Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Games and Economic Behavior 59 (2007) 213–239
GAMES and Economic Behavior
www.elsevier.com/locate/geb

Competitive burnout: Theory and experimental evidence ☆

J. Atsu Amegashie, C. Bram Cadsby *, Yang Song

Department of Economics, University of Guelph, Guelph, ON, Canada N1G 2W1

Received 25 August 2004

Available online 7 November 2006

Abstract

We examine equilibrium selection in a two-stage sequential elimination contest in which contestants compete for a single prize. This game has a continuum of equilibria, only one of which satisfies the Coalition-Proof Nash Equilibrium (CPNE) refinement. That equilibrium involves "burning out" by using all of one's resources in the first stage. It is Pareto-dominated by many other equilibria. We find that CPNE predicts well when four people compete, but not when eight people compete for two second-stage spots. Using a cognitive hierarchy (CH) framework, we show that when the number of players and the mean number of thinking steps are large, the CH prediction involves burning out. This provides a partial explanation of our results. We also develop a formal argument as to why CPNE logic is more compelling with more players. We conclude that more competition leads to higher bids, and that burning out is indeed a competitive phenomenon.

JEL classification: C72; C91; C92; D44

Keywords: All-pay auction; Burning out; Cognitive hierarchy; Coalition-proof Nash equilibrium; Contests; Experiment; Step thinking

☆ This paper is a revised version of "Competitive Burnout in the Laboratory: Equilibrium Selection in a Two-Stage Sequential Elimination Game".

* Corresponding author. Fax: +1 519 763 8497.
E-mail address: bcadsby@uoguelph.ca (C.B. Cadsby).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2006.08.009

1. Introduction

Contests are an important fact and pervasive aspect of economic life. A contest is a game in which players compete for a prize by making irreversible outlays. Elections, rent-seeking games, R & D races, competition for monopolies, litigation, wars, and sports are all contests.

A common feature of contests is that they involve multiple stages where the set of contestants is narrowed in successive stages of the contest until a winner is finally chosen. Another feature of contests is that the players may be constrained in terms of how much effort or outlay they can expend (e.g., Che and Gale, 1997, 1998; Gavious et al., 2002). In a sequential elimination contest with such a constraint, it may be rational for contestants to expend *all* their efforts in earlier stages, thus burning out and having nothing left to offer in subsequent stages. Amegashie (2004) shows that under certain conditions burning out in this manner may be equilibrium-consistent rational behavior even though the ultimate prize is won only if a contestant is successful in all stages including the final one.

However, the burning-out equilibrium is not the only equilibrium. There are also equilibria in which the players do not burn out. Indeed, there is a continuum of equilibria, many of which are Pareto-rankable. The presence of multiple Pareto-rankable equilibria suggests that it is desirable for the players to coordinate on Pareto-dominant equilibria. Since the burning-out equilibrium is Pareto-dominated by many other equilibria, it is never Pareto optimal to burn out.

Similar kinds of coordination problems are common in many economic contexts. A frequently-cited example is the case of team production. If low effort on the part of one worker reduces the marginal products of other team members, it may not be optimal for a worker to exert high effort when the efforts of another are low. In this case, the team may be stuck at a low-effort equilibrium even though all team members would be better off at a high-effort equilibrium.

Economists and game theorists have proposed solutions to equilibrium selection in such games. These include focal points (Schelling, 1960), belief-learning (Camerer and Ho, 1999), and Pareto dominance (Harsanyi and Selten, 1988). A growing area of research examines coordination games experimentally in order to shed light on the issue of equilibrium selection (e.g., Van Huyck et al., 1990, 1991, 2001; Camerer and Knez, 1994; Anderson et al., 2001; Berninghaus et al., 2002).[1] Generally, this literature finds that smaller groups reach more efficient equilibria than larger groups, especially when play is repeated with a fixed group of participants.

This paper contributes to this line of research by examining equilibrium selection in a two-stage sequential elimination contest in which a group of contestants competes to win a single prize. Only a subset of the participants survives the first stage. In the second stage, the survivors compete once more, with the winner taking home the prize. Like the weak-link team-production coordination game described above, the sequential elimination game has a continuum of Nash equilibria. In contrast to the weak-link coordination game, which has a continuum of Pareto-rankable equilibria, many but not necessarily all of the equilibria in the sequential elimination game are Pareto-rankable. A more significant contrast between the two games is that the main point of a sequential elimination contest is not cooperation to produce a high return for the group, but competition to win a single valuable prize. Thus, in the sequential elimination game, the equilibrium selected through some process of coordination by group members affects the earnings of the group as a whole even as its members compete for the ultimate prize. Is cooperation to maximize group welfare possible in such a competitive context?

[1] Chapter 7 of Camerer (2003) provides an excellent summary of this literature.

A refinement of Nash equilibrium, in particular the Coaliton-Proof Nash Equilibrium (CPNE) concept (Bernheim et al., 1987), suggests that the answer to this question is no. Garratt et al. (2005) find that CPNE has considerable predictive power when it exists in a game of coalition government formation. Gillette et al. (2003, 2004), however, find only limited support for the predictive power of CPNE when compared to that of an equilibrium that is strictly preferred by all agents. The unique CPNE in our game involves the exertion of maximum effort to the point of complete burnout during the first stage of the game, leaving no resources to utilize during the second stage. From the perspective of the competing participants, the burning-out CPNE is Pareto-dominated by many other equilibria in the game. Since the CPNE refinement is Pareto-dominated by many other equilibria, this is a challenging context in which to assess the predictive power of the refinement.

The burning-out equilibrium is somewhat puzzling because of its counter-intuitive prediction that active contestants expend *all* their energies or resources in stage one and thus have nothing left to offer in stage two. Recently, Parco et al. (2005) and Amaldoss and Rapoport (2005) both ran experiments based on an interesting, but rather different two-stage game.[2] In their game, no equilibrium predicts burning out. However, they nonetheless found that their contestants overspent in stage one relative to the equilibrium prediction. In our framework, it is consistent with equilibrium behavior for contestants to go much further and use up all of their resources in the first stage of a two-stage contest. Under what circumstances will we observe the behavior predicted by such an equilibrium, despite its inefficiency and seemingly myopic nature?

