

Overbidding and overspreading
in rent-seeking experiments:
Cost structure and prize allocation rules¹

Subhasish M. Chowdhury²
University of East Anglia

Roman M. Sheremeta³
Chapman University

Theodore L. Turocy⁴
University of East Anglia

January 28, 2014

¹We thank two anonymous referees and an Advisory Editor for valuable suggestions as well as the helpful comments of Klaus Abbink, Dan Levin, Phillip Reiss, Karl Wärneryd, participants at the 4th Maastricht Behavioral and Experimental Economics Symposium and seminar participants at the University of East Anglia, University of Gottingen, University of St Gallen, and Stockholm School of Economics. This research has been supported by a grant from the Centre for Behavioural and Experimental Social Science at the University of East Anglia. Any remaining errors are ours.

²Corresponding author, School of Economics, Centre for Behavioural and Experimental Social Science, and ESRC Centre for Competition Policy, University of East Anglia, Norwich NR4 7TJ, UK, Email: s.modak-chowdhury@uea.ac.uk

³Argyros School of Business and Economics, Chapman University, One University Drive, Orange, CA 92866, U.S.A.

⁴School of Economics, and Centre for Behavioural and Experimental Social Science, University of East Anglia, Norwich NR4 7TJ, UK

Abstract

We study experimentally the effects of cost structure and prize allocation rules on the performance of rent-seeking contests. Most previous studies use a lottery prize rule and linear cost, and find both overbidding relative to Nash equilibrium prediction and significant variation of efforts, which we term ‘overspreading.’ In a 2×2 design, we investigate the effects of allocating the prize by a lottery versus sharing it proportionally, and of a convex cost function versus linear costs of effort, while holding fixed the Nash equilibrium prediction for effort. We find that the share rule results in average effort closer to the Nash prediction and lower variation of effort. Combining the share rule with a convex cost function further enhances these results. Our findings indicate that a significant amount of subjects’ non-equilibrium behavior in contests can be explained by features of the experimental design. These results contribute towards design guidelines for contests based on behavioral principles that take into account implementation features of a contest that may not affect the Nash equilibrium prediction.

JEL Classifications: C72, C91, D72

Keywords: rent-seeking, contest, contest design, experiments, quantal response, overbidding

1 Introduction

Overbidding in rent-seeking contests (Tullock, 1980) is a robust phenomenon in the experimental literature. This phenomenon was first reported in experimental contest studies by Millner and Pratt (1989, 1991) and has since been replicated by many other experiments; for a comprehensive review see the survey of Dechenaux et al. (2012).¹ Average effort levels in contests generally exceed the Nash equilibrium prediction, in some cases by a wide enough margin that total expenditure by all contest participants exceeds the value of the prize. Moreover, contrary to the theoretical prediction of a unique pure strategy Nash equilibrium, experimental studies document that individual efforts are distributed on the entire strategy space, and individual behavior varies substantially across repeated plays of the game. We refer to this stylized fact as ‘overspreading.’

Over the last decade a number of studies have offered different explanations for overbidding and overspreading in rent-seeking contests.² Commonly-cited explanations for overbidding include noise and errors (Anderson et al., 1998; Shupp et al., 2013; Lim et al., 2014; Sheremeta, 2011); judgmental biases (Amaldoss and Rapoport, 2009; Sheremeta, 2011); a non-monetary utility for winning (Sheremeta, 2010; Price and Sheremeta, 2011); and evolutionarily stable behavior (Mago et al., 2012; Wärneryd, 2012). Overspreading is usually attributed to heterogeneity in subjects’ preferences towards losses (Kong, 2008), risk (Sheremeta, 2011), spitefulness (Herrman and Orzen, 2008), or winning (Sheremeta, 2010), as well as demographic differences (Price and Sheremeta, forthcoming).

In a standard lottery contest, all players exert effort in order to increase their probability of winning the prize. Higher efforts imply higher probability of winning, but they are also more costly. In equilibrium, the marginal benefit of effort is equal to the marginal cost. Therefore, a correct best-response computation requires experimental subjects to assess correctly marginal benefit, which depends on the probability of winning, and marginal cost, which depends on the convexity of the cost function. Any non-equilibrium behavior may thus simply come as a consequence of a difficult computational task (Wright, 1980; Simon, 1992; Rubenstein, 1998; Gigerenzer and Selten, 2001).

It has been well recognized, for example, that subjects may possess distorted perceptions of probabilities, which may lead to non-equilibrium behavior. As a consequence, many alternative theories have been proposed to account for such perceptions (Kahneman and Tversky, 1979; Quiggin, 1982; Chew, 1983; Tversky and Kahneman, 1992; Wilcox, 2011). Recent studies have tried to apply some of these theories to explain subject behavior in contests and auctions (Goeree et al., 2002; Baharad and Nitzan, 2008; Amaldoss and Rapoport, 2009). However, even after accounting for individual perceptions of probabilities, the aggregate patterns of overbidding and overspreading

¹Examples include Davis and Reilly (1998), Potters et al. (1998), Lim et al. (2014), Sheremeta (2010), Sheremeta (2011), and Sheremeta and Zhang (2010).

²The survey of Sheremeta (2013) discusses these in greater detail.

in lottery contests cannot be explained.

Another explanation for why subjects' behavior differs from theoretical predictions is based on flatness of payoff functions (Harrison, 1989; Goeree et al., 2002; Georganas et al., 2011). Harrison (1989), for instance, argues that overbidding in private value auctions relative to the Nash equilibrium may be due to the fact that the costs of such overbidding are rather small. By manipulating the cost of overbidding in the first-price and second-price winner pay auctions, Goeree et al. (2002) and Georganas et al. (2011) find support for this argument. Similarly, Müller and Schotter (2010) find that subjects overbid in private value all-pay auctions when the cost of bid function is linear but they actually underbid when the cost function is convex. The design of the experiment reported in the present study follows this approach by manipulating the relative costs of units of effort above versus below the equilibrium.

We examine whether the factors listed in the foregoing paragraphs can explain non-equilibrium behavior in contests, by manipulating design features of the context which do not affect the (risk-neutral) Nash equilibrium prediction. We consider four contest settings, organized in a 2×2 design. In one dimension, we vary whether the prize amount is indivisible and allocated stochastically, or whether it is shared proportionally; this manipulation speaks to hypotheses involving the salience of winning or limitations in reasoning about probability.³ In the other dimension, we vary whether the cost function is linear or convex in effort; the convex cost function induces an asymmetry in the amount of earnings foregone due to efforts in excess of the best response versus those foregone due to efforts less than the best response.

We find that in contests where the prize is shared proportionally, there is less overbidding and less overspreading. Average efforts are closer to the Nash equilibrium prediction, and there is lower variation in individual efforts. A convex cost function enhances these results under the share rule. However, we find that convex costs actually exacerbate overbidding with probabilistic allocation, which we attribute to knock-on effects driven by out-of-equilibrium play. These findings illustrate the importance of considering the behavioral drivers of out-of-equilibrium play for robust design of contests, and provide some first results for guidance of contest design along these lines.

