting these back into the player's objective[11] and rearranging gives

$$u_i(b,\theta_i) = \theta_i (F(\beta^{-1}(b)))^{n-1} (1 - \nu\theta_i + \nu\theta_i (F(\beta^{-1}(b)))^{n-1}) - b. \tag{5}$$

The first-order condition for maximisation of (5) with respect to b is[12]

$$\theta_i \left(1 - \nu\theta_i + 2\nu\theta_i (F(\beta^{-1}(b)))^{n-1}\right) (n-1) (F(\beta^{-1}(b)))^{n-2} f(\beta^{-1}(b)) \frac{\partial \beta^{-1}(b)}{\partial b} = 1. \tag{8}$$

In the symmetric equilibrium $b = \beta(\theta_i)$, this yields the first-order differential equation

$$\beta'(\theta_i) = \theta_i \left(1 - \nu\theta_i + 2\nu\theta_i (F(\theta_i))^{n-1}\right) (n-1) (F(\theta_i))^{n-2} f(\theta_i). \tag{9}$$

[11] Note that in our model using the modified mean-variance approach due to Blavatskyy (2010) would lead to qualitatively the same results since the mean absolute semideviation is $r(b,\theta_i) = (F(\beta^{-1}(b)))^{n-1}(1-(F(\beta^{-1}(b)))^{n-1})\theta_i$.

[12] The second-order condition for the case of two players is

$$\frac{2\theta_i\nu - 4\theta_i\nu F(\theta_i) - 1}{\theta_i^2 f(\theta_i)(1 - \theta_i\nu + 2\theta_i\nu F(\theta_i))^2} < 0. \tag{6}$$

The condition for the general case of $n > 2$ players is more involved and relegated to the appendix. In our model $\theta_i\nu \le \frac{1}{2} < 1$, therefore, a sufficient condition for the general second-order condition to hold is that

$$F(\theta)^{n-1} \ge \frac{1}{2} - \frac{1}{2\nu(\theta+1)}. \tag{7}$$

This condition (7) generally holds if ν is sufficiently small such that $\theta\nu \le \frac{1}{2}$ for all θ in the support of the type distribution $F(\cdot)$, but otherwise imposes a restriction on the distribution of types and/or the number of players.

This differential equation together with the boundary condition $\beta(0) = 0$ is solved (through repeated integration by parts) by the bidding function

$$
\begin{aligned}
\beta_{\mathsf{mv}}(\theta_i) =& \theta_i (F(\theta_i))^{n-1} - \int_0^{\theta_i} (F(\vartheta))^{n-1}\, d\vartheta - \nu\theta_i^2 (F(\theta_i))^{n-1} + \\
& \nu\theta_i^2 (F(\theta_i))^{2(n-1)} + \nu \int_0^{\theta_i} 2\vartheta (F(\vartheta))^{n-1}\, d\vartheta - \nu \int_0^{\theta_i} 2\vartheta (F(\vartheta))^{2(n-1)}\, d\vartheta \\
=& \beta_{\mathsf{rn}}(\theta_i) - \nu \left(\theta_i^2 \left(F(\theta_i)^{n-1} - (F(\theta_i))^{2(n-1)} \right) - \int_0^{\theta_i} 2\vartheta \left((F(\vartheta))^{n-1} - (F(\vartheta))^{2(n-1)} \right) d\vartheta \right) \\
=& \beta_{\mathsf{rn}}(\theta_i) - \nu \int_0^{\theta_i} \vartheta^2 (n-1) (F(\vartheta))^{n-2} f(\vartheta) \left(1 - 2(F(\vartheta))^{n-1} \right) d\vartheta.
\end{aligned}
\tag{10}
$$

Two comments are in order before we proceed with our analysis. First, note that, from (9), β_{mv} is a strictly increasing function for any $\nu > -1$ confirming our tentative monotonicity conjecture. Second, it is worthwhile to pay some attention to the second order condition of this problem. Intuitively, as players become very variance averse, bidding zero along with the resulting certain payoff of zero becomes more attractive. However, there cannot exist a pure strategy equilibrium in which all types bid zero. Therefore, in order to be able to solve for the desired increasing bidding function, we need to bound the variance aversion parameter. Inequality (7) illustrates how the distribution of types places an upper bound on the degree of variance-aversion which we may consider and vice versa. Interpreting $H_\nu(\theta) = \frac{1}{2} - \frac{1}{2\nu(\theta+1)}$ as a c.d.f. with support in $[\frac{1-\nu}{\nu}, +\infty]$ and an atom of size $\frac{1}{2}$ at infinity, condition (7) states that H_ν first order stochastically dominates F^{n-1}. As ν increases H_ν increases in the sense of first order stochastic dominance and (7) fails as ν grows too large. Higher degrees of variance aversion may be considered when F places sufficient mass on low types.

In order to confirm that the strategies (10) are indeed the symmetric equilibrium, we show that no profitable deviation exists. A player's utility from any bid b is provided in (5). Consider a player of type θ who bids $\beta_{mv}(z), z \in [0,1]$ against opponents who bid according to β_{mv} and their true type.[13] Then by (5) and (10)

$$
\begin{aligned}
u_i(\beta_{mv}(z), \theta) =& \theta(1-\nu\theta) F^{n-1}(z) + \nu\theta^2 F^{2n-2}(z) - z F^{n-1}(z) + \int_0^z F^{n-1}(\vartheta) d\vartheta \\
& + \nu \int_0^z \vartheta^2 (n-1) F^{n-2}(\vartheta) f(\vartheta) (1 - 2F^{n-1}(\vartheta)) d\vartheta.
\end{aligned}
$$

Hence, setting $z = \theta$ (i.e., bidding according to (10)) is optimal if

$$
\begin{aligned}
\frac{\partial u_i(\beta_{mv}(z), \theta)}{\partial z} &= \theta(1-\nu\theta)(n-1) F^{n-2}(z) f(z) + \nu\theta^2 2(n-1) F^{2n-3}(z) f(z) - F^{n-1}(z) \\
&\quad - (n-1) z F^{n-2}(z) f(z) + F^{n-1}(z) + \nu z^2 (n-1) F^{n-2}(z) f(z) (1 - 2F^{n-1}(z)) \\
&= (\theta - z)(n-1) f(z) F^{n-2}(z) (1 - \nu(\theta + z) + 2\nu(\theta + z) F^{n-1}(z))
\end{aligned}
$$

[13] Bids below $\beta_{mv}(0)$ are not feasible and any bid $b > \beta_{mv}(1)$ is strictly dominated by $\beta_{mv}(1)$ as b yields strictly lower expected payoff $\mu(b, \theta) < \mu(\beta_{mv}(1), \theta)$ without reducing the payoff variance as winning probabilities are identical under both strategies.

is positive for $z < \theta$, negative for $z > \theta$ and zero at $z = \theta$. This is the case if the factor $(1 - \nu(\theta + z) + 2\nu(\theta + z)F^{n-1}(z))$ is non-negative for all $(\theta, z) \in [0, 1]^2$ and always holds if $\nu \leq \frac{1}{2}$.

Our next result shows that low types submit lower bids under mean-variance preferences, while high types submit higher bids under mean-variance preferences than if they were to maximise expected payoff only.

Proposition 1 (Single-crossing). *Either $\beta_{mv}(\theta) \leq \beta_{rn}(\theta)$ for all $\theta \in [0, 1]$ or there exists a $\hat{\theta}$ in the support of F such that $\beta_{mv}(\theta) \leq \beta_{rn}(\theta)$ for $\theta \leq \hat{\theta}$, $\beta_{mv}(\hat{\theta}) = \beta_{rn}(\hat{\theta})$ and $\beta_{mv}(\theta) > \beta_{rn}(\theta)$ for $\theta > \hat{\theta}$.*

Proof. Note that the symmetric equilibrium strategy can be written as

$$\beta_{\text{mv}}(\theta) = \beta_{\text{rn}}(\theta) - \nu \int_0^\theta G(\vartheta)H(\vartheta)d\vartheta,$$

where $G(\vartheta) = \vartheta^2(n-1)(F(\vartheta))^{n-2}f(\vartheta)$ and $H(\vartheta) = 1 - 2(F(\vartheta))^{n-1}$. F is a cumulative distribution function with density f, therefore $G(\vartheta) \geq 0$ for all $\vartheta \in [\underline{\theta}, \bar{\theta}]$. $H(\vartheta)$ is a continuous and decreasing function with $H(\underline{\theta}) = 1$ and $H(\bar{\theta}) = -1$. Hence, $\int_0^\theta G(\vartheta)H(\vartheta)d\vartheta > 0$ for sufficiently small $\theta > 0$ and if $\int_0^{\bar{\theta}} G(\vartheta)H(\vartheta)d\vartheta = 0$ for any $\hat{\theta} > 0$, then $\int_0^\theta G(\vartheta)H(\vartheta)d\vartheta < 0$ for all $\theta > \hat{\theta}$. □

This result is qualitatively in line with Propositions 1 and 2 in Fibich, Gavious, and Sela (2006). The intuition is that low-valuation bidders expect to lose in a symmetric equilibrium and therefore decrease their bids in order to keep their variance low. High-valuation bidders, by contrast, are likely to win and therefore increase their bids in line with variance compression. Proposition 1 says that there is only a single type of bidder who submits the same equilibrium bid for any value of ν.[14] From (10) it is evident that this type $\hat{\theta}$ does not depend on the variance aversion parameter ν, but is fully characterised by the distribution of types, F, and number of players, n.