Our experimental results show little evidence of cooperation to maximize group welfare. Furthermore, they indicate that the predicted burning-out result is more likely to emerge when there are more players. This contrasts with the CPNE prediction of burning-out regardless of the number of players. We examine this puzzle using a cognitive hierarchy (CH) model, recently developed by Camerer et al. (2004). In that model, players engage in differing numbers of thinking steps, while overconfidently believing that other players engage in fewer thinking steps. We show that when the number of players and the mean number of thinking steps are both sufficiently large, the CH predicts burning out by using all of one's resources in the first stage.

We estimate the mean number of thinking steps based on the experimental data from the first two periods of each session of our eight-period experimental game. We find that it is very close to zero in the initial period for both four-player and eight-player treatments. In the eight-player case, it is substantially higher in the second period. After the first two periods, the CH model is less relevant because players learn about the behavior and beliefs of others as they experience more periods of play. Indeed, the predictions of the CH model are often inconsistent with the results of later periods.

In the later periods, CPNE is not a good predictor of behavior when four people compete for two second-stage spots, but it does predict well when eight people compete for the two available spots. We provide an analysis of this result, arguing that the logic of CPNE is more likely to affect equilibrium selection when the number of players is large since there is more chance that two or more players will deviate from a lower to a higher bid.

[2] The main differences between the model examined by both Parco et al. (2005) and Amaldoss and Rapoport (2005) compared with the model tested here are: (a) they do not use an all-pay auction; (b) they use identical contestants, while our contestants have different valuations; (c) their players compete with only a subset of the contestants in stage one, meeting the other winners of the stage one contests in stage two, while ours meet all the contestants in stage one, playing the subset of players who are successful at stage one in stage two; (d) their game has neither a burnout equilibrium nor an equilibrium at which each player bids zero in stage one; and (e) their game does not have multiple equilibria.

In the next section, we describe and analyze the two-stage sequential elimination game. Section 3 presents the experimental design and Section 4 discusses the results. Section 5 uses the cognitive hierarchy model of Camerer et al. (2004) to examine the relationship between the level of bids and the number of players in the early periods of the game. Section 6 presents a formal discussion based on CPNE of why burning-out occurs when there are eight players, but not when there are four players. Section 7 concludes the paper.

2. A two-stage sequential elimination game

In Amegashie (2004), the following game is presented. Consider N $\geqslant 3$ risk-neutral agents contesting for a prize with valuations commonly known to be $V_1 \geqslant V_2 \geqslant \cdots \geqslant V_{N-1} \geqslant V_N > 0$, where V_i is the valuation of the ith contestant, $i = 1, 2, \ldots, N-1, N$. The contest is divided into two stages. In the first stage, the F contestants with the highest bids or effort levels are chosen to compete in a second stage from which the ultimate winner is chosen, where $2 \leqslant F < N$. Ties are broken randomly in each stage. Formally, the contest success function in stage one is:

$$P_{1i} = \begin{cases} 1 \quad \text{if the } i\text{th contestant belongs to a unique set of constestants} \\ \quad \text{with the top } F \text{ effort levels,} \\ \max\{0, (F-g)/(r+1)\} \quad \text{if } g \text{ contestents bid higher than the } i\text{th contestant and} \\ \quad \text{this contestant ties with } r \text{ other contestants where } r+1 > F-g, \text{ and } g \geqslant 0, \end{cases}$$

where P_{1i} = the probability of advancing from stage one to stage two and e_i = the effort level of player i. In stage two, the contestant with the highest bid wins. Note that the contest in each stage is an all-pay auction.[3]

Following Che and Gale (1997, 1998) and Gavious et al. (2002), suppose all contestants face a common budgetary or effort constraint or cap, $B > 0$. These papers give examples of caps in contests: caps on campaign contributions, salary caps in US professional sports,[4] and caps on how fast Formula 1 racing cars can move. Also, a cap on effort arises because human beings naturally have a limit on how much effort they can expend.

Suppose B can be allocated between the two stages. Let e_i and x_i be the bid or effort levels of the ith contestant in stages one and two respectively, where $e_i + x_i \leqslant B$. We assume that e_i and x_i also represent the cost of expending effort, i.e. the cost function of effort is linear. In each stage, the contestants move simultaneously. Let $P_{1i}(\tilde{e}) = P_{1i}(e_1, e_2, \ldots, e_N)$ and $P_{2i}(\tilde{x}) = P_{2i}(x_1, x_2, \ldots, x_F)$ be the success probabilities of the ith contestant in stages one and two respectively. Denote the equilibrium success probabilities by $P_{1i}^*(\tilde{e}^*)$ and $P_{2i}^*(\tilde{x}^*)$ for the ith contestant.

In stage two, the equilibrium expected payoff of the ith contestant, conditional on making it to that stage, is $\Pi_{2i}^* = P_{2i}^*(\tilde{x}^*)V_i - x_i^*$. Focusing on a subgame perfect Nash equilibrium and applying backward induction, the equilibrium payoff to the ith contestant in stage one is $\Pi_{1i}^* = P_{1i}^*(\tilde{e}^*)\Pi_{2i}^* - e_i^*$.

The solution to this game is summarized in the following proposition:

Proposition 1. *Consider a two-stage contest where the contest in each stage is an all-pay auction and the contestants have valuations commonly known to be* $V_1 \geqslant V_2 \geqslant \cdots \geqslant V_{N-1} \geqslant$

[3] See Baye et al. (1996) and Clark and Riis (1998) for analyses of all-pay auctions.

[4] As noted by Gavious et al. (2002), in the year 2000, NFL teams faced a salary cap of $62,172,000. This was a cap on the aggregate amount they could spend on their top 51 salaried players.