2 Theoretical background

We study a rent-seeking contest game following Tullock (1980). There are N players, indexed by i . There is a prize, worth $V > 0$. Each player i simultaneously and independently chooses an effort level $e_i \in [0, V]$. In a standard abuse of notation, we will write $e_{-i} = \sum_{j \neq i} e_j$ to denote the

³Another avenue to limiting the role of chance in understanding outcomes is to compare the lottery mechanism with the all-pay auction, in which the participant with the highest effort wins with certainty. This has been done by, for example, Potters et al. (1998). Moving to the all-pay auction results in a qualitative change in the strategic structure of the game, as equilibria in common-value all-pay auctions generally involve randomization.

sum of efforts of other players. The cost of effort is given by a function $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which we assume to be continuous, monotonically increasing, and twice differentiable in effort. Irrespective of the outcome of the contest, all players forgo the cost of effort.

The contest success function is given by

$$p_i(e_i, e_{-i}) = \begin{cases} \frac{e_i}{e_i + e_{-i}} & \text{if } e_i + e_{-i} > 0 \\ \frac{1}{N} & \text{otherwise} \end{cases} \quad (1)$$

This function can be interpreted as either the probability that player i wins the prize, if it is allocated indivisibly to just one player, or the proportion of the prize awarded to player i . In either case, the expected payoff to player i can be written as

$$E(\pi_i) = p_i V - c(e_i). \quad (2)$$

Szidarovszky and Okuguchi (1997) show the existence and uniqueness of equilibrium of this game, where the equilibrium effort level e^* is given by the solution to the equation

$$c'(e^*)e^* = V \frac{(N-1)}{N^2}. \quad (3)$$

If the prize is awarded proportionally, then this pure-strategy equilibrium does not depend on the risk attitude of the players. If the prize is awarded using a lottery, then this is the pure-strategy equilibrium assuming all players are risk-neutral.⁴

3 Experimental design and procedures

We implemented a two-dimensional factorial design. In one dimension, we varied the contest success function, using the probabilistic (P) or share (S) rules for awarding the prize. In the second dimension, we varied the cost function, using the standard linear (L) cost function or a convex (C) cost function. We therefore had four treatments, labeled PL , PC , SL , and SC .

In each treatment, we conducted 3 independent sessions. Sessions consisted of 12 participants, who participated in 30 contests. Each session investigated only one treatment, so all comparisons

⁴Theoretically, if players are risk-averse, the direction of the change in equilibrium effort in contests is ambiguous. Hillman and Katz (1984) and Skaperdas and Gan (1995), for example, show that risk-averse players should exert lower effort. On the other hand, Konrad and Schlesinger (1997) and Cornes and Hartley (2003) show that the direction of the change in equilibrium effort caused by an increase in risk aversion may be ambiguous. Experimentally, however, there is more agreement, since most studies document that risk-averse subjects exert significantly lower effort than risk-neutral or risk-seeking subjects. (Millner and Pratt, 1991; Anderson and Freeborn, 2010; Sheremeta and Zhang, 2010; Sheremeta, 2011)

are across subjects. All sessions used student subjects at the Centre for Behavioural and Experimental Social Science at University of East Anglia. The computerized experimental sessions were run using z-Tree (Fischbacher, 2007).

Participants were matched into groups of $N = 4$, with random anonymous rematching after each contest. The value of the prize in all contests was $V = 80$ experimental francs. In each contest, participants simultaneously selected an effort level between 0 and 80. In sessions employing the linear cost function, the cost of effort was $c(e) = e$. In sessions with the convex cost function, the cost function $c(e) = \frac{e^2}{30}$ was used. Both cost functions lead to an equilibrium prediction for effort of $e^* = 15$. Subjects were informed about the group size, prize value, and the cost structure. After each period, a summary screen provided information about the one's own effort and the total effort of the group, the outcome of the contest (win or loss, or proportion of the prize), and one's own payoff in the period.

At the conclusion of the experiment, 5 of the 30 periods were chosen at random for payment. The earnings were converted into British pounds at the rate of 40 francs to £1. All subjects also received a participation fee of £15 to cover potential losses. Sessions lasted about an hour inclusive of instructions and payment. The average amount earned was £15.20, with individual earnings ranging from £2.00 to £20.40.

In both the linear and convex cost treatments, the choice variable for participants was the amount of effort, as opposed to the cost of effort. We made this decision to ensure comparability of the strategy space between the treatments. We conveyed information about the relationship between effort choices and costs by means of a table in the instructions⁵ enumerating each effort choice and its corresponding cost.⁶ The use of the table device permitted participants, in principle, to make decisions either based on the amount of effort to invest, or the amount of cost of effort. We therefore did not pursue a treatment in which cost rather than effort was framed as the decision variable.

4 Results

[Table 1 about here.]

Table 1 displays summary statistics on the distribution of efforts by treatment. Table 2 breaks out the mean and standard deviation of efforts by each session. In terms both of the aggregate point estimates, and in the ordering of sessions in Table 2, employing the share rule instead of the

⁵The instructions are available in Appendix A.

⁶In the linear cost sessions, the presence of this table may have seemed superfluous to participants, insofar as the effort and cost columns were identical. As a check on whether this device had any behavioral effect, we confirm in our analysis that we replicate results from a previous experiment.

probability rule decreases the average effort, both with linear and with convex costs, with the effect being much more pronounced in the convex cost case. The treatment effect of using the convex cost structure relative to linear is not as clear. With the share rule, convex costs decrease mean and median effort levels; however, with the probability rule, convex costs slightly increase the mean and median effort level, although as will be seen, the effect is not large enough to be statistically significant.

[Table 2 about here.]

The average costs of effort in Table 1 imply that on average participants were making losses in expectation in *PL*, *PC*, and *SL*. As this implies that there were opportunities in the experiment for participants to significantly increase their earnings, a next question is whether there is any evidence of a change in behavior based on experience over time in the experiment. Figure 1 displays the time series of average effort levels for each of the four treatments. There is a modest decrease in the average effort over time, but not a strong convergence towards the equilibrium prediction. The ordering of the treatments in terms of average efforts is stable throughout the experiment.

[Figure 1 about here.]

We formalize the analysis of treatment effects and time trends on average bidding levels by estimating a panel regression model

$$o_{it} = e_{it} - 15 = \beta_0 + \beta_1 S + \beta_2 C + \beta_3 t + u_i + \varepsilon_{it}. \quad (4)$$

The dependent variable, o_{it} , is the excess effort relative to the Nash equilibrium prediction. We include indicator variables for whether the share rule is employed, $S = 1$, and whether the convex cost structure is employed, $C = 1$. We use a random-effects error structure by subject, to account for the multiple decisions made by individual subjects. Standard errors are clustered at the session level to account for session effects.

[Table 3 about here.]

Table 3 reports the estimation results of the models. We report four specifications. Each specification focuses on comparing treatment pairs which differ only in one dimension. Each specification is estimated separately for the first 15 periods and last 15 periods, respectively. Based on these regressions, we can formally support the following results on average overbidding levels.