Our proposition 1 shows that variance-aversion alone is sufficient to observe the bidding behaviour found by Fibich, Gavious, and Sela (2006) for general risk-aversion as modelled by a utility function that is concave in payoffs. It is not obvious that the simple mean-variance approach would yield this qualitatively identical finding to the full account of all moments. The main advantage we obtain from knowing equilibrium bidding functions is our ability to further investigate the impact that variance-aversion has on auction characteristics that are relevant from a designer's perspective such as revenue or bidder participation. While the general approach of Fibich, Gavious, and Sela (2006) results in the same qualitative observations regarding the bidding function, it does not allow for such an analysis. Fibich, Gavious, and Sela (2006) use perturbation analysis, therefore considering only infinitesimally risk averse (i.e., almost risk neutral) players. Our analysis generalises the analysis to players with a significant variance aversion parameter.

Corollary 1. *As the number of participating bidders n expands,*

1. *the convexity of the bidding function $\beta_{mv}(\theta)$ increases, i.e., low types decrease their bids and high types increase their bids relative to the case with a lower number of bidders;*

[14] This distortion of bids seems to correspond to experimental evidence. Both Barut, Kovenock, and Noussair (2002) and Noussair and Silver (2006) report bidding behaviour along these lines in all-pay auctions with private valuations. Moreover, gender differences in competitions à la Gneezy, Niederle, and Rustichini (2003) can be explained by our (a)symmetric risk-aversion result.

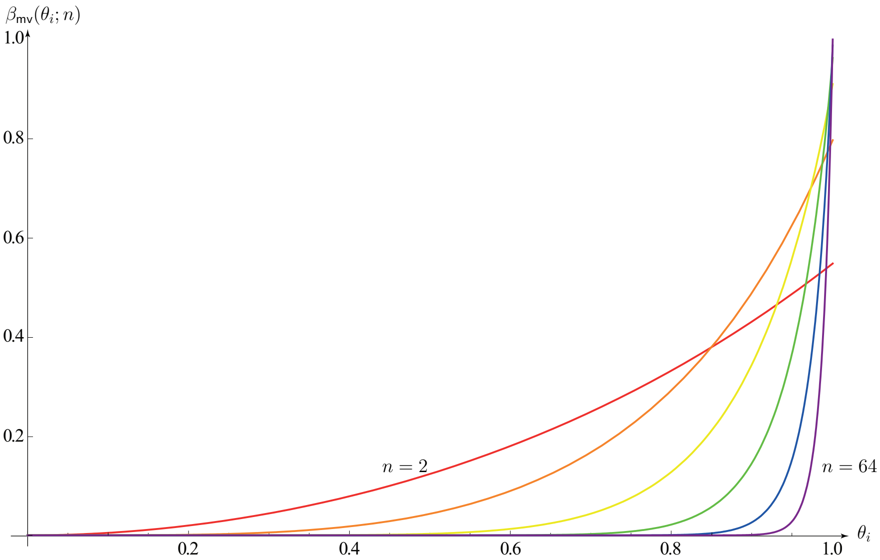

Figure 1: Equilibrium bidding functions for the all-pay auction under mean-variance preferences for uniformly distributed types, $\nu = 1/4$, and $n \in \{2, 4, 8, 16, 32, 64\}$ players, respectively (sorted in the colours of the rainbow from red to violet).

2. *the type $\hat{\theta}$ who issues the same bid under mean-variance and risk-neutral von Neumann-Morgenstern preferences shifts to the right.*

Proof. Consider the derivative of (9) with respect to n

$$\theta f(\theta) F(\theta)^{n-3} \left[(1 - \theta\nu) F(\theta)(1 - \kappa) + 2\theta\nu F(\theta)^n (1 - 2\kappa)\right] \tag{11}$$

where $\kappa = -(n-1)\log(F(\theta))$. Notice that $\log(F(\theta)) \leq 0$ and $\log(F(\theta))$ is strictly increasing in θ with $\log(F(\theta)) \to -\infty$ as θ approaches the lower bound of the support of its distribution and $\log(F(\theta)) \to 0$ as θ approaches the upper bound of the support of its distribution. Therefore, for sufficiently small θ, (11) becomes negative. Similarly, for θ sufficiently large, (11) is positive. □

Corollary 1 shows that the comparative statics of equilibrium bidding in the number of participants in an all-pay auction qualitatively do not differ whether bidders are risk neutral or symmetrically variance averse. However, as shown in the second part of the corollary, the convergence to a distribution where only high types submit positive bids is faster the more variance averse the bidders are.

3.1.1 Examples

The uniform distribution

In the following, we exemplarily illustrate our findings for the case of n players when values are drawn from a uniform distribution over the interval $[0, 1]$. In this case, the expression for the objective of a bidder with mean-variance preferences simplifies to

$$u_i(b, \theta_i) = \theta_i(\beta^{-1}(b))^{n-1}(1 - \nu\theta_i + \nu\theta_i(\beta^{-1}(b))^{n-1}) - b \tag{12}$$

which determines the symmetric equilibrium bidding functions as

$$b^* = \beta(\theta_i) = \frac{n-1}{n}\theta_i^n + \nu\left(\frac{n-1}{n}\theta_i^{2n} - \frac{n-1}{n+1}\theta_i^{n+1}\right) + C \tag{13}$$

for some constant C which is zero because a type-0 will not make a positive bid. Figure 2 compares this equilibrium bidding behaviour with that under standard risk-neutral von Neumann-Morgenstern preferences for two players. As seen before, the bidding behaviour of low-intermediate valuation types is more aggressive under expected payoff maximisation while high types submit higher bids under mean-variance preferences.

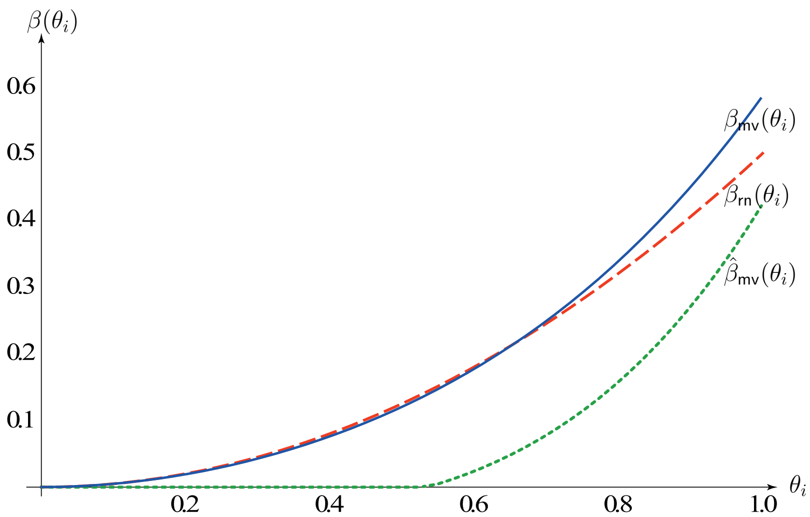

Figure 2: Equilibrium bidding functions for the all-pay auction under risk-neutrality (dashed, $\beta(\theta_i) = \theta_i^2/2$) and mean-variance preferences ($\nu = 1/2$, solid). The dotted bidding function results under mean-variance preferences if the prize itself is risky $\hat{\varepsilon} - \hat{\delta} = \frac{1}{4}$.

Other distributions

The table below displays the symmetric, two-player equilibrium bidding functions for the most commonly used distribution functions when players have mean-variance preferences and values are i.i.d. according to the specified distribution.

Distribution	$F(\theta)$	$\beta(\theta)$
Uniform[0,b]	θ/b	$\theta_i^2(\theta_i\nu(3\theta_i - 2b) + 3b)/6$
Power[0,1]	θ^α	$\dfrac{\alpha\left(\theta_i^{\alpha+1}(-(\alpha+1)\theta_i\nu + \alpha + 2) + (\alpha+2)\nu\theta_i^{2\alpha+2} - (\alpha+2)\left(0^{2\alpha+2}\nu + 0^{\alpha+1}\right)\right)}{(\alpha+1)(\alpha+2)}$
Beta(2,2)	$\dfrac{\int_0^\theta u(1-u)du}{\int_0^1 u(1-u)du}$	$\theta_i^3(3\theta_i((2\theta_i(5((7\theta_i - 20)\theta_i + 14)\theta_i + 14) - 35)\nu - 35) + 140)/70$
Quadratic-U	$4(\theta - {}^1/_2)^3 + 4({}^1/_2)^3$	$\theta_i^2((\theta_i(8(\theta_i(3(7\theta_i - 20)\theta_i + 70) - 42)\theta_i + 105) - 14)\theta_i\nu + 14(3\theta_i - 4)\theta_i + 21)/14$

Table 1: Two-player equilibrium bidding functions.