V_N. If $F \geqslant 2$ contestants are chosen in the first stage to compete in the second stage and all the contestants face a common budget (effort) constraint, B, which can be allocated between the two stages, then a given equilibrium effort allocation $(e^*, B - e^*)$ between the two stages induces a corresponding equilibrium number of active contestants, K, such that $\Pi_i^* = (F/K)[(1/F)V_i - (B - e^*)] - e^* \geqslant 0$ for $e^* \in [0, B]$, $i = 1, 2, \ldots, K - 1, K$ and $\Pi_i^* = (F/(K+1))[(1/F)V_i - (B - e^*)] - e^* < 0$ for $e^* \in [0, B]$, $i = K+1, K+2, \ldots, N-1, N$ and $F < K \leqslant N$. In any equilibrium, the active contestants $i = 1, 2, \ldots, K - 1, K$ bid e^* in stage one and $B - e^*$ in stage two, while the rest bid zero in each stage.[5]

Proof. In any equilibrium the expected payoff for the ith active player is $\Pi_i^* = (F/K)[(1/F) \times V_i - (B - e^*)] - e^* \geqslant 0$, $i = 1, 2, \ldots, K - 1, K$. If $F \geqslant 3$, a player who deviates from this equilibrium by bidding marginally more than e^* in stage one guarantees entry to stage two, but will then lose in stage two with certainty since he/she will be joined by, at least two players who, having bid e^* in stage one, have bigger caps in stage two. There exists a pure-strategy equilibrium in the stage-two subgame in which the players with the bigger cap in stage two will bid their cap. This will yield an expected payoff lower than the equilibrium expected payoff for the player who deviated. If $F = 2$, a player who deviates by bidding marginally more than e^* in stage one guarantees entry to stage two, but will be joined by a player who bid e^* in stage one and hence has a bigger cap in stage two. In this case, there is no equilibrium in pure strategies in the stage-two subgame. However, in any mixed-strategy equilibrium in stage two, the player with the smaller cap will get a zero expected payoff,[6] which is less than the expected payoff in the symmetric equilibrium in which everyone bids e^* in stage one. Hence, it is not profitable for any player to deviate by bidding more than e^* if $F = 2$. A player who bids less than e^* in stage one will lose with certainty in that stage, yielding an expected payoff lower than the equilibrium expected payoff. Hence there is no profitable deviation from the equilibrium stated in the proposition for an active player. The players $i = K + 1, \ldots, N - 1, N$, have no incentive to participate if $[F/(K + 1)][(1/F)V_i - (B - e^*)] - e^* < 0$ for $e^* \in [0, B]$. □

According to Proposition 1, different values of e^* may induce different numbers of active contestants, K. If K and e^* vary simultaneously, a Pareto ranking of the different equilibria is not straightforward. For the sake of exposition, we initially investigate the Pareto ranking of equilibria that share a common number of active participants, K. For a given K, all such equilibria can be ranked by noting that $\delta\Pi_i^*/\delta e^* = F/K - 1 < 0$. Hence the equilibrium with the lowest e^* gives the highest expected payoff and the equilibrium with the highest e^* gives the lowest expected payoff for $i = 1, 2, \ldots, K - 1, K$. This of course implies that the burning-out equilibrium in which $e^* = B$, the highest possible e^*, is Pareto-dominated by all other equilibria with the same number of active participants, K, since each of those equilibria has an $e^* < B$.

As indicated above, a general Pareto ranking of the different equilibria is less straightforward when comparing equilibria with different Ks. When equilibria with different Ks exist, the burning-out equilibrium may not be Pareto-dominated by all other equilibria.[7] However, there

[5] Equilibria may also exist in which a player with a lower valuation is active (i.e., bids a positive amount in at least one of the stages) while a player with a higher valuation bids nothing in either stage. The existence of such an equilibrium requires that the difference in valuations between these two players be sufficiently small. We do not focus on such equilibria. Note also that we assume that if a player is indifferent between participating in the contest and staying out, he will participate.

[6] For a proof of this result, see Appendix A available at http://www.uoguelph.ca/~jamegash/burn_out_GEB.pdf.

[7] For a proof of this result, see Appendix B1 available at http://www.uoguelph.ca/~jamegash/burn_out_GEB.pdf.

will always be many equilibria, including all of those with the same number of participants as the burning-out equilibrium, that will Pareto-dominate burning out. For the parameters used in our experimental treatments, all equilibria Pareto-dominate burning out.

If we apply the Coalition-Proof Nash Equilibrium (CPNE) refinement,[8] which allows for joint deviations, the burning-out equilibrium, in which $e^* = B$, is the only surviving pure-strategy equilibrium. To see this, consider an equilibrium in which all the contestants in stage one bid $e^* < B$. Suppose a group of M contestants deviate by bidding marginally more than e^* in stage one.[9] If $M = F \geqslant 2$, then they are all guaranteed entry to stage two. Their payoff will be $\Pi_i^d = (1/F)V_i - B > 0$. It is easy to show that $\Pi_i^d > \Pi_i^*$ as long as $(1/F)V_i - (B - e^*) > 0$ which is true for all active players. Note that such a deviation by the $M = F$ players is immune to further deviations by sub-coalitions of this deviating group, since each coalition member's probability of success in stage one is already at a maximum (i.e., 1). Hence, there exists a profitable joint deviation from any equilibrium where $e^* < B$.[10] Neither a single nor joint deviation is feasible at $e^* = B$. Thus, $e^* = B$ is the unique pure-strategy CPNE. This leads to the following proposition:

Proposition 2. *Consider a two-stage contest where the contest in each stage is an all-pay auction and the contestants have valuations commonly known to be $V_1 \geqslant V_2 \geqslant \cdots \geqslant V_{N-1} \geqslant V_N$. If $F \geqslant 2$ contestants are chosen in the first stage to compete in the second stage and all the contestants face a common budget (effort) constraint, B, which can be allocated between the two stages, then there exists a continuum of symmetric pure-strategy Nash equilibria in which each active contestant bids $e^* \in [0, B]$ in stage one and $B - e^*$ in stage two but $e^* = B$ is the only coalition-proof Nash equilibrium.*

We experimentally investigate the following issues. First, how does the value of the prize affect the effort or bid level? Given K active contestants bidding $e = e^*$ with $e^* \in [0, B]$, a risk neutral player i should bid e^* in stage one and $B - e^*$ in stage two if $(F/(K+1))[(1/F)V_i - (B - e^*)] - e^* \geqslant 0$ and should bid zero in both stages if $(F/(K+1))[(1/F)V_i - (B - e^*)] - e^* < 0$. Actual players need not be risk-neutral. Nonetheless, for each player there should be a critical valuation level consistent with his/her level of risk aversion that would induce a bid of e^* rather than zero.

Second, do we observe Pareto-preferred equilibria, or do we find the burning out predicted by the CPNE refinement, despite the fact that this unique CPNE is Pareto-dominated by other pure-strategy Nash equilibria? Under what if any circumstances will players allocate all their efforts to stage one when there is another stage ahead? Will there be a process of convergence to the burning-out CPNE over the rounds of a finitely repeated game?