Result 1. *The share rule significantly reduces average overbidding when using convex costs. When using linear costs, overbidding is not significantly different under the share rule versus probability rule.*

Support. In specification (1) comparing PC and SC , average overbidding is lower by more than 11 in SC in both halves of sessions; this is significant at the 1% level. PC is roughly three times farther, in effort choice terms, from the equilibrium than SC . Turning to specification (2), average overbidding is lower in PL than SL , but the difference is not significant at the 5% level. \square

Result 2. *When using the share rule, overbidding is significantly lower with convex costs. When using the probability rule, there is no significant difference in overbidding between the linear cost and convex cost treatments.*

Support. In specification (4) comparing SC and SL , convex costs significantly decrease effort levels in both halves of the experiment. Specification (3) comparing PC and PL reveals the opposite pattern; the treatment effect of convex costs is not statistically significant, and moreover the sign of the point estimate changes to indicate convex costs increase average efforts. \square

Result 3. *There is evidence of learning and adjustment over time towards lower overbidding in all treatments at the aggregate level, with most of the adjustment occurring within the first half of the experiment.*

Support. In all four regression model specifications using all data, we find that in the first half of the experiment, the coefficient on the period number is negative in sign and statistically significant at the 1% level, with magnitudes ranging from -0.27 to -0.65. In the second half of the experiment, the point estimates are generally smaller in magnitude, and are not significantly different from zero at the 5% level. \square

We now turn to examining the variability of effort choices, which we term ‘overspreading,’ both within and across subjects. Figure 3 displays the histograms of effort levels in the second half of the experiment for each treatment. The histograms illustrate that focusing only on average levels misses out on much of the richness of the observed behavior. In addition, both factors in the experimental design have a qualitative effect on the distribution of effort choices. Subjects do respond both to changes in the allocation rule and changes in the cost structure, indicating both are taken into account in the subjects’ decision-making processes.

[Figure 2 about here.]

We now decompose the extent to which heterogeneity in effort choices arises from variation in how individual subjects behave, versus systematic differences across subjects. In view of the regressions in Table 3, by the second half of sessions, subjects have settled down into patterns of behavior that do not demonstrate a significant time trend. Whatever learning takes place about the

rules of the game and the possible behavior of others appears to be complete, at least in aggregate, by this time. Therefore, we consider choices in periods 16-30 only in the analysis below.⁷

To visualize the data in a way that captures both these sources of variation, Figure 2 presents a collection of boxplots, one for each subject, capturing the distribution of effort choices over the last half of the experiment.⁸ Subjects are ordered by increasing median effort choices, which are indicated by diamonds. Therefore, by focusing on the diamonds, one can read off the cumulative empirical distribution of the median choices across subjects, while by focusing on the boxplots themselves, one can get a sense of the degree of within-subject variation in behavior.

The boxplots suggest that the effects of the share rule and of convex costs on variability in efforts parallel the treatment effects on averages established earlier. We construct two measures to capture variation across subjects and within subjects. Let m_{ij} denote the median effort of subject j in session i .⁹ Then, our across-subject measure of variability in session i , V_i^A , is the standard deviation of m_{ij} over all subjects j in session i . Alternatively, let s_{ij} denote the standard deviation of effort of subject j in session i . Then, our within-subject measure of variability in session i , V_i^W , is the median of s_{ij} over all subjects j in session i . Table 2 reports these measures for each session.

As a starting point, the procedures and instructions used in the current study were adapted from those used by Sheremeta (2010). We therefore compare the results of the *PL* treatment to the results reported therein. The experiments differed in the subject pool (undergraduate students at Purdue University in the United States versus students at the University of East Anglia in the United Kingdom) and the number of experimental francs in the endowment and value of the prize (120 versus 80). The instructions in the current study differed only in the use of an explicit table summarizing the cost of each level of effort.

Result 4. *The PL treatment replicates the results of Sheremeta (2010) in terms of effort levels and variability of efforts both across subjects and within subjects.*

Support. We renormalize the effort levels in Sheremeta (2010) from the original $[0, 120]$ scale onto our $[0, 80]$ scale. We take the session as the unit of independent observation; Sheremeta (2010) conducted 6 sessions and our data contains 3. A Mann-Whitney-Wilcoxon (MWW) test on the null hypothesis that the median effort by session is equal between the studies cannot be rejected (p -value 0.36), and a MWW test on the null hypothesis that the standard deviation of effort by session is equal between the studies also cannot be rejected (p -value 0.20).

⁷Cutting the experiment in two at the midpoint here, and in the earlier regressions, is arbitrary. Moving the cut-point a few periods in either direction does not affect the conclusions.

⁸In these boxplots, the boxes cover the range from the lower quartile to the upper quartile of efforts. The “whiskers” indicate the adjacent values, as defined by Tukey (1977). The upper whisker is the highest observed effort which is no more than 1.5 times the interquartile range above the upper quartile; the lower whisker is defined analogously. Dots outside the whiskers indicate outliers.

⁹For these measures, we report statistics using data from the second half of the sessions.

Examining the underlying sources of variability in efforts, a MWW test on the null hypothesis that the across-subject measure V_i^A is equal between the studies cannot be rejected (p -value 0.20), and a MWW test on the null hypothesis that the within-subject measure V_i^W is equal between the studies cannot be rejected (p -value 0.44). \square

We now turn to the analysis of treatment effects within our experiment on across-subject and within-subject variability.

Result 5. *The share rule reduces variability both within subjects and across subjects.*

Support. Using a MWW test, we test the null hypothesis that the across-subject variability measure V_i^A is equal in sessions using the share rule and sessions using the probability rule. This null hypothesis is rejected (p -value 0.02). Also using a MWW test, we test the null hypothesis that the within-subject variability measure V_i^W is equal in sessions using the share rule and sessions using the probability rule. This null hypothesis is also rejected (p -value 0.03). \square

As with the results on average efforts, the effect of convex costs on variability in efforts is less clear-cut.

Result 6. *Point estimates indicate the use of convex costs tends to reduce variability both within subjects and across subjects, although the effect is not statistically significant. Sessions using both convex costs and the share rule tend to exhibit the lowest variability by both measures.*

Support. Using a MWW test, we test the null hypothesis that the across-subject variability measure V_i^A is equal in sessions using convex costs and those using linear costs. The sign of the MWW test statistic indicates variability is lower in the convex costs sessions, but the result is not significant at standard levels (p -value 0.11).

For the within-subject variability test, again we use a MWW test, with the null hypothesis that the within-subject measure V_i^W is equal in sessions using convex costs and those using linear costs. The sign of the test statistic indicates lower variability with convex costs, but again the test is not significant (p -value 0.34).

Treatment *SC* exhibits the lowest variability by both measures. The three *SC* sessions have the lowest across-subject variability, and the lowest, second-lowest, and fourth-lowest within-subject variability (with a *SL* session ranking third). Convex costs have a more clear-cut behavioral impact in the share rule setting, compared to the probability rule. \square

[Figure 3 about here.]

Nash equilibrium predictions rest crucially on the assumptions that players have correct beliefs about the distributions of strategy choices, and, given these beliefs, choose responses which

maximize their own payoffs. The non-degenerate distributions of effort levels observed in contests provide *prima facie* evidence that indeed subjects are not best-responding expected-earnings maximizers who have correct beliefs about the behavior of other participants in their cohort. While our rejection of point predictions of Nash equilibrium is neither novel nor surprising, our design allows us to look more closely at the Nash equilibrium assumptions to infer whether and how they fail.

In our design, subjects get feedback about the overall spending in their group in each period. Therefore, we begin by supposing that subjects have at least an approximate sense of the distribution of effort levels being chosen, therefore retaining the correct beliefs assumption, while relaxing the assumption of expected earnings maximization. One commonly-used model which captures these assumptions is the quantal response equilibrium (QRE) of McKelvey and Palfrey (1995). In a QRE, a player evaluates the expected payoff of each strategy choice inclusive of an additive noise term. We follow the standard in random-utility models by using the logit form of QRE. This form has one free parameter, $\lambda \in [0, \infty)$, which is a precision parameter; larger values of λ correspond to a smaller variance in the noise term in payoff evaluation.