3.2 Two asymmetric bidders

This section presents our results on all-pay auctions between bidders who are not identical in terms of their risk preferences. Consider the following uniform, two-players setup featuring asymmetric degrees of variance aversion ν_i where player $i \in \{1,2\}$ maximises

$$u_i(\mu(b,\theta_i),\sigma^2(b,\theta_i)) = \mu(b,\theta_i) - \nu_i\sigma^2(b,\theta_i),\ \ \nu_i \in \mathbb{R}_+. \tag{14}$$

We consider the particular case of $\nu_1 = 0$ and $\nu_2 = \nu$, i.e., bidder 1 is risk-neutral while bidder 2 is variance-averse. Therefore, we get

$$\begin{aligned} u_1(\theta_1,b_1) &= \beta_2^{-1}(b_1)\theta_1 - b_1, \\ u_2(\theta_2,b_2) &= \beta_1^{-1}(b_2)\theta_2 - \nu\left(\theta_2^2\beta_1^{-1}(b_2)(1-\beta_1^{-1}(b_1))\right) - b_2 \end{aligned} \tag{15}$$

with the pair of first-order conditions

$$\begin{aligned} \frac{\partial u_1(\theta_1,b_1)}{\partial b_1} &= \frac{1}{\beta_2'(\beta_2^{-1}(b_1))}\theta_1 - 1 = 0 \\ &\Leftrightarrow\ \beta_2'(\beta_2^{-1}(b_1)) = \theta_1, \end{aligned} \tag{16}$$

$$\begin{aligned} \frac{\partial u_2(\theta_2,b_2)}{\partial b_2} &= \frac{1}{\beta_1'(\beta_1^{-1}(b_2))}\left(\theta_2 - \nu\theta_2^2(1-2\beta_1^{-1}(b_2))\right) - 1 = 0 \\ &\Leftrightarrow\ \beta_1'(\beta_1^{-1}(b_2)) - 2\nu\theta_2^2\beta_1^{-1}(b_2) = \theta_2 - \nu\theta_2^2. \end{aligned} \tag{17}$$

In equilibrium, $b_1 = \beta_1(\theta_1)$ and $b_2 = \beta_2(\theta_2)$. Thus, we substitute $\beta_1^{-1}(b_1) = \theta_1$ into (16) to obtain

$$\beta_2'(\beta_2^{-1}(b)) = \beta_1^{-1}(b). \tag{18}$$

Taking the derivative of $\beta_1^{-1}(b)$ and applying (16) gives

$$\beta_1'(\beta_1^{-1}(b)) = \frac{\beta_2'(\beta_2^{-1}(b))}{\beta_2''(\beta_2^{-1}(b))} \tag{19}$$

where we use b as variable from the joint support of $\beta_1(\cdot)$ and $\beta_2(\cdot)$.[15] Substituting (19) and (18) into (17) yields the following second-order differential equation in β_2

$$\frac{\beta_2'(\beta_2^{-1}(b))}{\beta_2''(\beta_2^{-1}(b))} - 2\nu\theta^2\beta_2'(\beta_2^{-1}(b)) = \theta - \nu\theta^2. \tag{20}$$

This differential equation can be solved using the boundary condition $\beta_2(0) = 0$ to obtain

$$\begin{aligned} \beta_2(\theta_2) &= \frac{1}{2\sqrt{1+4c\nu}}\Big[\sqrt{1+4c}\left(-1+\theta\nu+\sqrt{1+\theta\nu(-2+\theta\nu+4c\theta\nu)}+\log(2)\right) \\ &+ \log\left(1-\sqrt{1+4c}\right) - \sqrt{1+4c}\log\left(1-\theta\nu+\sqrt{1+\theta\nu(-2+\theta\nu+4c\theta\nu)}\right) \\ &- \log\left(1-\nu\left(\theta+4c\theta+\sqrt{1+4c}\sqrt{\tfrac{1}{\nu^2}+\tfrac{\theta}{\nu}(-2+\theta\nu+4c\theta\nu)}\right)\right)\Big] \end{aligned} \tag{21}$$

for yet undetermined constant of integration c. In order to solve for the first player's bidding function, we solve (18) for

$$\beta_1(\theta) = \beta_2\left((\beta_2')^{-1}(\theta)\right) \tag{22}$$

[15] The standard argument applies that in the two-player, all-pay auction the supports of both players' bidding functions coincide.

where

$$(\beta_2')^{-1}(\theta) = \frac{\theta}{\nu(\theta - \theta^2 + c)}. \tag{23}$$

From an argument similar to the one used in a standard (risk-neutral) all-pay auction follows that the two bidding functions $\beta_1(\cdot)$ and $\beta_2(\cdot)$ must share the same support. Intuitively, in equilibrium, no player's type can submit a strictly higher bid than the other player's highest type as such a bid does not affect the payoff variance (the player wins with certainty), but strictly lowers the expected payoff. Setting $\beta_1(1) = \beta_2(1)$ implies that the only possible value for the constant of integration is

$$c = \frac{1}{\nu}. \tag{24}$$

Substituting this constant into (21), we obtain the following pair of bidding functions

$$\begin{aligned}
\beta_1(\theta) &= \frac{\log(1-\theta^2\nu+\theta\nu)}{2\nu} - \frac{\theta^2}{\theta\nu(\theta-1)-1} \\
&+ \frac{1}{2\sqrt{\nu(4+\nu)}}\left(\log\left(1-\sqrt{\tfrac{4+\nu}{\nu}}\right) - \log\left(\frac{\sqrt{\frac{4+\nu}{\nu}}+\theta\left(4+\theta\left(\nu+\sqrt{\nu(4+\nu)}\right)\right)-1}{\theta\nu(\theta-1)-1}\right)\right), \\
\beta_2(\theta) &= \frac{1}{2\nu}\left[\theta\nu - 1 + \sqrt{1-2\theta\nu+\theta^2\nu(4+\nu)} + \sqrt{\tfrac{\nu}{4+\nu}}\log\left(\sqrt{\tfrac{4+\nu}{\nu}}-1\right)\right. \\
&- \log\left(\tfrac{1}{2}\left(1-\theta\nu+\sqrt{1-2\theta\nu+\theta^2\nu(4+\nu)}\right)\right) \\
&- \left.\sqrt{\tfrac{\nu}{4+\nu}}\log\left(\theta(4+\nu)-1+\sqrt{\tfrac{(4+\nu)(1-2\theta\nu+\theta^2\nu(4+\nu))}{\nu}}\right)\right],
\end{aligned} \tag{25}$$

which are illustrated for the case of $\nu = 1/2$ in figure 3.

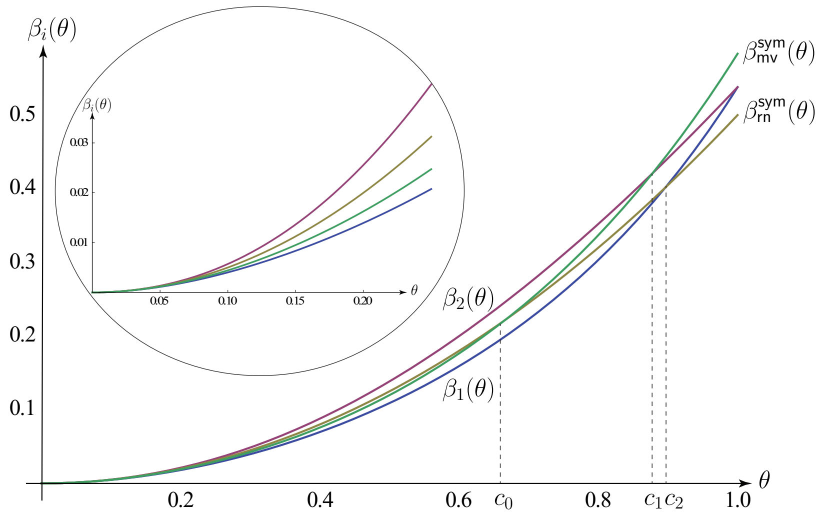

Figure 3: Comparison of asymmetric and symmetric bidding under mean-variance preferences for $\nu = 1/2$.

As the figure shows, each positive risk-neutral player type bids less than the corresponding type of her variance-averse opponent. While the ν-variance-averse, asymmetric bidder with bidding function $\beta_2(\cdot)$ always bids more than symmetric risk-neutral bidders $\beta^{\mathsf{sym}}_{\mathsf{rn}}$, the asymmetric risk-neutral bidder with bidding function $\beta_1(\cdot)$ (competing with a variance-averse player) bids up to a cutoff-type c_2 below the symmetric risk-neutral bidders and, for types higher than c_2, she bids above. Similarly, the asymmetric variance-averse bidder (competing with a risk-neutral bidder) bids up to a cutoff-type c_1 above the symmetric ν-variance-averse bidders ($\beta^{\mathsf{sym}}_{\mathsf{mv}}$) and bids below for types higher than c_1. Both properties are qualitatively similar to the single-crossing property with cutoff $\hat{\theta} = c_0$ from proposition 1. The generally high bids of the variance-averse bidder cause low types of the risk-neutral bidder to bid less in comparison to their strategy when faced with risk-neutral opponents. High types of the risk-neutral bidder, on the other hand, increase their bid in reaction to their variance-averse opponent's strategy.