Third, will the feedback received between rounds make a difference to the convergence process? Whether or not winning bids are announced at the end of each stage makes no difference to Nash equilibrium predictions. Nash equilibria are based on consistent beliefs, beliefs

[8] See Bernheim et al. (1987) for a discussion of CPNE.

[9] In the experiment, only integer bids were permitted. Thus, in our experimental context, a bid marginally more than e^* may be interpreted as a bid of $e^* + 1$.

[10] Notice that a deviation by $M > F$ players to bid more than e^* is not immune to further deviations by a sub-coalition of F players. A deviation is also not profitable for $M < F$ players because they will be joined by at least one player who has a bigger cap in stage two. In any case, to show that any equilibrium with $e^* < B$ is not CPNE, we only need to show that there exists a coalition size which can deviate profitably.

that are simply confirmed with announcements of winning bids. However, a number of recent papers have suggested that the type of feedback provided between periods of play can significantly affect bids in first-price sealed bid auctions despite having no effect on Nash predictions (e.g., Neugebauer and Selten, 2003; Ockenfels and Selten, 2005). In particular, Neugebauer and Selten (2003) found that bids were significantly higher when winning bids were revealed to participants than when they were told only whether they had won the auction or not.[11] They attributed this result to an asymmetry that arises when only winning bids are revealed. Losers receive a clear signal about how much more they should have bid to win the auction. However, winners do not receive an analogous signal about how much less they could have bid without losing the auction. Neugebauer and Selten (2003) argue that this asymmetric revelation of winning bids pushes bids upward over repeated rounds of play. Similarly, we hypothesize that announcing successful bids might promote higher stage-one bids as our two-stage all-pay auction unfolds, leading to faster convergence to the burning-out CPNE.

Fourth, how does the number of players affect the equilibrium. Earlier experimental studies of coordination games have shown that coordination on Pareto-superior outcomes is harder to sustain with more players. For example, Camerer and Knez (1994) argue that coordination on Pareto-superior outcomes in their minimum-effort coordination game was difficult to sustain for more than two players because forming beliefs about the behavior of other players becomes more complex with larger numbers. While two players only have to worry about each other's beliefs, the introduction of additional players forces everyone to think about the beliefs that each player has about the others in order to predict behavior. In the above analysis the uniqueness of the burning-out CPNE is independent of the number of players. However, the predictive power of the burning-out CPNE may depend on the number of players, since the higher the number of players, the more likely it is that some coalition of $F \geqslant 2$ players will deviate from a non-burning-out equilibrium, as discussed more formally in our theoretical analysis of the results in Section 6.

3. Experimental design

We ran twelve sessions with participants who were undergraduate students at the University of Guelph. They were recruited in the University Centre. A thirteenth session was run using economics professors at the University of Guelph. Participants received a \$3.00 Canadian show-up fee. The rest of their earnings depended on their performance in the game. Average earnings were \$13.20 Canadian, equal to about \$10.00 US, inclusive of the show-up fee. The sessions lasted about one hour.

Upon entering the room, participants were asked to take a seat and were assigned a player number. Written instructions were distributed.[12] The instructions were then read aloud while participants followed along on their own copies. The experiment lasted for eight periods, each of which was divided into two stages. At the beginning of each period, each participant was asked to draw an envelope containing an information slip from a box held by the experimenter. The randomly selected information slip told each participant his/her potential prize value. There were four different prize values. Participants were also told the prize values assigned to the other

[11] In contrast, Dufwenberg and Gneezy (2002) found that when the entire vector of bids was announced, this feedback affected behavior in the lab. However, announcing only the winning bid did not affect behavior relative to announcing nothing between periods of play.

[12] The instructions are available in Appendix C at http://www.uoguelph.ca/~jamegash/burn_out_GEB.pdf.

players. The potential prize values determined the monetary payoff of each participant if he/she won the prize at the end of stage two. The information slip also indicated that each participant had an endowment of 50 tokens, some or all of which could be used to place bids in stages one and two. Each token was worth two cents Canadian. Any tokens that were not used in either stage could be cashed in at the end of the game.

In stage one, participants were given the opportunity to bid any integer amount of tokens between zero and their budgetary caps of 50. After writing their bids in the designated space on their information slips, participants raised their hands and the experimenter collected the slips. Participants understood that once bids were placed, the amount bid would not be returned, regardless of whether or not they won the prize. The two participants with the highest bids were then privately informed that they would move on to stage two. Ties were broken randomly by a draw. Other participants were informed privately that they would not be moving on. Their earnings for the period were 50 tokens minus their stage-one bids.

The two participants who reached stage two were then given the opportunity to bid any amount of tokens from zero up to whatever number of tokens remained after their stage-one bids by writing the desired amount in the designated space on their information slips. The participants who had not reached stage two were asked to write zero in the designated space so that it would not be obvious which two players were still in the game. The person who placed the highest stage-two bid was then privately informed that he/she had won the prize, which was worth the amount that had been indicated on his/her information slip. As in stage one, a random draw was used to determine the final winner if both participants bid the same amount.

At the end of each period, the information slips were returned to each participant, indicating his/her earnings for the period. Earnings were equal to the 50-token endowment plus the payoff from playing the game. Thus, the earnings of the final winner consisted of the 50-token endowment, minus the tokens bid in stages one and two, plus the prize value drawn at the beginning of the period. The earnings of the other participants consisted of the 50-token endowment, minus the bid or bids placed during the period.

At the beginning of a new period, each participant drew a new information slip at random, containing a new prize value. Tokens from earlier periods could not be used in the new period. Each participant began each period with exactly 50 tokens.

We ran four treatments, which are summarized in Table 1.

Treatment 1: Four persons, no announcement of winning bids: In the first treatment, four persons participated in the game. Participants were informed at the end of stage one whether they would advance to stage two. However, they were not given any information about the level of the successful bids. Similarly, at the end of stage two, continuing participants were informed whether they had won the prize. However, they were not told the level of the winning bid.

Table 1
Summary of treatments

	4-person group	8-person group
Without announcement (eight periods)	3 sessions with students	3 sessions with students
With announcement (eight periods)	3 sessions with students 1 session with economics professors (excluded from statistical analysis)	3 sessions with students

Treatment 2: Four persons, announcement of winning bids: Once again in treatment 2, four persons participated in the game. However, in this treatment, the two stage-one bids of those moving on to stage two were publicly announced after stage one and the stage-two bid of the final winner was publicly announced after stage two.