To take logit QRE to the data, we again focus on the last 15 periods of the experiment, both because QRE assumes players have accurate beliefs about the play of others, and because we treat QRE as a static concept, so we avoid a confound with any early-period learning and adjustment. We estimate λ by maximum likelihood, pooling across all subjects across all sessions.¹⁰

Figure 3 displays the QRE fits superimposed over the histogram of choices, and reports the corresponding values of λ for each. Table 4 reports the fitted λ values and corresponding log-likelihoods for each treatment individually, as well as fits in which we estimate QRE fits restricting λ to be identical for pairs of treatments with a common factor, and finally a restriction with a common λ for all treatments combined.

We also construct a measure Q of the quality of the QRE fit obtained. QRE generates uniform randomization over all strategies when $\lambda = 0$; therefore, the worst possible log-likelihood that can result from a QRE fit is the log-likelihood of the data against the uniform distribution; call this log-likelihood $\ln L_u$. The best possible log-likelihood would occur if the distribution of choices were exactly a QRE distribution for some parameter λ ; call this log-likelihood $\ln L_m$. Then, we define Q as

$$Q = \frac{\ln L - \ln L_u}{\ln L_m - \ln L_u}. \quad (5)$$

Q will always be a number in the interval $[0, 1]$, with higher values of Q corresponding to better fits. Values for Q for the fits are also included in Table 4. We use Q as a convenient summary to

¹⁰As in all logit models, λ is cardinal, having units equal to the unit of measurement of payoffs. To facilitate comparison with other papers reporting logit QRE estimates, we express λ in units of US dollars, using the exchange rate $\pounds 1 = \$1.60$ in force at the time of the experiments.

compare the relative quality-of-fit for each treatment.

The values of Q and λ communicate different information. Q captures the extent to which the empirical distribution of the data match with the predictions of QRE; roughly, this indicates how good the QRE model is at organizing the overall features of the distribution. The value of λ quantifies the relationship between the frequency with which non-best-response efforts are played, and the earnings foregone due to that suboptimal play relative to the best response. Because the game has a unique Nash equilibrium which is in pure strategies, as λ becomes large, QRE predicts play will be concentrated at effort levels very close to the Nash equilibrium. Therefore, if one obtains a large value of λ from the fit, this entails that Q is likely also to be large. However, the relationship for smaller λ is less clear. For example, in treatment PL , we obtain a best-fit $\lambda = 0.267$ and $Q = 0.463$, while in treatment PC , we obtain $\lambda = 0.099$ and $Q = 0.525$. Interpreted qualitatively, this suggests that QRE is a slightly more satisfactory model in PC than in PL for capturing the features of the distribution of play; however, in order for QRE to accommodate this, players in PC must be making larger average optimization errors than in PL .

[Table 4 about here.]

Result 7. *Logit QRE organizes behavior in the share rule treatments better than in the probability rule treatments. In addition, the precision estimates for the share rule treatment QRE are larger, capturing that effort levels are clustered more closely to the Nash equilibrium.*

Support. In Table 4, the quality-of-fit measure Q is higher by a substantial margin for each fit using share rule data than for the corresponding fit using the probability rule. The share rule data are much more consistent with the hypothesis underlying QRE relating the frequency of an effort choice to its expected payoff.

The fitted λ values for QRE in the share rule are also large. The fit for SC has $\lambda = 2.488$ and that for SL has $\lambda = 0.807$, which indicates a high degree of precision in best responses. By way of comparison, Lim et al. (2014) report $\lambda \approx 0.57$ for four-player probability-rule contests, which figure is comparable insofar as both estimates use US dollars as the unit of payoff.¹¹ \square

There is a correlation between the characteristics of the QRE fits and design features of each treatment. QRE is an equilibrium concept, in that players are assumed to have correct beliefs about the distribution of other players' decisions; it relaxes the assumption that players always choose a best reply and instead allows deviations from the best reply, with "mistakes" having larger payoff consequences being made less frequently. In the share rule, feedback about payoff implications of behavior is received without noise, whereas in the probability rule the ex-post realized payoffs have

¹¹Both the estimated λ and the corresponding Q are lowered because of the frequency of efforts of 80 in SL . Figure 2 illustrates that one participant accounts for the majority of instances in which 80 was chosen.

a substantial random component; this would be one explanation why the estimated λ values under the share rule are larger. In addition, the random component of the outcome in the probability rule gives scope for other motivations to operate in the decision-making process of participants, including attitudes towards risk, loss aversion, or a motivation to win the contest. This would make the expected payoff a less salient factor in determining effort choices, which would be captured by a smaller λ value.

The effect of convex costs relative to linear costs is more subtle. The fits reported in Table 4 show that the *PC* has a better quality-of-fit than *PL*, and *SC* a better quality-of-fit than *SL*. However, if we impose a common λ on *PL* and *SL*, and on *PC* and *SC*, the quality-of-fit for the linear cost treatment is better. This is driven by some characteristics of behavior in *PC* which cannot be accommodated by a simple QRE model. In treatment *PC*, convex costs eliminate high effort choices, as effort choices above 49 result in a certain loss. However, there is evidence of a knock-on effect of also removing very low effort choices. In treatment *PL*, 8 of the 36 subjects have median choices below 1, so that in effect they choose not to enter the contest a majority of the time. Conversely, in treatment *PC* most participants make effort choices in the interval between 20 and 30. This behavior cannot be accommodated by QRE; with both the linear and convex cost functions, QRE predicts the modal effort level will be below the Nash equilibrium level of 15.

5 Discussion

Experimental studies of rent-seeking contests find both overbidding and overspreading: efforts are on average significantly higher than the risk-neutral Nash equilibrium prediction, and there is significant variation in effort levels within and across subjects. A number of studies suggest that overbidding in contests can be explained by noise and errors, judgmental biases, a non-monetary utility of winning, and/or evolutionarily stable behavior. Overspreading has been attributed to subjects heterogeneous preferences towards losses, risk, spitefulness, and winning, as well as demographic differences.

In this study, we show how features of the contest environment which do not change the Nash equilibrium prediction nevertheless have significant implications for both overbidding and overspreading in the contest. Specifically, in a 2×2 design, we investigate the effects of allocating the prize by a lottery versus sharing it proportionally, and of a convex cost function versus a linear cost of effort, while holding fixed the Nash equilibrium. We find that the share rule results in average efforts closer to the Nash prediction, and lower variation in individual efforts. Combining the share rule with a convex cost function further enhances these results.

Our findings speak to several puzzles in the literature. First, experimental studies on rank-order tournaments (Lazear and Rosen, 1981) find that there is almost no overbidding and average efforts

are usually consistent with Nash equilibrium (Bull et al., 1987). This is in sharp contrast to the findings from lottery contests (Sheremeta, 2013). Our results suggest that this disparity can be explained by the fact that experiments on rank-order tournaments employ convex cost of effort – which, in that setting, is often needed to obtain a Nash equilibrium in pure strategies (Lazear and Rosen, 1981; Cason et al., 2012) – while experiments on lottery contests employ linear cost of effort. Specifically, we find significant overbidding with linear costs (as in most experiments on lottery contests), but that convex costs can contribute to lowering bidding levels. Second, our findings help to clarify why a subset of studies on rent-seeking contests, including Baik et al. (1999), and Linster et al. (2001) find little to no overbidding. While the majority of experimental studies employ lottery rule to determine the winner of the contest (Dechenaux et al., 2012), the two aforementioned studies employ a share rule. The results of our experiment show that the share rule results in average efforts closer to the Nash equilibrium prediction, although overbidding is not entirely eliminated.