4 Revenue valuation

The classical reference for revenue valuation in winner-pay auctions under risk aversion is Holt (1980) who discusses a procurement setup. Revenue equivalence between the standard auction formats breaks down with risk averse bidders. While second-price bidders maintain their dominant strategies of bidding their values, first-price competitors increase their bids with respect to the standard, risk-neutral case. This is due to the fact that raising one's bid in a first-price auction can be seen as partial insurance against loosing. From a risk averse seller's point of view, the first-price auction is preferable to a second-price format because it exposes the seller to less revenue risk.[16]

In this section we limit attention to uniformly distributed bidder valuations because our objective lies in the derivation of a series of concrete revenue ranking results. The results, however, are qualitatively similar for the other distributions listed in the table of section 3.1.1. The seller's expected revenue R depends on the bidder's preferences. In the case of risk-neutral bidders with von Neumann-Morgenstern preferences, the seller expects to earn

$$\mathbb{E}[\mathsf{R}_{\mathsf{rn}}] = n \int_0^1 \frac{n-1}{n}\theta^n \, d\theta = \frac{n-1}{n+1}. \tag{26}$$

If the bidders exhibit mean-variance preferences, then the seller can expect

$$\begin{aligned} \mathbb{E}[\mathsf{R}_{\mathsf{mv}}] &= n \int_0^1 \frac{n-1}{n}\theta^n + \nu \left(\frac{n-1}{n}\theta^{2n} - \frac{n-1}{n+1}\theta^{n+1} \right) d\theta \\ &= \frac{n-1}{n+1} + \nu \left(\frac{n-1}{2n+1} - \frac{n(n-1)}{(n+1)(n+2)} \right). \end{aligned} \tag{27}$$

The revenue limit for $n \to \infty$ is $1 - \nu/2$.

Holding the number of players fixed, expected revenue is strictly increasing in ν for the two-players case and strictly decreasing in ν for all $n > 2$. This is illustrated in figure 4 and summarised in corollary 2.

Corollary 2. *For $n = 2$, the expected revenue is strictly greater if bidders exhibit mean-variance preferences than if they are expected payoff maximisers. For all other n, this relationship is reversed.*

A consequence of this corollary is that the auctioneer's revenue in the (first-price) all-pay auction dominates that of the second-price auction for two bidders but is strictly lower for all other numbers of auction participants. The reason is that the dominant-strategy equilibrium revenue of a second-price auction is invariant with respect to the bidders' variance aversion ν and, hence, second-price

[16] For references, see Milgrom (2004, p123).

revenue is given by (26). To see this more formally, consider a bidder's mean-variance optimisation problem in the second-price auction

$$\max_b \int_0^{\beta^{-1}(b)} (\theta_i - \beta(\theta_j))\, d\theta_j - \nu \left(\int_0^{\beta^{-1}(b)} ((\theta_i - \beta(\theta_j)) - \mu)^2\, d\theta_j + \int_{\beta^{-1}(b)}^1 (0 - \mu)^2\, d\theta_j \right). \quad (28)$$

Supplying the candidate $\beta(\theta_i) = \theta_i$ and $\beta^{-1}(b) = b$ gives the first-order condition

$$(b - \theta_i)\left(\nu^2(b(b - 2\theta_i - 1) + \theta_i) - 1\right) = 0 \quad (29)$$

in which the first term is solved by the equilibrium $\beta(\theta_i) = b = \theta_i$. (The second term only gives solutions outside of [0,1].) This argument can be made independent of the distribution of bidder valuations.

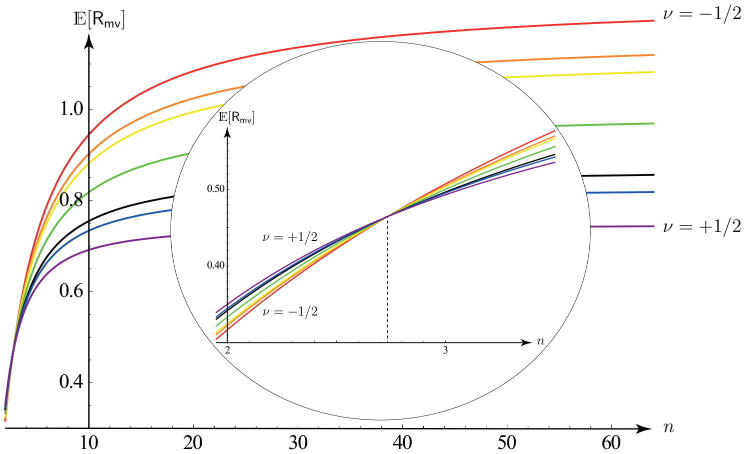

Figure 4: Revenue for $\nu \in \{-\frac{1}{2}, -\frac{1}{3}, -\frac{1}{4}, 0, \frac{1}{4}, \frac{1}{3}, \frac{1}{2}\}$ players, respectively (sorted in the colours of the rainbow from red to violet) and $\nu = 0$ in black for $n \in [2, 64]$. The analysis of variance seeking behaviour (i.e., negative ν) is presented in section 5.2.

If the seller himself also considers the revenue variance in addition to the revenue's mean, then his preference may be reversed. In the case of risk-neutral bidding, the seller's revenue variance is

$$\mathbb{V}[R_{rn}] = n \int_0^1 \left(\beta_{rn}(\theta) - \frac{\mathbb{E}[R_{rn}]}{n} \right)^2 d\theta = \frac{n(n-1)^2}{(2n+1)(n+1)^2}. \quad (30)$$

The bidders behaving according to mean-variance preferences cause a revenue variance of

$$\begin{aligned} \mathbb{V}[R_{mv}] &= n \int_0^1 \left(\beta_{mv}(\theta) - \frac{\mathbb{E}[R_{mv}]}{n} \right)^2 d\theta \\ &= \frac{n(n-1)^2}{(1+2n)^2} \left(\frac{1+2n}{(1+n)^2} + \nu \frac{7 + n(21 - 4n(n-3))}{(1+n)^2(2+n)(1+3n)} \right. \\ &\qquad \left. + \nu^2 \frac{74 + n(151 + 8n(n-3)(n-1))}{(2+n)^2(3+2n)(2+3n)(1+4n)} \right). \end{aligned} \quad (31)$$

The rate of $\mathbb{V}[R_{mv}]/\mathbb{V}[R_{rn}]$ is given by

$$1 + \frac{(7 + (7-2n)n)\nu}{(2+n)(1+3n)} + \frac{(1+n)^2(74 + n(151 + 8(n-3)(n-1)n))\nu^2}{(2+n)^2(1+2n)(3+2n)(2+3n)(1+4n)}. \quad (32)$$

For $n = 2, 3, 4$, this ratio is greater than one and increasing in ν. For $n \geq 5$ the variance ratio is decreasing in ν and below one. This implies the following corollary.

		$n = 2$	$n = 3, 4$	$n \geq 5$
Corollary 3.	$\mathbb{E}[R]$	$\mathbb{E}[R_{mv}] > \mathbb{E}[R_{rn}]$	$\mathbb{E}[R_{mv}] < \mathbb{E}[R_{rn}]$	$\mathbb{E}[R_{mv}] < \mathbb{E}[R_{rn}]$
	$\mathbb{V}[R]$	$\mathbb{V}[R_{mv}] > \mathbb{V}[R_{rn}]$	$\mathbb{V}[R_{mv}] > \mathbb{V}[R_{rn}]$	$\mathbb{V}[R_{mv}] < \mathbb{V}[R_{rn}]$

Therefore, for $n = 3, 4$, a seller with both types of preferences will prefer bidders maximising expected payoff. In all other cases, a variance-averse seller may prefer bidders with mean-variance preferences, where the exact ranking depends on the degree of the seller's variance aversion.

4.1 Optimal reserve prices

Corollary 2 shows that with only two bidders the expected revenue of an all-pay auction is strictly increasing in the degree of variance aversion, although the opposite relationship is true for any other number of players. We now introduce an exogenous reserve price $p_r > 0$ into our revenue analysis to show that this ordering can be reversed by choosing an optimal reserve price.[17]

In the symmetric equilibrium, either one of two symmetric, variance-averse players will participate in the auction if their utility at bidding p_r equals

$$u_i(\mu(b = p_r, \theta_i), \sigma^2(b = p_r, \theta_i)) = \mu(p_r, \theta_i) - \nu\sigma^2(p_r, \theta_i) = 0, \tag{33}$$

where

$$\mu(p_r, \theta_i) = \theta_i\theta_i - p_r, \text{ and } \sigma^2(p_r, \theta_i) = \theta_i(\theta_i(1 - \theta_i))^2 + (1 - \theta_i)(\theta_i\theta_i)^2. \tag{34}$$

The first participating type in the contest with reserve price p_r, θ_r who solves (33) is implicitly defined by

$$p_r = \theta_r^2 + (\theta_r - 1)\theta_r^3\nu. \tag{35}$$

The solution to this equation, $\theta_r^{-1}(p_r)$, gives this type as a function of the reserve price.[18] Only if a player's type θ_i is at least as high as $\theta_r^{-1}(p_r)$, she will participate in the auction. Her maximisation problem (39) gives the bidding function $\beta_r(\theta)$ as equivalent of (42) as solution to

$$\beta_r(\theta) = \int_{\theta_r^{-1}(p_r)}^{\theta} \vartheta(1 - \nu\vartheta + 2\nu\vartheta^2)\, d\vartheta + p_r. \tag{36}$$

The seller's expected revenue when setting reserve price p_r is now

$$\mathbb{E}[\mathsf{R}_{\mathsf{mv}}(p_r)] = 2\int_{\theta_r^{-1}(p_r)}^{1} \beta_r(\vartheta)\, d\vartheta; \tag{37}$$

it is shown for various ν in figure 5.