Treatment 3: Eight persons, no announcement of winning bids: In treatment 3, eight persons participated in the game. As in treatment 1, successful bids were not announced.

Treatment 4: Eight persons, announcement of winning bids: In treatment 4, eight persons participated in the game. As in treatment 2, successful bids were announced.

In treatments 1 and 2, the prize values assigned randomly to the four participants were set at 100, 170, 230 and 300 tokens. Consider a risk-neutral participant who believed the other three participants would also behave as if they were risk-neutral. If such a participant drew the possibility of winning the 100-token prize, Proposition 1 indicates that he/she would bid zero in both stages for all non-zero equilibria since $(F/K+1)[(1/F)V_i-(B-e^*)]-e^*<0$, for $e^*\in(0,B]$ in this case. If $e^*=0$, then $(F/K+1)[(1/F)V_i-(B-e^*)]-e^*=0$. Given the assumption that a player who is indifferent between participating in the contest and staying out will participate, the risk-neutral player with a valuation of 100 tokens would bid zero in stage one and $B=50$ in stage two. However, if such a risk-neutral participant drew the possibility of winning one of the other three prizes, Proposition 1 indicates that he/she would bid $0\leqslant e^*\leqslant B$ in equilibrium in stage one and $x^*=B-e^*$ in stage two since $(F/K+1)[(1/F)V_i-(B-e^*)]-e^*>0$ in these cases. The available equilibria for the two treatments with four participants are summarized in the top half of Table 2. The burning-out CPNE in which $K=3$ and $e^*=B=50$ is the worst equilibrium in that it is Pareto-dominated by all of the other Nash equilibria, while $K=3$ and $e^*=1$ is the Pareto-optimal equilibrium in the four-player case.[13]

In treatments 3 and 4, the prize values were doubled relative to treatments 1 and 2 in order to hold expected earnings constant across the four- and eight-person treatments. The prize values were accordingly set at 200, 340, 460 and 600 tokens. Each of these prize values was randomly assigned to two of the eight participants. Employing the same reasoning as above, risk neutrality implies a bid of zero for those drawing the 200-token prize in stages one and two when $20<e^*\leqslant B$. When $16.667<e^*\leqslant 20$, the equilibrium calls for one of the players with the 200-token valuation to bid e^* in stage one and $B-e^*$ in stage two, while the other bids zero in both stages. Both of the players with the 200-token valuations will bid e^* in stage one and $B-e^*$ in stage two in any equilibrium in which $0\leqslant e^*\leqslant 16.667$. Those drawing any of the other prize values will

Table 2
Summary of the complete set of Nash equilibria

Total number of players, N	Number of active players, K	Symmetric bids of active players in stage one	Identity of non-active players
4	4	$e^*=0$	None
4	3	$0<e^*\leqslant 50$	Player 4
8	8	$0\leqslant e^*\leqslant 16.667$	None
8	7	$16.667<e^*\leqslant 20$	Player 7 or 8 but not both
8	6	$20<e^*\leqslant 50$	Players 7 and 8

Note. The stage-two bid $x^*=B-e^*$ for all active players.

[13] For proofs, see Appendix B2 at http://www.uoguelph.ca/~jamegash/burn_out_GEB.pdf.

place a bid of $0 \leqslant e^* \leqslant B$ in equilibrium in stage one and $x^* = B - e^*$ in stage two. The available equilibria for the two treatments with eight players are summarized in the bottom half of Table 2. The burning-out CPNE in which $K = 6$ and $e^* = B = 50$ is the worst equilibrium in that it is dominated by all other equilibria in this case.[14] With eight players, there are two Pareto-optimal equilibria that are not themselves Pareto-rankable: $K = 8$, $e^* = 0$ and $K = 6$, $e^* = 21$.[15]

Three sessions of each treatment were run using undergraduate student participants and were analyzed in a two-by-two factorial design framework. One session of treatment 2 was run using economics professors. As discussed above, we hypothesized that both announcements of the winning bids and larger numbers of players might facilitate convergence to the burning-out CPNE. In the case of announcements, we conjectured that if everyone learned how much those moving on to stage two had bid in stage one, it might encourage attempts to bid even higher. In the case of eight-person versus four-person competitions, we reasoned that more competitors would increase the likelihood of coalition formation and defection, pushing bids higher.

4. Results

We focus our analysis on the stage-one bids. The CPNE refinement calls for all participants for whom the prize value is sufficiently large to burn out by bidding their entire 50-token endowment in the first stage. Participants for whom the prize value is not large enough to justify bidding withdraw from the contest by bidding zero. Of course, any outcome in which all active participants bid a common amount in stage one is consistent with a Nash equilibrium. The CPNE is Pareto inferior to all of the other pure-strategy Nash equilibria in both the four- and eight-person treatments.

Figures 1 to 5 present representative results from five of the 13 experimental sessions, one from each of the student treatments as well as the one session with economics professors as participants. The bars in the figures indicate the bids placed by the individual participants in the first stage of each period. The bars are ordered by prize value from lowest to highest in each period. The participant numbers appear beneath the figure. Asterisks indicate bids of zero.

Table 3 reports the mean bids and bid standard deviations for active players in the final period of each session.[16] The Pareto-optimal equilibria ($e^* = 1$, $K = 3$ in the four-person treatments; $e^* = 0$, $K = 8$ or $e^* = 21$, $K = 6$ in the eight-person treatments) were not achieved in any of the experimental sessions. The economics professors playing the four-person announcement treatment, illustrated in Fig. 3, came closest, converging to a bid of about $e^* = 20$, $K = 3$, which was nonetheless still a whopping 19 tokens above the Pareto-optimal equilibrium bid for the four-person case.

Mean active bids in the final period of the eight-person sessions were all within 3.5 tokens of the CPNE burning-out equilibrium. Standard deviations were less than one in all but one eight-person session. While eight-person sessions converged to a bid very close to the CPNE, four-person sessions did not. Mean bids were dramatically lower in all but one four-person session.

[14] The proof is in Appendix B3 at http://www.uoguelph.ca/ ~jamegash/burn_out_GEB.pdf.

[15] The proof is in Appendix B4 at http://www.uoguelph.ca/~jamegash/burn_out_GEB.pdf.