Our study also contributes to a rapidly growing literature on proportional-prize contests. For example, Cason et al. (2010, 2012) examine entry into proportional-prize and single-prize contests, as well as their performance. In contrast to our study, their main focus is on the contest design aspects of different compensation schemes. Morgan et al. (2012) also examine entry into different contests, and find that subjects sort themselves into risky types and safe types. Moreover, subjects' entry decisions are more consistent with theory in contests employing a share rule. Most closely related to our study are working papers by Fallucchi et al. (2013), Masiliunas et al. (2012), and Shupp et al. (2013), who examine how the use of the share rule affects individual behavior in contests. Consistent with those studies we find the share rule encourages behavior which is closer to the Nash prediction. We find additionally that the presence of the convex cost function alongside the share rule is most effective in reducing both overbidding and overspreading.

Finally, our findings contribute to the literature on contest design. Some contest settings arise naturally and are not amenable to design in the implementation of the prize allocation or cost rules. However, when design is possible, our results provide guidance on design from behavioral principles. The experimental literature has shown that behavior in games can vary as a function of design parameters, even when the Nash equilibrium is independent of those parameters (see the elegant review of Goeree and Holt (2001) for a selection of examples). In contests with the probability rule and linear costs, behavior does not appear to be well-organized by equilibrium. Effective design of a contest game in which behavior in the baseline is inconsistent with equilibrium requires an understanding of the drivers of non-equilibrium play. Our results in the *PC* treatment illustrate this. An obvious hypothesis would be that making very aggressive effort choices prohibitively expensive would rein in aggressive play. But equally, the convex cost function lowers the cost of smaller effort levels; participants who might otherwise sit out a contest in the face of aggressive

co-players may now find it worthwhile to choose positive levels of effort. These effects operate in opposite directions, and in our data we find that the effect of increased participation is at least that of the reined-in aggressive players. Our results, then, illustrate that work on design in contests, in the lab and the field, should be informed by an account of the determinants of (non-equilibrium) behavior.

The structure of our design is not able to tease out specifically why the combination of the share rule and a convex cost function reduces overbidding and overspreading. With a proportional prize, there is no longer a clear winner; this would reduce the impact of any non-monetary utility of winning. The proportional prize eliminates a potentially significant amount of objective risk.¹² Proportional prizes may enhance learning incentives, as adjustments in effort choices map concretely to earnings through the share rule, rather than the more abstract probability distribution over earnings implied by the probability rule. Each of these explanations would suggest we would observe less overbidding and less overspreading under the share rule, but our design cannot distinguish to what degree each consideration is playing a role. All these explanations would suggest that the manipulation of using convex costs, which are intended to make more expensive efforts greater than the earnings-maximizing best reply, might be more effective with the share rule, insofar as the share rule eliminates other potentially salient motivators (winning), lowers considerations due to risk, and provides more direct feedback. Our results indicate that further work to understand how these behavioral factors interact would be interesting and useful.

References

- W. Amaldoss and A. Rapoport. Excessive expenditure in two-stage contests: Theory and experimental evidence. In F. Columbus, editor, *Game Theory: Strategies, Equilibria, and Theorems*. Nova Science Publishers, Hauppauge, NY, 2009.
- L. A. Anderson and B. A. Freeborn. Varying the intensity of competition in a multiple prize rent seeking experiment. *Public Choice*, 143:237–254, 2010.
- S. P. Anderson, J. K. Goeree, and C. A. Holt. Rent seeking with bounded rationality: An analysis of the all-pay auction. *Journal of Political Economy*, 106:828–853, 1998.
- E. Baharad and S. Nitzan. Contest efforts in light of behavioral considerations. *Economic Journal*, 118:2047–2059, 2008.

¹²Strategic uncertainty as to the behavior of other players remains, although our results provide evidence that a reduction in strategic uncertainty emerges as well with proportional prizes.

- K. H. Baik, T. L. Cherry, S. Kroll, and J. F. Shogren. Endogenous timing in a gaming tournament. *Theory and Decision*, 47:1–21, 1999.
- C. Bull, A. Schotter, and K. Weigelt. Tournaments and piece rates: An experimental study. *Journal of Political Economy*, 95:1–33, 1987.
- T. N. Cason, W. A. Masters, and R. M. Sheremeta. Entry into winner-take-all and proportional-prize contests: An experimental study. *Journal of Public Economics*, 94:604–611, 2010.
- T. N. Cason, R. M. Sheremeta, and W. A. Masters. Winner-take-all and proportional-prize contests: Theory and experimental results. Working paper, 2012.
- S. H. Chew. A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica*, 51:1065–1092, 1983.
- R. C. Cornes and R. Hartley. Risk aversion, heterogeneity and contests. *Public Choice*, 117:1–25, 2003.
- D. Davis and R. Reilly. Do many cooks always spoil the stew? An experimental analysis of rent-seeking and the role of a strategic buyer. *Public Choice*, 95:89–115, 1998.
- E. Dechenaux, D. Kovenock, and R. M. Sheremeta. A survey of experimental research on contests, all-pay auctions, and tournaments. Chapman University Working Paper, 2012.
- F. Fallucchi, E. Renner, and M. Sefton. Information feedback and contest structure in rent-seeking games. *European Economic Review*, 64:223–240, 2013.
- U. Fischbacher. z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics*, 10:171–178, 2007.
- S. Georganas, D. Levin, and P. McGee. Do irrelevant payoffs affect behavior when dominant strategy is available: Experimental evidence from second-price auctions. Working paper, 2011.
- G. Gigerenzer and R. Selten. *Bounded Rationality: The Adaptive Toolbox*. MIT Press, 2001.
- J. Goeree, C. Holt, and T. Palfrey. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104:247–272, 2002.
- J. K. Goeree and C. A. Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 91(5):1402–1422, 2001.

- G. W. Harrison. Theory and misbehavior of first-price auctions. *American Economic Review*, 79: 749–762, 1989.
- B. Herrman and H. Orzen. The appearance of homo rivalis: Social preferences and the nature of rent-seeking. Working paper, University of Nottingham, 2008.
- A. L. Hillman and E. Katz. Risk-averse rent seekers and the social cost of monopoly power. *Economic Journal*, 94:104–110, 1984.
- D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47:263–291, 1979.
- X. Kong. Loss aversion and rent-seeking: An experimental study. Working paper, University of Nottingham, 2008.
- K. A. Konrad and H. Schlesinger. Risk aversion in rent-seeking and rent-augmenting games. *Economic Journal*, 107:1671–1683, 1997.
- E. P. Lazear and S. Rosen. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89:841–864, 1981.
- W. Lim, A. Matros, and T. L. Turocy. Bounded rationality and group size in Tullock contests: Experimental evidence. *Journal of Economic Behavior & Organization*, 2014.
- B. G. Linster, R. L. Fullerton, M. McKee, and S. Slate. Rent-seeking models of international competition: An experimental investigation. *Defence and Peace Economics*, 12:285–302, 2001.
- S. D. Mago, A. C. Savikhin, and R. M. Sheremeta. Facing your opponents: Social identification and information feedback in contests. ESI Working Paper, 2012.
- A. Masiliunas, F. Mengel, and J. P. Reiss. (Strategic) uncertainty and the explanatory power of Nash equilibrium in Tullock contests. Working paper, 2012.
- R. D. McKelvey and T. R. Palfrey. Quantal response equilibrium for normal form games. *Games and Economic Behavior*, 10:6–38, 1995.
- E. L. Millner and M. D. Pratt. An experimental investigation of efficient rent-seeking. *Public Choice*, 62:139–151, 1989.
- E. L. Millner and M. D. Pratt. Risk aversion and rent-seeking: An extension and some experimental evidence. *Public Choice*, 69:81–92, 1991.