As evident from the figure, the revenue-maximising reserve price p_r^* is lower when players are more variance-averse.

Corollary 4. *The highest expected revenue achievable by optimally setting a reserve price in a symmetric two-player all-pay auction is decreasing in the degree of variance aversion. The revenue maximizing reserve prize, p_r^*, is decreasing in the degree of bidders' variance aversion, ν.*

[17] Whether sellers actually set reserve prices optimally is debatable, Davis, Katok, and Kwasnica (2011) investigate this question in the context of winner-pay auctions in a laboratory experiment and find amongst other possible explanations that risk-aversion can explain parts of the observed data.

[18] As the explicit form of $\theta_r^{-1}(p_r)$, (59), is rather unappealing it is relegated to the appendix (as are all following expressions which are based on it).

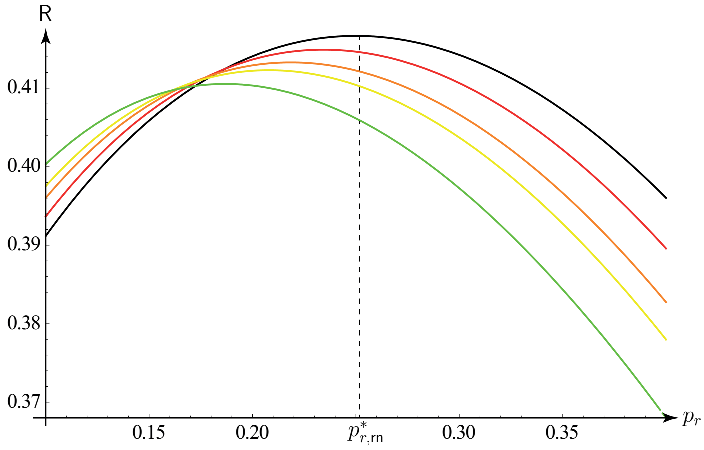

Figure 5: Seller's expected revenue (in the colours of the rainbow from red $\nu \in \{1/2, 1/3, 1/4, 1/8\}$ green) as a function of the reserve price p_r under variance-averse bidding compared to the case of risk-neutral bidding $\nu = 0$ (black).

5 Generalisations

5.1 Exogenous shocks

Another way to influence investments more subtly—and also in environments where it is not feasible to set reserve prices or restrict the number of contest participants—is via the controlled disclosure of information. To motivate the analysis of exogenous noise we return to the stylised example from the introduction. Consider an R&D company engaging in costly research outlays in order to obtain some (patentable) innovation first among a group of competitors. The endogenous variance is grounded, as before, in the uncertain spread between certain outlays and probabilistic winning. The exogenous component may be seen as market uncertainty in case of winning: the competitors in the smartwatch war cannot be entirely certain about the market perception and success of their future products. In fact, firms often invest in market research long before launching new products in order to get a more precise signal of future success in new product markets. For the purpose of analysing the effect of the precision of signals, we extend the basic model analysed in section 3 with exogenous noise. In a similar fashion to Esö and White (2004) we consider the possibility that the final value of the object may further be influenced by an exogenous shock, $\varepsilon \sim W(0, \hat{\varepsilon}^2)$, which is distributed over some compact interval with mean zero and variance $\hat{\varepsilon}^2$. Similarly, a player's valuation of the state in which she does not win the object may be subject to another independently distributed exogenous shock $\delta \sim L(0, \hat{\delta}^2)$.[19]

After realising their own (expected) valuations of the object, θ_i, all players simultaneously submit their bids, b_i, $i \in \mathcal{N}$. The player with the highest bid receives the object and all players forgo their bids. Consider expected valuations distributed $\mathbb{E}[\mathrm{R}] \sim W[\theta_i, \hat{\varepsilon}^2]$, where the distribution W is

[19] As an example of an interaction in which a discrimination between shocks is natural consider patent races where firms are already in the market. Here, the success of one firm certainly affects the market prospects of the losers.

elliptical, i.e., completely determined by mean θ_i and variance $\hat{\varepsilon}^2 \in [0,1]$.[20] Similarly, we allow for the case that a player's valuation, if she does not win the object, is subject to another exogenous shock $\delta \sim L(0, \hat{\delta}^2)$. In the case of R&D competition, $\hat{\delta}^2$ reflects the uncertainty in the company's forecast of the residual demand after the innovation. The exogenous shocks realise after the auction stage has ended and player i's payoff is given by

$$\pi_i(b_i, b_{-i}; \theta_i) = \begin{cases} \theta_i + \varepsilon - b_i & \text{if } b_i > b_j \forall j \neq i, \\ \frac{1}{m}(\theta_i + \varepsilon) + \frac{m-1}{m}\delta - b_i & \text{if } i \in Q = \{j \in \mathcal{N} | b_j = \max_{k \in \mathcal{N}} b_k\}, m = |Q|, \\ \delta - b_i & \text{if } \exists j : b_i < b_j. \end{cases}$$

Notice that neither of these exogenous noise variables have any effects on equilibrium bidding behaviour in the standard model of buyers with risk-neutral von Neumann-Morgenstern utility, who simply maximise expected payoffs. In the following we discuss how bidding strategies of buyers with mean-variance preferences are affected by noisy signals. The player's objective changes from (1) into

$$u_i(\mu(b, \theta_i), \sigma^2(b, \theta_i)) = \mu(b, \theta_i) - \nu \left(\sigma^2(b, \theta_i) + \hat{\varepsilon}^2 F(\beta^{-1}(b))^{n-1} + \hat{\delta}^2 (1 - F(\beta^{-1}(b))^{n-1}) \right) \quad (38)$$

for $\nu \in [0, \frac{1}{2}]$. Inserting back the expressions for the mean and variance (4), her objective is

$$\begin{aligned} u_i(b, \theta_i) = (F(\beta^{-1}(b)))^{n-1} &\left[\theta_i \left(1 - \nu\theta_i + \nu\theta_i (F(\beta^{-1}(b)))^{n-1}\right) - \nu\hat{\varepsilon}^2 \right) \right] \\ &+ (1 - (F(\beta^{-1}(b))))^{n-1} \hat{\delta}^2 - b. \end{aligned} \quad (39)$$

The first-order condition for maximisation of this expression with respect to b is

$$\begin{aligned} &\left[\theta_i \left(1 - \nu\theta_i + 2\nu\theta_i (F(\beta^{-1}(b)))^{n-1}\right) - \nu(\hat{\varepsilon}^2 - \hat{\delta}^2) \right] \times \\ &(n-1)(F(\beta^{-1}(b)))^{n-2} f(\beta^{-1}(b)) \frac{\partial \beta^{-1}(b)}{\partial b} = 1. \end{aligned} \quad (40)$$

In the symmetric equilibrium $b = \beta(\theta_i)$ for all $i \in \mathcal{N}$, this yields the first-order differential equation

$$\begin{aligned} \beta'(\theta_i) = \theta_i &\left(1 - \nu\theta_i + 2\nu\theta_i (F(\theta_i))^{n-1}\right) (n-1)(F(\theta_i))^{n-2} f(\theta_i) \\ &- (n-1)(F(\theta_i))^{n-2} f(\theta_i) \nu(\hat{\varepsilon}^2 - \hat{\delta}^2). \end{aligned} \quad (41)$$

This differential equation is solved by the bidding function[21]

$$\hat{\beta}_{\mathsf{mv}}(\theta_i) = \begin{cases} 0 & \text{if } \theta_i \leq \theta_0 \\ \beta_{\mathsf{mv}}(\theta_i) - (F(\theta_i))^{n-1} \nu(\hat{\varepsilon}^2 - \hat{\delta}^2) & \text{if } \theta_i > \theta_0 \end{cases} \quad (42)$$

[20] Elliptical distributions are a generalisation of the normal family containing, among others, the uniform, Student-t, Logistic, Laplace, symmetric stable, and Normal distributions. A detailed presentation of these distributions is available in Fang, Kotz, and Ng (1987).