[16] We defined active players as those bidding more than one. There were three instances of players bidding one in the last period of a session. None of these three players bid zero in any of the other periods. Thus, it is possible that they did not understand that they were permitted to bid zero, and thus bid one rather than zero when they did not want to compete for the prize.

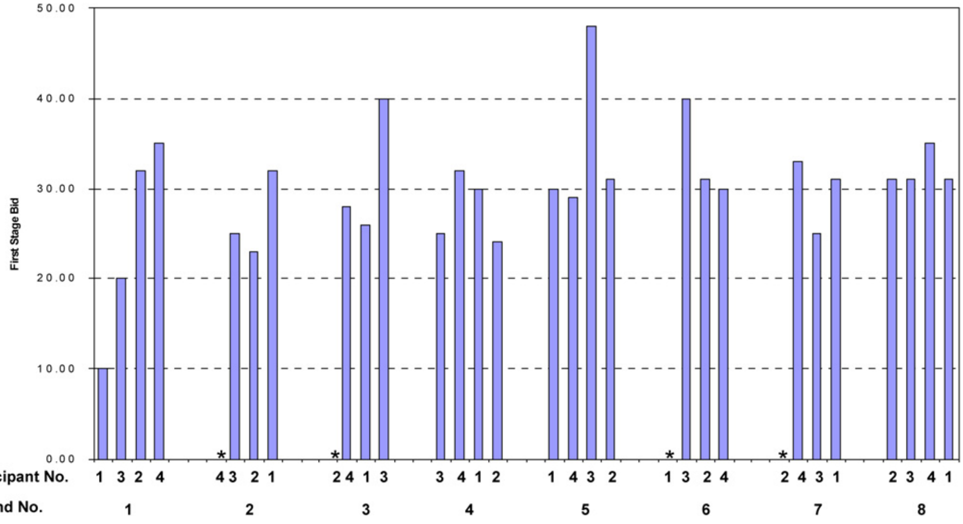

Fig. 1. Four persons without announcement.

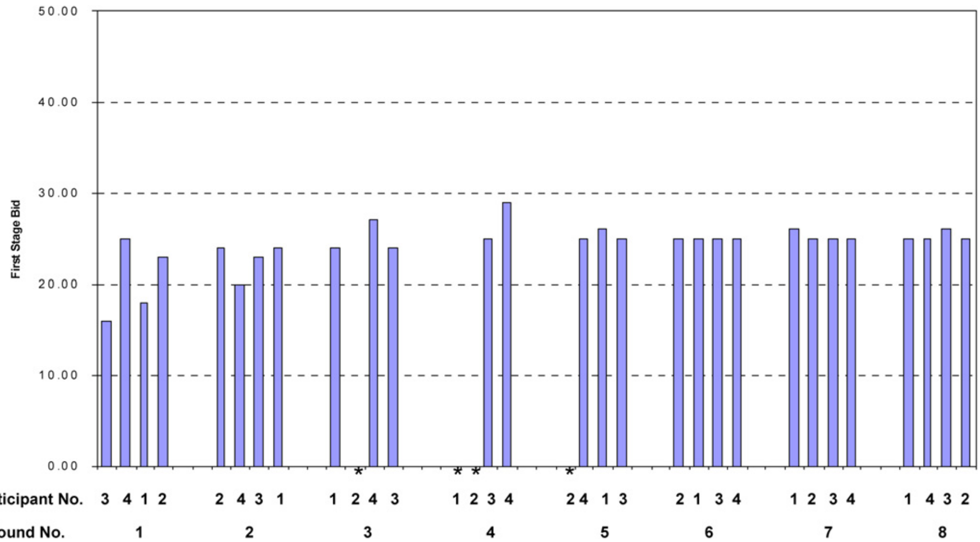

Fig. 2. Four persons without announcement.

Standard deviations of active bids were greater than one in four of the six four-person student sessions, indicating less convergence to one of the Nash equilibria.

The figures also indicate that some participants placed zero bids. However, only in the case of the economics professors did the bidding behavior suggest reasonably consistent risk neutrality. In every period except period two, the professor who drew the lowest prize value of 100 bid zero. In both periods two and eight, the professor who drew the second-lowest prize value of 170 also placed a zero bid, showing some risk aversion. In the student sessions, some participants who

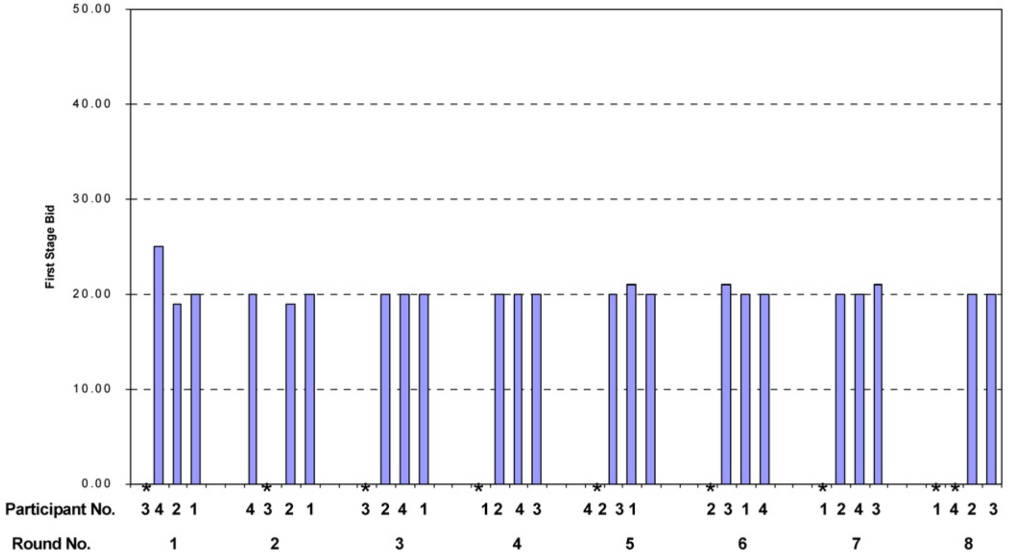

Fig. 3. Four professors without announcement.