- J. Morgan, H. Orzen, M. Sefton, and D. Sisak. Strategic and natural risk in entrepreneurship: An experimental study. Working paper, 2012.
- W. Müller and A. Schotter. Workaholics and dropouts in organizations. *Journal of the European Economic Association*, 8:717–743, 2010.
- J. C. Potters, C. G. de Vries, and F. van Winden. An experimental examination of rational rent seeking. *European Journal of Political Economy*, 14:783–800, 1998.
- C. R. Price and R. M. Sheremeta. Endowment effects in contests. *Economics Letters*, 111:217–219, 2011.
- C. R. Price and R. M. Sheremeta. Endowment origin, demographic effects and individual preferences in contests. *Journal of Economics and Management Strategy*, forthcoming.
- J. Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3: 323–343, 1982.
- A. Rubenstein. *Modeling Bounded Rationality*. MIT Press, 1998.
- R. M. Sheremeta. Experimental comparison of multi-stage and one-stage contests. *Games and Economic Behavior*, 68:731–747, 2010.
- R. M. Sheremeta. Contest design: An experimental investigation. *Economic Inquiry*, 49:573–590, 2011.
- R. M. Sheremeta. Overbidding and heterogeneous behavior in contest experiments. *Journal of Economic Surveys*, 27:491–514, 2013.
- R. M. Sheremeta and J. Zhang. Can groups solve the problem of over-bidding in contests? *Social Choice and Welfare*, 35:175–197, 2010.
- R. Shupp, R. M. Sheremeta, D. Schmidt, and J. Walker. Resource allocation contests: Experimental evidence. *Journal of Economic Psychology*, 39:257–267, 2013.
- H. Simon. *Models of Bounded Rationality*. MIT Press, 1992.
- S. Skaperdas and L. Gan. Risk aversion in contests. *Economic Journal*, 105:951–962, 1995.
- F. Szidarovszky and K. Okuguchi. On the existence and uniqueness of pure Nash equilibrium in rent-seeking games. *Games and Economic Behavior*, 18:135–140, 1997.
- J. W. Tukey. *Exploratory data analysis*. Addison-Wesley, Reading MA, 1977.

- G. Tullock. Efficient rent seeking. In J. M. Buchanan, R. D. Tollison, and G. Tullock, editors, *Toward a theory of the rent-seeking society*, pages 97–112. Texas A&M University Press, College Station, TX, 1980.
- A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5:297–323, 1992.
- K. Wärneryd. The evolution of preferences for conflict. *Economics Letters*, 116:102–104, 2012.
- N. T. Wilcox. 'Stochastically more risk averse:' A contextual theory of stochastic discrete choice under risk. *Journal of Econometrics*, 162:89–104, 2011.
- W. Wright. Cognitive information processing biases: Implications for producers and users of financial information. *Decision Sciences*, 11:284–298, 1980.

A Instructions

We present as a baseline the instructions for convex costs with each rule. Relative to these instructions, the only change for linear costs was the formula for the cost of each bid, and the corresponding bid cost table.

PC treatment

GENERAL INSTRUCTIONS

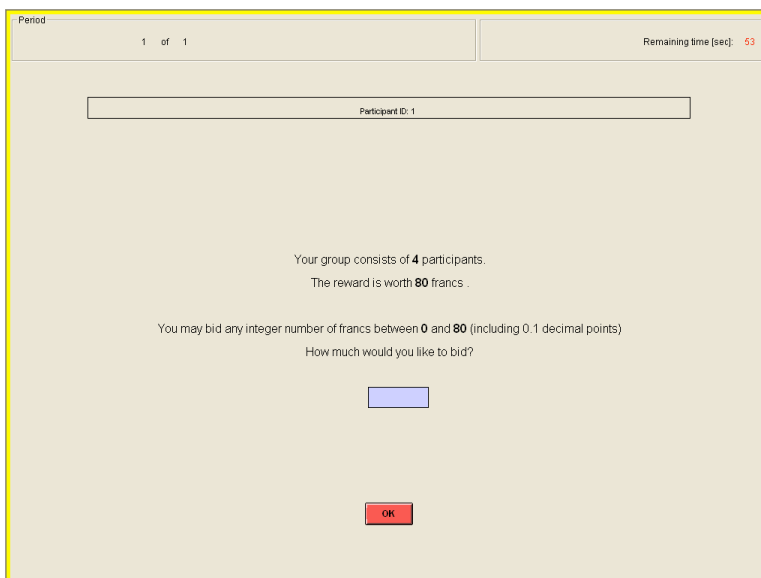
This is an experiment in the economics of strategic decision making. Various research agencies have provided funds for this research. The instructions are simple. If you follow them closely and make appropriate decisions, you can earn an appreciable amount of money.

The currency used in the experiment is francs. Francs will be converted to British Pounds at a rate of 20 francs to 1 pound. You have already received a **£15.00 participation fee**. At the end of today's experiment, you will be paid in private and in cash. **12** participants are in today's experiment.

It is very important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

YOUR DECISION

The experiment consists of **30** decision-making periods. At the beginning of each period, you will be randomly and anonymously placed into a group of **4 participants**. The composition of your group will be changed randomly every period. Each period, you may bid for an **80 francs reward**. You may bid any number between **0** and **80** (including 0.1 decimal points). An example of your decision screen is shown below.



The screenshot shows a decision screen with a light beige background. At the top left, it says "Period 1 of 1". At the top right, it says "Remaining time [sec]: 53". Below this is a horizontal bar labeled "Participant ID: 1". The main text in the center reads: "Your group consists of 4 participants. The reward is worth 80 francs. You may bid any integer number of francs between 0 and 80 (including 0.1 decimal points). How much would you like to bid?". Below the text is a blue rectangular input field. At the bottom center is a red button labeled "OK".

YOUR EARNINGS

For each **bid** there is an associated **cost**. A table is attached to these instructions: each possible bid is given in column A, and its cost is given in column B. Note that as bids rise from 0 to 80, costs rise. The cost of a bid can be also calculated using the following formula:

$$\text{Cost of your bid} = \frac{(\text{Your bid})^2}{30}.$$

The more you bid, the more likely you are to receive the reward. The more the other participants in your group bid, the less likely you are to receive the reward. Specifically, your chance of receiving the reward is given by your bid divided by the sum of all 4 bids in your group:

$$\text{Chance of receiving the reward} = \frac{\text{Your bid}}{\text{Sum of all 4 bids in your group}}.$$

You can consider the amounts of the bids to be equivalent to numbers of lottery tickets. The computer will draw one ticket from those entered by you and the other participants, and assign the reward to one of the participants through a random draw. If you receive the reward, your earnings

for the period are equal to the reward of 80 francs minus the cost of your bid. If you do not receive the reward, your earnings for the period are equal to 0 francs minus the cost of your bid. In other words, your earnings are:

If you receive the award: Earnings = Reward - Cost of your bid = 80 - Cost of your bid

If you do not receive the award: Earnings = 0 - Cost of your bid

AN EXAMPLE

Lets say participant 1 bids 10 francs, participant 2 bids 15 francs, participant 3 bids 0 francs, and participant 4 bids 40 francs. Therefore, the computer assigns 10 lottery tickets to participant 1, 15 lottery tickets to participant 2, 0 lottery tickets to participant 3, and 40 lottery tickets for participant 4. Then the computer randomly draws **one lottery ticket out of 65** (10 + 15 + 0 + 40). As you can see, participant 4 has the **highest chance** of receiving the reward: **0.62 = 40/65**. Participant 1 has a **0.15 = 10/65** chance, participant 2 has a **0.23 = 15/65** chance, and participant 3 has a **0 = 0/65** chance of receiving the reward.