[21] Note that the bidding function (42) constitutes an equilibrium in both cases, whether a zero bid is interpreted as abstaining from the contest and hence results in a winning probability of zero, or the winner is determined through efficient tie-breaking (i.e., ties are broken in favour of the player with the highest valuation for the prize) in the event that all players bid zero, which happens with strictly positive probability if $\hat{\varepsilon}^2 - \hat{\delta}^2 > 0$ and F has full support $[0, 1]$. Which of these cases is more appropriate depends on the exact environment to be modelled.

where $\beta_{\mathsf{mv}}(\theta_i)$ is defined in (10) and the 'cutoff type' θ_0 is implicitly defined as the solution to

$$\begin{aligned}&\theta_0(F(\theta_0))^{n-1} - \int_0^{\theta_0} (F(\vartheta))^{n-1} d\vartheta \\ &-\nu \int_0^{\theta_0} \vartheta^2(n-1)(F(\vartheta))^{n-2} f(\vartheta) \left(1 - 2(F(\vartheta))^{n-1}\right) d\vartheta - (F(\theta_0))^{n-1}\nu(\hat{\varepsilon}^2 - \hat{\delta}^2) = 0\end{aligned} \tag{43}$$

for which a closed form solution is generally unavailable. As we restrict bids to be non-negative, the resulting bidding function is still invertible over the relevant region. Similarly to the common practice of normalising the valuation of the outside option to zero, (42) shows that there is a degree of freedom to normalise the variance of one of the two possible outcomes. In the remainder we therefore normalise $\hat{\delta}^2 \equiv 0$ for simplicity.

Corollary 5. *Introducing exogenous noise $\hat{\varepsilon}^2 > 0$ on the prize rotates the optimal bidding schedule down, causing low-type bidders to abstain from participating in the auction.*

Consequently, it lies in the interest of an effort maximizing contest designer to minimise $\hat{\varepsilon}^2 - \hat{\delta}^2$, i.e., reveal as much information about the prize as feasible, while possibly keeping the loser's payoff uncertain.

5.1.1 Example

We conclude this section with our usual uniform, two-bidders example. Consider the equilibrium bidding function

$$b^* = \hat{\beta}(\theta_i) = \frac{\theta_i^2}{2} - \nu \left(\frac{\theta_i^3}{3} - \frac{\theta_i^4}{2} + \theta_i \hat{\varepsilon}^2 \right). \tag{44}$$

The bidding behaviour this suggests for the case of a stochastic prize parameterised by $\hat{\varepsilon}^2 = {}^1\!/_4$, is shown as dotted line in figure 2. Consider now a case in which we auction two objects valued $\theta_1 > \theta_2$ with exogenous prize variance $\hat{\varepsilon}^1 > \hat{\varepsilon}^2$. If (full demand) bidders submit separate bidding functions for each object, then we can get cases where the bid for the high-value/high-risk object is below that of the low-value/low-risk object as exhibited by the dashed pair of bidding functions relative to their solid counterparts in figure 6.

5.2 Variance seeking

Up to this point of our analysis we have focussed on players who aim to avoid variance in their payoffs. Although risk aversion is a much more common assumption in economic studies, empirical studies have found risk seeking behavior in multiple environments including investment games and inventive activities (see e.g., Åstebro, 2003). In this section, we point out that the approach using mean-variance preferences is equally applicable when players are variance seeking and briefly discuss the effects of such preferences on our results in section 3. Variance seeking preferences are captured by a negative parameter, $\nu \in (-1/2, 0)$, in our basic model (1). Following the same analytical steps as before, we find that (10) still describes the symmetric equilibrium bidding functions when n players have symmetric mean-variance preferences with variance seeking parameter $\nu \in (-1/2, 0)$.

Proposition 2 shows that in the equilibrium of the symmetric all-pay auction with variance seeking players low types bid more and high types bid less than their counterparts in the symmetric all-pay auction with risk neutral players.

Proposition 2 (Single-crossing II)**.** *In an all-pay auction with n symmetric players, who each maximise* (1) *with $\nu \in (-1/2, 0)$,*

1. *the strategies*

$$\beta_{mv}(\theta_i) = \beta_{rn} - \nu \int_0^{\theta_i} \vartheta^2 (n-1)(F(\vartheta))^{n-2} f(\vartheta) \left(1 - 2(F(\vartheta))^{n-1}\right) d\vartheta \qquad (45)$$

   constitute the symmetric equilibrium, and

2. *either $\beta_{mv}(\theta) \geq \beta_{rn}(\theta)$ for all $\theta \in [0,1]$ or there exists a $\hat{\theta}$ in the support of F such that $\beta_{mv}(\theta) \geq \beta_{rn}(\theta)$ for $\theta \leq \hat{\theta}$, $\beta_{mv}(\hat{\theta}) = \beta_{rn}(\hat{\theta})$ and $\beta_{mv}(\theta) < \beta_{rn}(\theta)$ for $\theta > \hat{\theta}$.*

Proof. The proof is the same as in section 3. □

Figure 6 illustrates this observation graphically. It further shows that from a contest designer's point of view it is strongly beneficial to introduce even a small amount of exogenous noise when players have variance seeking preferences. In our model this effect can be achieved both by a higher degree of uncertainty about the value of the prize and by less uncertainty in the event of losing the contest.

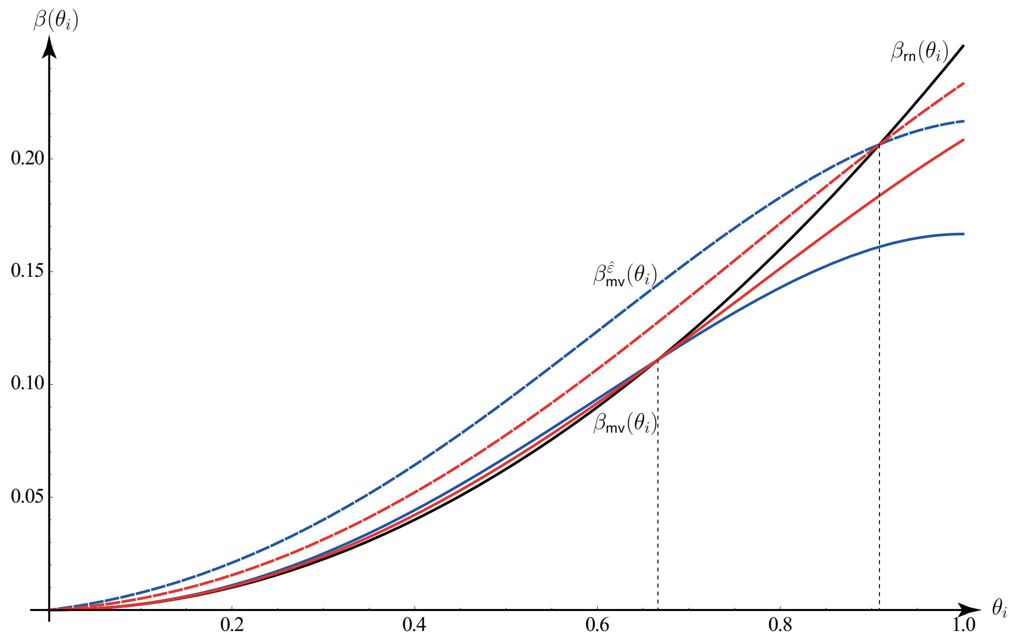

Figure 6: Equilibrium bidding functions for the all-pay auction under risk-neutrality (solid black, β_{rn}) and risk-seeking behaviour ($\nu \in \{-1/4, -1/2\}$, solid, red and blue, resp.). The dotted bidding functions exhibit risky prizes $\hat{\varepsilon} - \hat{\delta} = 1/10$.

5.3 Type-dependent variance preferences

In this section we further generalise our model to account for systematic differences in risk preferences. In section 3.2 we analyse bidding behaviour of two players who ex-ante differ in their risk attitudes. We take these variance aversion parameters as exogenously determined and assume that prize values are independently drawn from the same distribution. This model accurately captures bidding strategies when the mean-variance preference parameters are fixed. Contrary to this classical understanding of preferences, empirical and experimental studies suggest that the same player may indeed show different risk attitudes when facing different tasks and in particular when operating in different payoff spaces. In the spirit of Post and Levy (2005), who test a model of locally risk-seeking agents whose risk attitudes differ in bear and bull markets, we generalise our model such that risk preferences systematically differ with a player's valuation of the prize. In order to study this case, we change the simple objective (1) to

$$u_i(\mu(b,\theta_i),\sigma^2(b,\theta_i);\nu,\lambda,t) = \mu(b,\theta_i) - \nu(\theta_i - t)^\lambda \sigma^2(b,\theta_i), \tag{46}$$

for $\lambda \in \{2n+1 | n \in \mathbb{N}\} \cup \{0\}$ and reference point $t \in (0,1)$.[22] Note that setting $\lambda = 0$ takes us back to the case of type-independent mean-variance preferences as displayed in objective (1); in all other cases, for $\nu > 0$, types $\theta < t$ enjoy variance while types $\theta > t$ are variance averse. Type $\theta = t$ is risk neutral. A higher parameter λ represents more pronounced differences in variance preference towards the extremes, while a lower parameter λ corresponds to more uniformly increasing variance aversion. For the case where $\nu < 0$, this interpretation is reversed.