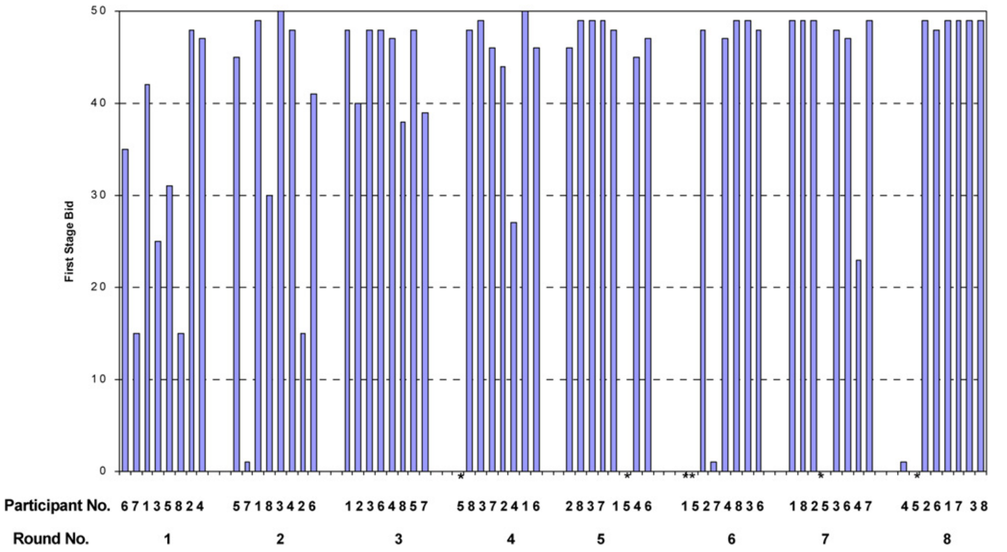

Fig. 4. Eight persons without announcement.

drew low prize values bid positive amounts, while some who drew higher prize values bid zero. Thus, there is evidence of both risk-averse and risk-loving behavior.

If participants had different attitudes toward risk, the prize value required to produce a level of expected earnings high enough to warrant a positive bid at a given e^* would differ from person to person. However, one would nonetheless expect the overall probability of a positive bid to be higher, the higher the prize value drawn. In fact, those drawing the lowest prize bid zero 30% of the time, those drawing the second lowest prize bid zero 15% of the time, those drawing the

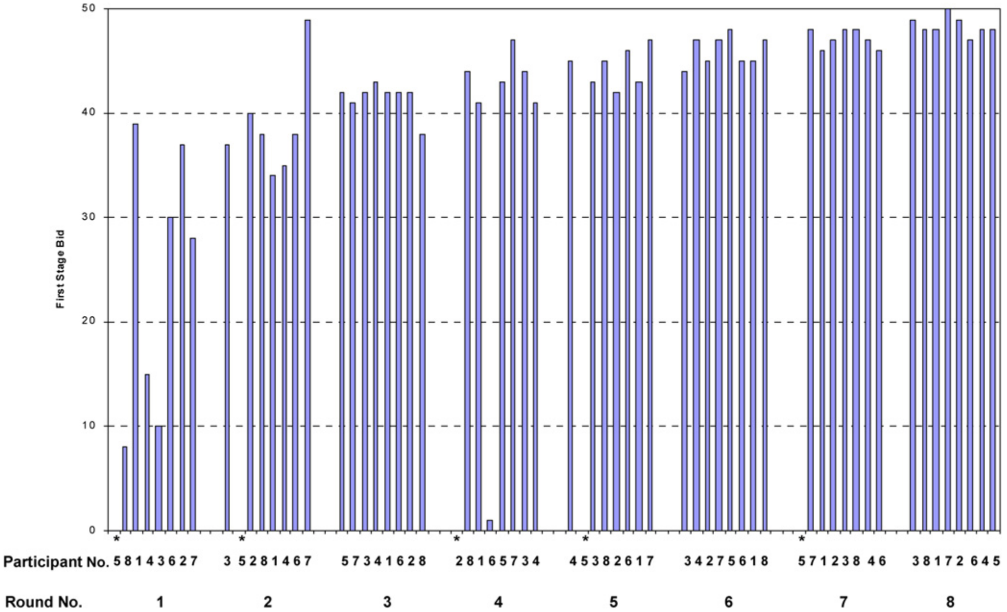

Fig. 5. Eight persons without announcement.

Table 3
Period-eight mean bids and standard deviations in various samples

Sample	Winning bids announced	Sample size	Mean active bid in period eight	Active bid std dev in period eight
1	No	4	32.00	2.00
2	No	4	27.75	16.52
3	No	4	32.33	1.53
4	Yes	4	25.25	0.50
5	Yes	4	45.67	0.58
6	Yes	4	35.50	5.26
7	No	8	48.83	0.41
8	No	8	46.71	2.69
9	No	8	48.57	0.52
10	Yes	8	50.00	0.00
11	Yes	8	50.00	0.00
12	Yes	8	48.38	0.92
Econ. profs.	Yes	4	20.00	0.00

second highest prize bid zero 6% of the time, and those drawing the highest prize bid zero just 4% of the time. These observations indeed suggest a positive relationship between the probability of a positive bid and the prize value drawn. To examine this issue more formally, we use a three-level hierarchical logit model and estimate it using the data from the 12 student sessions.[17] The binary dependent variable is equal to one if a positive bid is placed and zero if a zero bid is placed.

[17] Raudenbush and Bryk (2002), and Snijders and Bosker (1999) both provide excellent discussions of hierarchical linear and logit models (also called mixed models or random-effects models) incorporating both fixed and random effects.

We hypothesize that the probability of a positive bid will be positively related to the prize value drawn, while controlling for the period of play, possible treatment effects, and random effects related to both individual participants and particular sessions.

Level 1 is a logit model, defined for each individual participant i in every session s over the eight periods of play t:

$$\log\left[P_{tis}/(1-P_{tis})\right] = \pi_{0is} + \pi_{1is}(PER_t) + \pi_{2is}(NV_{tis}), \tag{1}$$

where P_{tis} is the probability of a positive bid in period t for individual i in session s, PER_t is the period number minus eight in period t, NV_{tis} is the normalized prize value in period t for individual i in session s, and the πs are individual-level coefficients. Subtracting eight from the period number allows the effect of treatment variables that may interact with the period of play to be tested during the last period of the game when convergence to an equilibrium is most likely to have occurred. The prize value is normalized to correspond with the expected earnings it represents by dividing prize values by the number of participants in the session, either four or eight.