Assume that the computer assigns the reward to participant 4, then the earnings of participant 4 for the period are $26.67 = 80 - 53.33$, since the reward is 80 francs and the cost of bid of 40 is 53.33 as shown on your Cost of Bid table. Similarly, the earnings of participant 1 are $-3.33 = 0 - 3.33$, participant 2 are $-7.5 = 0 - 7.5$, and participant 3 are $0 = 0 - 0$.

At the end of each period, your bid, the sum of all 4 bids in your group, your reward, the cost of your bid, and your earnings for the period are reported on the outcome screen as shown below. Once the outcome screen is displayed you should record your results for the period on your **Personal Record Sheet** under the appropriate heading.

The screenshot shows a web-based interface for an experiment. At the top, it indicates 'Period 1 of 1' and 'Remaining time [sec]: 10'. Below this, a box labeled 'Participant ID: 4' is shown. The main content area displays the following data:

Your bid:	15.0
Sum of all 4 bids in your group:	65.0
Your reward:	80.0
Cost of your bid:	7.5
Your earnings for this period:	72.5

An 'OK' button is located at the bottom of the screen.

IMPORTANT NOTES

You will not be told which of the participants in this room are assigned to which group. At the beginning of each period you will be randomly re-grouped with other three participants to form a 4-person group.

For calculation purposes, you can use the “Calculator” button at the lower left side of your screen. A calculator will appear in your screen after clicking the button.

At the end of the experiment we will randomly choose **5 of the 30** periods for actual payment for this experiment using a computer program. You will sum the total earnings for these 5 periods and convert them to a British pound payment.

Are there any questions?

SC treatment

GENERAL INSTRUCTIONS

This is an experiment in the economics of strategic decision making. Various research agencies have provided funds for this research. The instructions are simple. If you follow them closely and make appropriate decisions, you can earn an appreciable amount of money.

The currency used in the experiment is francs. Francs will be converted to British Pounds at a rate of 20 francs to 1 pound. You have already received a **£15.00 participation fee**. At the end of today's experiment, you will be paid in private and in cash. **12** participants are in today's experiment.

It is very important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

YOUR DECISION

The experiment consists of **30** decision-making periods. At the beginning of each period, you will be randomly and anonymously placed into a group of **4 participants**. The composition of your group will be changed randomly every period. Each period, you may bid for a **share of an 80 francs reward**. You may bid any number between **0** and **80** (including 0.1 decimal points). An example of your decision screen is shown below.

YOUR EARNINGS

For each **bid** there is an associated **cost**. A table is attached to these instructions: each possible bid is given in column A, and its cost is given in column B. Note that as bids rise from 0 to 80, costs rise. The cost of a bid can be also calculated using the following formula:

$$\text{Cost of your bid} = \frac{(\text{Your bid})^2}{30}.$$

The more you bid, the higher is your share of the reward. The more the other participants in your group bid, the lower is your share of the reward. Specifically, the computer will allocate to you a share of the 80 francs reward according to your share of the sum of all 4 bids in your group. In other words, your share is:

$$\text{Share} = 80 \times \frac{\text{Your bid}}{\text{Sum of all 4 bids in your group}}$$

Your earnings for the period are equal to the share of the 80 francs reward minus the cost of your bid. In other words, your earnings are:

$$\text{Earnings} = \text{Share} - \text{cost of your bid}$$

AN EXAMPLE

Lets say participant 1 bids 10 francs, participant 2 bids 15 francs, participant 3 bids 0 francs, and participant 4 bids 40 francs. Then the sum of all 4 bids is 65 (10 + 15 + 0 + 40). As you can see, participant 4 receives the highest share of 80 francs reward: $49.2 = 80 \times 40 / 65$. Participant 1

receives a share of $12.3 = 80 \times 10/65$, participant 2 receives $18.5 = 80 \times 15/65$, and participant 3 receives $0 = 80 \times 0/65$.

The earnings of participant 1 for the period are $8.97 = 12.3 - 3.33$, since the share of the reward is 12.3 and the cost of bid of 10 is 3.33 as shown on your Cost of Bid table. Similarly, the earnings of participant 2 are $11 = 18.5 - 7.5$, participant 3 are $0 = 0 - 0$, and participant 4 are $-4.1 = 49.2 - 53.33$.

At the end of each period, your bid, the sum of all 4 bids in your group, your share, the cost of your bid, and your earnings for the period are reported on the outcome screen as shown below. Once the outcome screen is displayed you should record your results for the period on your **Personal Record Sheet** under the appropriate heading.

Period	
1 of 1	Remaining time [sec]: 20
Participant ID: 2	
Your bid:	40.0
Sum of all 4 bids in your group:	65.0
Your share:	49.2
Cost of your bid:	53.3
Your earnings for this period:	-4.1
OK	

Cost table (common to both treatments)

Table – Cost of Bid

Column A	Column B	Column A	Column B	Column A	Column B
Bid	Cost of Bid	Bid	Cost of Bid	Bid	Cost of Bid
0	0.00	30	30.00	60	120.00
1	0.03	31	32.03	61	124.03
2	0.13	32	34.13	62	128.13
3	0.30	33	36.30	63	132.30
4	0.53	34	38.53	64	136.53
5	0.83	35	40.83	65	140.83
6	1.20	36	43.20	66	145.20
7	1.63	37	45.63	67	149.63
8	2.13	38	48.13	68	154.13
9	2.70	39	50.70	69	158.70
10	3.33	40	53.33	70	163.33
11	4.03	41	56.03	71	168.03
12	4.80	42	58.80	72	172.80
13	5.63	43	61.63	73	177.63
14	6.53	44	64.53	74	182.53
15	7.50	45	67.50	75	187.50
16	8.53	46	70.53	76	192.53
17	9.63	47	73.63	77	197.63
18	10.80	48	76.80	78	202.80
19	12.03	49	80.03	79	208.03
20	13.33	50	83.33	80	213.33
21	14.70	51	86.70		
22	16.13	52	90.13		
23	17.63	53	93.63		
24	19.20	54	97.20		
25	20.83	55	100.83		
26	22.53	56	104.53		
27	24.30	57	108.30		
28	26.13	58	112.13		
29	28.03	59	116.03		