Taking the derivative with respect to player i's bid b, we obtain the first-order condition

$$\frac{2\theta_i^2\nu(n-1)\beta^{(-1)}(b)^{2n}(\theta_i-t)^\lambda - \beta^{(-1)}(b)^3\beta'\left(\beta^{(-1)}(b)\right) - \theta_i(n-1)\beta^{(-1)}(b)^{n+1}\left(\theta_i\nu(\theta_i-t)^\lambda - 1\right)}{\beta^{(-1)}(b)^3\beta'\left(\beta^{(-1)}(b)\right)} \tag{47}$$

in which we assume that the unknown function $\beta(\theta_i)$ is invertible. Setting zero and supplying the symmetric equilibrium condition $b = \beta(\theta_i)$, this simplifies to

$$\beta'(\theta_i) = (n-1)\theta_i^{n-1}\left(1 + (\theta_i - t)^\lambda \theta_i \nu \left(2\theta_i^{n-1} - 1\right)\right). \tag{48}$$

A type-zero player still optimally bids zero. The first-order differential equation (48) is then solved for boundary condition $\beta(0) = 0$ by

$$\beta_{\nu,\theta}(\theta_i;\nu,\lambda,t) = \frac{n-1}{n}\theta_i^n - \nu(n-1)t^{n+1}(-t)^\lambda \int_0^{\frac{\theta_i}{t}} x^n(1-x)^\lambda\left(1 - 2t^{n-1}x^{n-1}\right)dx. \tag{49}$$

This formulation allows for changes of variance-preferences based on endogenous, type-dependent considerations. Our main motivation in this section is to illustrate that the introduction of mean-variance considerations into the standard auction analysis allows for simple analytical predictions even in relatively complex environments in which solutions based on full von Neumann-Morgenstern preferences seem entirely out of reach.

[22] The limiting cases of $t \in \{0,1\}$ are already covered by the previous analysis.

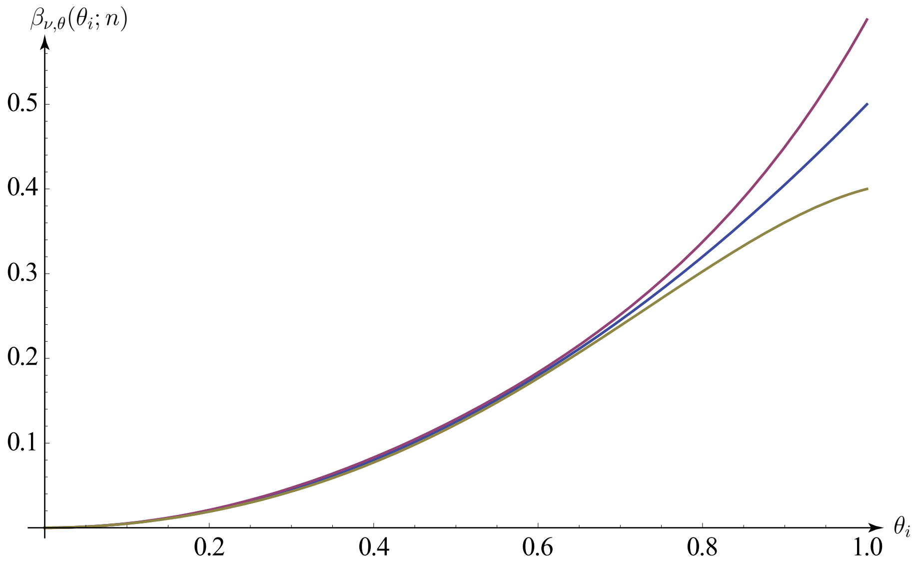

Figure 7: Type-dependent equilibrium bidding functions $\beta_{\nu,\theta}(\theta; 0, 0, 1/2)$ (blue),$\beta_{\nu,\theta}(\theta; 1/2, 1, 1/2)$ (red), $\beta_{\nu,\theta}(\theta; -1/2, 1, 1/2)$ (gold) for two players. The bidding functions do not intersect for $\theta > 0$.

6 Concluding remarks

We present first results for the study of all-pay auctions if buyers or sellers are endowed with mean-variance preferences. We fully characterise the symmetric equilibrium bidding functions of the all-pay auction with n identical bidders when bidders maximise an additively separable function of their expected payoff and payoff variance. Our first proposition shows that consideration of mean-variance preferences suffices to derive the qualitative properties of the bidding function which Fibich, Gavious, and Sela (2006) obtain in their analysis of a similar environment but considering any von Neumann-Morgenstern utility function which entails risk aversion. These qualitative properties seem to correspond well to experimental data. As such, all our results based on this bidding function appear relevant even when payoff distributions are not fully characterised by their first two moments. Eisenhuth and Ewers (2012) observe similar effects on equilibrium bidding behaviour in their analysis of loss-averse players. Most standard models of loss-aversion are equivalent to our model of mean-variance preferences because the wedge between the winning and losing payoff is exogenously determined by a player's type in the all-pay auction.

In our model of mean-variance preferences, players choose a strategy that maximises the difference between their expected payoff and the payoff variance, which is weighted by a parameter, ν, representing the players' degree of variance aversion. One major advantage of this approach is that we obtain closed form solutions for the bidding functions with just a single parameter representing risk aversion. Thus, we can perform comparative statics. Furthermore, this functional form allows us to relax the standard assumption of identical preferences. We exemplary solve for the bidding functions in an all-pay auction with one expected payoff maximiser and one bidder with mean-variance preferences. In contrast to the symmetric equilibrium, we find that the mean-variance bidder of a given type always bids more than her risk neutral opponent of the same type. Although the analysis is only provided for the case of two bidders, the result would look similar if more general sets of

n_1 risk neutral and n_2 mean-variance bidders were competing. Similarly, we conjecture that the qualitative findings from our benchmark case would carry over if the first bidder type was not risk neutral, but just less variance averse than her opponent.

Having obtained the (symmetric) equilibrium bidding function we then turn to the seller's perspective and consider effects of the number of bidders, their degree of variance aversion, and an optimally set reserve price. Corollary 2 shows that the influence of variance aversion on expected revenue depends on the number of players. In particular, we find that considering $n \geq 3$ reverses the ranking found for the two-player case. This finding suggests that under risk-aversion the generalisation from the two-player case to the general case may not always be as intuitive as it is often the case under risk neutrality.

With the exception of the analysis of bidding behaviour of n ex-ante identical players, much of our analysis focuses on the case of valuations that are i.i.d. draws from the uniform distribution over $[0, 1]$. The resulting simplification of otherwise lengthy expressions and the possibility to analytically obtain solutions has caused us to make this assumption. However, qualitatively similar results can be obtained for other standard distributions.

Appendix

A. Generalised mean-variance preferences

Throughout the paper we use simple preferences of the form

$$u_i(\mu(b, \theta_i), \sigma^2(b, \theta_i)) \; = \mu(b, \theta_i) - \nu\sigma^2(b, \theta_i), \text{ with } \nu \in [0, 1]. \tag{50}$$

In this appendix, we generalise this expression to the form mean plus a monomial function of the variance as follows

$$u_i(\mu(b, \theta_i), \sigma^2(b, \theta_i)) \; = \mu(b, \theta_i) - \nu\sigma^2(b, \theta_i)^\kappa, \text{ with } \kappa > 0 \tag{51}$$

allowing for both concave and convex transformations of the variance term. Going through the same steps as in the derivation of the original symmetric bidding function of subsection 3.1 we obtain, for two bidders with uniformly distributed valuations, the equilibrium bidding function

$$\beta(\theta) = \frac{\theta^2}{2} - \kappa\nu^{\kappa+1} \left[B(\theta; 3\kappa, \kappa) - 2B(\theta; 3\kappa + 1, \kappa)\right] \tag{52}$$

where $B(\theta; a, b)$ represents the incomplete Euler Beta function defined as

$$B(\theta; a, b) = \int_0^\theta t^{a-1}(1-t)^{b-1} dt. \tag{53}$$

This illustrates that, although we choose a particularly simple form of preferences for presentational purposes, nothing stops us from using more involved functional preference representations.

B. Second-order condition

The second-order condition is obtained by twice differentiating the objective (2) and supplying (9) for $\frac{\partial \beta^{-1}(b)}{\partial b}$ and $\frac{\partial^2 \beta^{-1}(b)}{\partial b^2} = -\frac{\beta''(\theta)}{\beta'(\theta)^3}$. The resulting expression simplifies to

$$F(\theta_i)^{5-2n} \frac{F(\theta_i)^2(\theta_i\nu - 2\theta_i\nu F(\theta_i) - 1)\left(1 - 2\theta_i\nu + 4\theta_i\nu F(\theta_i)^{n-1}\right) + \theta_i f(\theta_i) D}{(n-1)^2\theta_i^2 f(\theta_i)\left(F(\theta_i) - \theta_i\nu F(\theta_i) + 2\theta_i\nu F(\theta_i)^n\right)^3} \tag{54}$$

where

$$\begin{aligned} D = &(\theta_i\nu - 1)F(\theta_i)((n-2)(1-\theta_i\nu) + 2(n-3)\theta_i\nu F(\theta_i)) \\ &- 2\theta_i\nu F(\theta_i)^n((2n-3)(1-\theta_i\nu) + 4(n-2)\theta_i\nu F(\theta_i)), \end{aligned} \tag{55}$$

which, under the sufficient condition (7), is negative.