Figure 1: Average efforts over time, by treatment

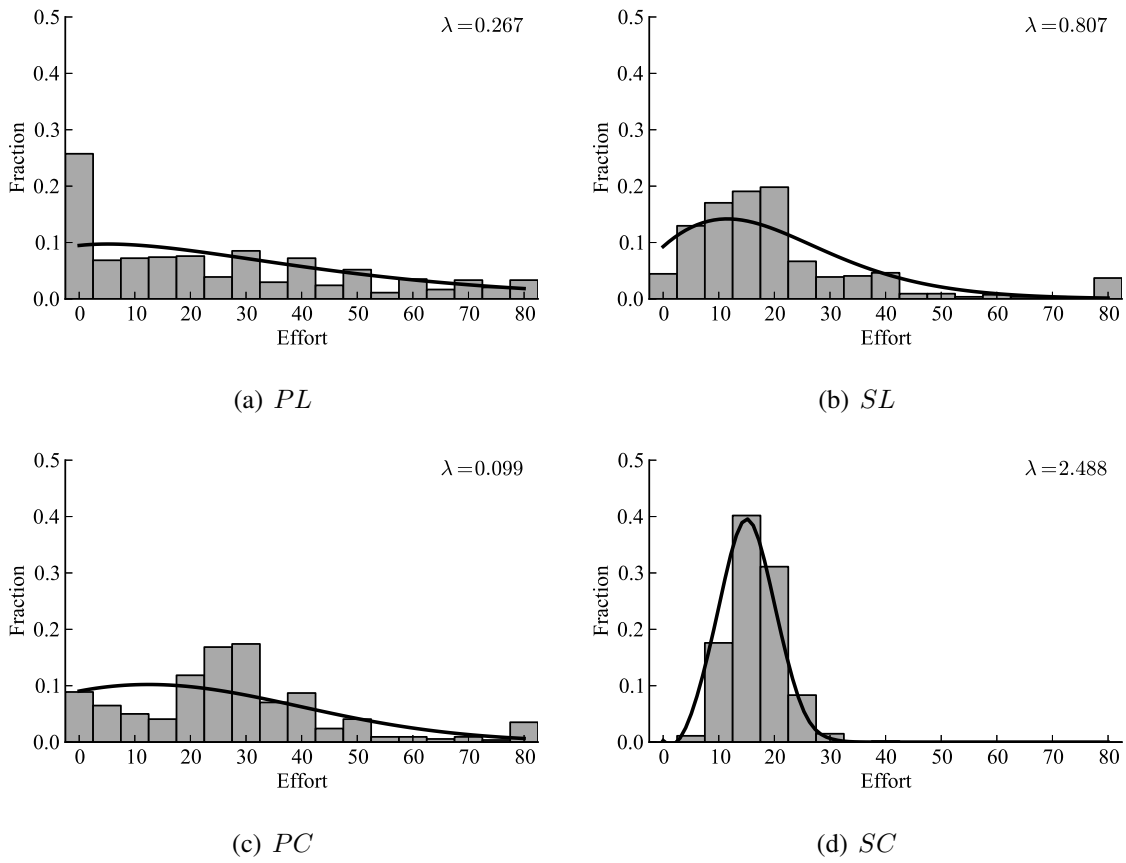


Figure 3: Histogram of all effort choices, last 15 periods. The curved lines represent the best logit quantal response equilibrium fit to the respective empirical distributions.

Effort	Linear			Convex		
	Nash	<i>PL</i>	<i>SL</i>	Nash	<i>PC</i>	<i>SC</i>
Mean	15.0	26.2	23.0	15.0	29.0	17.7
Median	15.0	20.0	17.5	15.0	26.0	16.3
SD	0.0	24.1	20.0	0.0	18.6	8.6

Cost	Linear			Convex		
	Nash	<i>PL</i>	<i>SL</i>	Nash	<i>PC</i>	<i>SC</i>
Mean	15.0	26.2	23.0	7.5	39.4	12.9
Median	15.0	20.0	17.5	7.5	22.5	8.8
SD	0.0	24.1	20.0	0.0	48.0	19.9

Table 1: Summary statistics on effort and cost for each of the four treatments

Session	Treatment	Effort		Variability	
		Mean	SD	V^A	V^W
7	<i>SC</i>	16.8	8.6	2.6	4.0
9	<i>SC</i>	17.1	5.8	3.4	2.1
11	<i>SL</i>	18.6	13.6	8.1	3.1
8	<i>SC</i>	19.1	10.6	4.1	1.8
6	<i>PL</i>	20.5	22.9	22.5	7.1
12	<i>SL</i>	24.4	23.1	21.6	8.1
1	<i>PC</i>	25.4	12.0	9.7	7.4
10	<i>SL</i>	25.9	21.3	8.7	7.6
4	<i>PL</i>	28.0	25.0	18.1	19.4
2	<i>PC</i>	29.1	14.5	10.4	9.4
5	<i>PL</i>	30.1	23.3	23.5	13.8
3	<i>PC</i>	32.3	25.6	22.4	15.9

Table 2: Summary statistics by session, sorted by mean effort, for all participants and all periods. SD is the standard deviation of all effort choices, irrespective of participant. V^A is a measure of across-subject variability in effort, and V^W a measure of within-subject variability in effort, as defined in the text, for the second half of the experiment.

Dependent variable, $o_{it} = e_{it} - 15$	PC & SC	PL & SL	PC & PL	SC & SL
Periods 1-15 - Specification	(1)	(2)	(3)	(4)
S (share) (1 if share and 0 if probability)	-11.50*** (2.39)	-1.39 (4.15)		
C (convex costs) (1 if convex and 0 if linear)			3.25 (3.76)	-6.87*** (2.96)
period (period trend)	-0.28*** (0.07)	-0.64*** (0.14)	-0.27*** (0.06)	-0.65*** (0.14)
constant	17.66*** (1.99)	17.26*** (3.33)	14.30*** (3.17)	15.99*** (2.63)
Observations	1080	1080	1080	1080
Dependent variable, $o_{it} = e_{it} - 15$	PC & SC	PL & SL	PC & PL	SC & SL
Periods 16-30 - Specification	(1)	(2)	(3)	(4)
S (share) (1 if share and 0 if probability)	-11.04*** (1.56)	-5.00* (2.61)		
C (convex costs) (1 if convex and 0 if linear)			2.34 (2.64)	-3.71** (1.51)
period (period trend)	-0.18* (0.11)	-0.11 (0.27)	-0.35 (0.25)	0.05 (0.09)
constant	16.74*** (3.54)	12.78 (8.02)	18.22*** (7.63)	3.97* (2.05)
Observations	1080	1080	1080	1080

Table 3: Panel estimation of treatment effects. Robust standard errors in parentheses. *** indicates significance at the 1% level, ** at 5%, * at 10%. All models included a random effects error structure, with the individual subject as the random effect; standard errors were clustered at the session level.

	L	C	$L&C$
	$\lambda = 0.267$	$\lambda = 0.099$	$\lambda = 0.123$
P	$\ln L = -1453.37$	$\ln L = -1420.38$	$\ln L = -2895.88$
	$Q = 0.463$	$Q = 0.525$	$Q = 0.438$
	$\lambda = 0.807$	$\lambda = 2.488$	$\lambda = 1.298$
S	$\ln L = -1299.52$	$\ln L = -768.95$	$\ln L = -2138.70$
	$Q = 0.705$	$Q = 0.961$	$Q = 0.823$
	$\lambda = 0.429$	$\lambda = 0.210$	$\lambda = 0.265$
$P\&S$	$\ln L = -2796.73$	$\ln L = -2592.03$	$\ln L = -5424.54$
	$Q = 0.535$	$Q = 0.467$	$Q = 0.466$

Table 4: Summary of QRE fits. The body of the table reports fitted λ values, corresponding log-likelihoods, and a measure of quality-of-fit Q for each of the four individual treatments. The margins report fits where λ is constrained to be the same for both treatments in the corresponding row or column. The lower-right cell reports the QRE fit pooling all data from all treatments.