C. Explicit forms used in revenue derivation

This appendix shows the explicit form of some of the equations kept for presentation reasons from the main text. Define

$$A = \sqrt[3]{27p_r\nu^2 - 2 - 72p_r\nu + \sqrt{4(12p_r\nu - 1)^3 + (2 + 9p_r(8-3\nu)\nu)^2}}, \tag{56}$$

$$B = \sqrt{3 - \frac{8}{\nu} + \frac{4\ 2^{1/3}(12p_r\nu - 1)}{\nu A} - \frac{2\ 2^{2/3}A}{\nu}}, \tag{57}$$

and

$$C = \sqrt{3}B - 3 - \sqrt{6}\sqrt{3 - \frac{8}{\nu} + \frac{2\ 2^{1/3}(1 - 12p_r\nu)}{\nu A} + \frac{2^{2/3}A}{\nu} - \frac{3\sqrt{3}(\nu - 4)}{\nu B}}. \tag{58}$$

Then the inverse of (35) is given explicitly as

$$\theta_r^{-1}(p_r) = \frac{3 - \sqrt{3}B + \sqrt{6}\sqrt{3 - \frac{8}{\nu} + \frac{2\ 2^{1/3}(1-12p_r\nu)}{\nu A} + \frac{2^{2/3}A}{\nu} - \frac{3\sqrt{3}(\nu-4)}{\nu B}}}{12}. \tag{59}$$

The explicit version of (36) is

$$\beta_r(\theta_i) = \frac{3p_r + \theta_i^2\left(3 - 2\nu\theta_i + 3\nu\theta_i^2\right) + \nu\left(\frac{C}{12}\right)^3}{6}. \tag{60}$$

Finally, the explicit version of the revenue of the all-pay auction with two symmetrically variance-averse players with parameter ν and reserve price p_r (37) is

$$\mathsf{Rev}_{\mathsf{mv}}(p_r) = \frac{10 + 39p_r + \nu + 3p_rC + (4+\nu)\left(\frac{C}{12}\right)^3 - \left(\frac{C}{4}\right)^2}{30}. \tag{61}$$

In principle, the derivative of the last expression with respect to p_r gives an explicit version of the revenue-optimal reserve price p_r^*. This derivation is not shown here for reasons of economy of space.

References

Amann, E., and W. Leininger (1996): "Asymmetric All-Pay Auctions with Incomplete Information: The Two-Player Case," *Games and Economic Behavior*, 14, 1–18.

Åstebro, T. B. (2003): "The Return to Independent Invention: Evidence of Unrealistic Optimism, Risk Seeking or Skewness Loving?," *Economic Journal*, 113, 226–39.

Bagnoli, M., and S. G. Watts (2000): "Chasing Hot Funds: The Effects of Relative Performance on Portfolio Choice," *Financial Management*, 29, 31–50.

Barut, Y., D. Kovenock, and C. Noussair (2002): "A Comparison of Multiple-Unit All-Pay and Winner-Pay Auctions Under Incomplete Information," *International Economic Review*, 43, 675–708.

Blavatskyy, P. R. (2010): "Modifying the Mean-Variance Approach to Avoid Violations of Stochastic Dominance," *Management Science*, 56(11), 2050–2057.

Chiu, H. (2010): "Consistent mean-variance preferences," *Oxford Economic Papers*, 63, 398–418.

Cingottini, I., and D. Menicucci (2006): "On the profitability of reducing competition in all-pay auctions with risk averse bidders," *Economic Letters*, 91, 260–6.

Cornes, R., and R. Hartley (2012a): "Loss aversion in contests," *University of Manchester, Economics Discussion Paper*, EDP-1204.

——— (2012b): "Risk aversion in symmetric and asymmetric contests," *Economic Theory*, 51, 247–75.

Coyle, B. T. (1992): "Risk Aversion and Price Risk in Duality Models of Production: A Linear Mean-Variance Approach," *American Journal of Agricultural Economics*, 74, 849–59.

Dasgupta, P. (1986): "The Theory of Technological Competition," in *New Developments in the Analysis of Market Structure*, ed. by J. E. Stiglitz, and G. F. Mathewson, pp. 519–49. MIT Press, Cambridge, Mass.

Davis, A. M., E. Katok, and A. M. Kwasnica (2011): "Do Auctioneers Pick Optimal Reserve Prices?," *Management Science*, 57(1), 177–192.

Eichner, T. (2008): "Mean variance vulnerability," *Management Science*, 54(3), 586–593.

Eichner, T., and A. Wagener (2009): "Multiple Risks and Mean-Variance Preferences," *Operations Research*, 57, 1142–54.

Eisenhuth, R., and M. Ewers (2012): "Auctions with Loss Averse Bidders," Mimeo.

Esö, P., and G. Futó (1999): "Auction design with a risk averse seller," *Economic Letters*, 65, 71–74.

Esö, P., and L. White (2004): "Precautionary bidding in auctions," *Econometrica*, 72, 77–92.

Fang, K., S. Kotz, and K. Ng (1987): *Symmetric Multivariate and Related Distributions*. London: Chapman & Hall.

Fibich, G., A. Gavious, and A. Sela (2004): "Revenue equivalence in asymmetric auctions," *Journal of Economic Theory*, 115, 309–21.

——— (2006): "All-pay auctions with risk-averse players," *International Journal of Game Theory*, 34, 583–99.

Fu, Q., and J. Lu (2012): "The optimal multi-stage contest," *Economic Theory*, 51, 351–82.

Gneezy, U., M. Niederle, and A. Rustichini (2003): "Performance in competitive environments: Gender differences," *Quarterly Journal of Economics*, 118(3), 1049–74.

Holt, C. A. (1980): "Competitive bidding for contracts under alternative auction procedures," *Journal of Political Economy*, 88, 433–45.

Hu, A., S. Matthews, and L. Zou (2010): "Risk Aversion and Optimal Reserve Prices in First and Second-Price Auctions," *Journal of Economic Theory*, 145, 1188–202.

Ireland, N. (2004): "Risk Aversion and Aggression in Tournaments," *University of Warwick*, Working Paper.

Kaplan, T., and S. Zamir (2012): "Asymmetric first-price auctions with uniform distributions: analytic solutions to the general case," *Economic Theory*, 50, 269–302.

Kirkegaard, R. (2012): "A Mechanism Design Approach to Ranking Asymmetric Auctions," *Econometrica*, 80, 2349–64.

Kroll, Y., H. Levy, and H. M. Markowitz (1984): "Mean-Variance Versus Direct Utility Maximization," *Journal of Finance*, 39, 47–61.

Lazear, E., and S. Rosen (1981): "Rank Order Tournaments as Optimal Labor Contracts," *Journal of Political Economy*, 89, 841–64.

Lizzeri, A., and N. Persico (2000): "Uniqueness and existence of equilibrium in auctions with a reserve price," *Games and Economic Behavior*, 30(1), 83–114.

Markowitz, H. (1952): "Portfolio selection," *The Journal of Finance*, 7(1), 77–91.

——— (2012): "Meanvariance approximations to expected utility," *European Journal of Operations Research*, forthcoming.

Maskin, E., and J. Riley (1984): "Optimal auctions with risk averse buyers," *Econometrica*, 52(6), 1473–1518.

——— (2000): "Asymmetric auctions," *Review of Economic Studies*, 67, 413–38.

Matthews, S. (1983): "Selling to risk averse buyers with unobservable tastes," *Journal of Economic Theory*, 30, 370–400.

——— (1987): "Comparing auctions for risk averse buyers: A buyer's point of view," *Econometrica*, 55, 633–46.

Mayer, J. (1987): "Two-Moment Decision Models and Expected Utility Maximization," *American Economic Review*, 77, 421–30.

Milgrom, P. (2004): *Putting Auction Theory to Work*. Cambridge University Press, Cambridge, U.K.

Milgrom, P., and J. Roberts (1982): "Predation, reputation, and entry deterrence," *Journal of Economic Theory*, 27, 280–312.

Noussair, C., and J. Silver (2006): "Behavior in all-pay auctions with incomplete information," *Games and Economic Behavior*, 55, 189–206.

Parreiras, S. O., and A. Rubinchik (2010): "Contests with three or more heterogeneous agents," *Games and Economic Behavior*, 68(2), 703–715.

Post, T., and H. Levy (2005): "Does Risk Seeking Drive Asset Prices? A Stochastic Dominance Analysis of Aggregate Investor Preferences," *Review of Financial Studies*, 18(3), 925–53.

Robson, A. (2012): "Contests Between Players With Mean-Variance Preferences," *Griffith Business School Discussion Paper*, 2012-07.

Rothschild, M., and J. Stiglitz (1970): "Increasing Risk: I. A Definition," *Journal of Economic Theory*, 2, 225–43.

Saha, A. (1997): "Risk preference estimation in the nonlinear mean standard deviation approach," *Economic Inquiry*, 35, 770–82.

Sargent, T., and S. Heller (1987): *Macroeconomic Theory*. Academic Press, New York, 2 edn.

Sinn, H.-W. (1983): *Economic Decisions under Uncertainty*. North-Holland Publishing Company, Amsterdam.

Smith, J. L., and D. Levin (1996): "Ranking auctions with risk-averse bidders," *Journal of Economic Theory*, 68, 549–61.

Tesler, L. (1966): "Cutthroat competition and the long purse," *Journal of Law and Economics*, 9, 265–9.

Tsiang, S. C. (1972): "The Rationale of the Mean-Standard Deviation Analysis, Skewness Preference, and the Demand for Money," *American Economic Review*, 62, 354–71.