The level-2 model takes account of possible individual-specific random effects on the level-1 coefficients:

$$\begin{aligned} \pi_{0is} &= \beta_{00s} + \eta_{0is}, \\ \pi_{1is} &= \beta_{10s} + \eta_{1is}, \\ \pi_{2is} &= \beta_{20s} + \eta_{2is}, \end{aligned} \tag{2}$$

where the βs are session-level coefficients and the ηs represent individual-specific random effects.

The level-3 model takes account of possible session-specific treatment and random effects on the level-2 coefficients:

$$\begin{aligned} \beta_{00s} &= \gamma_{000} + \gamma_{001}(NA_s) + \gamma_{002}(8P_s) + \mu_{00s}, \\ \beta_{10s} &= \gamma_{100} + \gamma_{101}(NA_s) + \gamma_{102}(8P_s) + \mu_{10s}, \\ \beta_{20s} &= \gamma_{200} + \gamma_{201}(NA_s) + \gamma_{202}(8P_s) + \mu_{20s}, \end{aligned} \tag{3}$$

where the γs are level-3 coefficients and the μs represent possible session-specific random effects. The treatment dummy variable NA_s is equal to zero for sessions in which the winning bids are announced and one if they are not announced. The treatment dummy variable $8P_s$ is equal to zero for the four-person treatments and one for the eight-person treatments. Combining the three sets of equations, we estimate:

$$\begin{aligned} &\log\left[P_{tis}/(1-P_{tis})\right] \\ &\quad = \gamma_{000} + \gamma_{001}(NA_s) + \gamma_{002}(8P_s) + \gamma_{100}(PER_t) + \gamma_{101}(PER_t \times NA_s) \\ &\qquad + \gamma_{102}(PER_t \times 8P_s) + \gamma_{200}(NV_{tis}) + \gamma_{201}(NV_{tis} \times NA_s) + \gamma_{202}(NV_{tis} \times 8P_s) + \eta_{0is} \\ &\qquad + \eta_{1is}(PER_t) + \eta_{2is}(NV_{tis}) + \mu_{00s} + \mu_{10s}(PER_t) + \mu_{20s}(NV_{tis}). \end{aligned} \tag{4}$$

Table 4 reports the results. The prize value is positively related to the probability of a positive bid as hypothesized, rejecting the null hypothesis with a two-tailed p-value of 0.076, which corresponds to a one-tailed p-value of 0.038. We can thus reject the null in the direction of the hypothesized positive relationship. Neither the period variable nor either of the treatment variables or their interactions is significantly related to the probability of a positive bid. Thus,

Table 4
Positive versus zero bid results

Independent variables	Estimate	t value	$Pr > \|t\|$
Intercept	0.004332	0.004	0.997
No announcement (NA)	−0.635910	−0.516	0.618
8 Participants (8P)	−0.980644	−0.762	0.465
Adjusted period (PER)	0.008298	0.059	0.954
NA × PER	0.045866	0.331	0.748
8P × PER	−0.191543	−1.272	0.236
Normalized valuation (NV)	0.052823	1.999	0.076
NA × NV	0.037350	1.262	0.239
8P × NV	−0.009759	−0.320	0.756

Notes. Repeated measures three-level hierarchical logit model with random effects on intercept, period and normalized valuation, using full PQL (penalized quasi-likelihood) estimation. Equation estimated: Eq. (4).

the positive relationship between prize value and the probability of a positive bid appears to be invariant to both the period in which the prize is drawn and the four treatments. If we drop all of the insignificant variables, maintaining only NV and the individual-specific and session-specific random effects, the two-tailed p-value on NV falls to 0.001, strongly supporting the hypothesized relationship.[18]

We are primarily interested in how close participants came to the burning-out CPNE in the various treatments. The CPNE is consistent with some participants bidding zero in stage one if they determine that the expected gains from bidding are not worth the cost. Of course, if everyone bid zero in stage one, they would be playing a different Nash equilibrium. Nothing like this ever happened in any period of any session. In the CPNE, while some participants may bid zero, many others burn out by bidding their entire 50-token endowment in stage one of the game. Since a bid of either zero or 50 is consistent with the burning-out CPNE, we define $EQDIST = \mathrm{Min}(50 - Bid, Bid - 0)$ as the dependent variable in a three-level hierarchical linear model.

The level-1 model is defined over time t for each individual participant i in each session s to account for convergence over the course of the game as:

$$EQDIST_{tis} = \pi_{0is} + \pi_{1is}(PER_t) + \varepsilon_{tis}, \tag{5}$$

where ε_{tis} is an observation-specific disturbance term. The level-2 model takes into account the possibility of individual-level random effects:

$$\begin{aligned} \pi_{0is} &= \beta_{00s} + \eta_{0is}, \\ \pi_{1is} &= \beta_{10s} + \eta_{1is}. \end{aligned} \tag{6}$$

The level-3 model introduces the session-specific treatment effects, which are now our primary focus of interest, as well as session-specific random effects:

$$\begin{aligned} \beta_{00s} &= \gamma_{000} + \gamma_{001}(NA_s) + \gamma_{002}(8P_s) + \mu_{00s}, \\ \beta_{10s} &= \gamma_{100} + \gamma_{101}(NA_s) + \gamma_{102}(8P_s) + \mu_{10s}. \end{aligned} \tag{7}$$

18 If the data from the professor treatment is added to the estimation of Eq. (4), the two-tailed p-value becomes 0.019 and all the other variables remain insignificant. When the insignificant variables are dropped the two-tailed p-value becomes 0.000.

Initially, we included interaction effects between NA, the no-announcement dummy, and 8P, the eight-person dummy at level 3. These effects were very far from significance and therefore dropped from the model. Combining Eqs. (5), (6), and (7), we estimate:

$$EQDIST_{tis} = \gamma_{000} + \gamma_{001}(NA_s) + \gamma_{002}(8P_s) + \gamma_{100}(PER_t) + \gamma_{101}(PER_t \times NA_s) + \gamma_{102}(PER_t \times 8P_s) + \eta_{0is} + \eta_{1is}(PER_t) + \mu_{00s} + \mu_{10s}(PER_t) + \varepsilon_{tis}. \quad (8)$$

Table 5 outlines the results. It is important to remember that there are eight periods in the game and that PER is defined as the period number minus eight. Thus, the estimated intercept and coefficients on both NA and 8P are calculated with respect to the last period. The intercept is equal to about 14.5 and highly significant ($p = 0.000$), indicating that in the last period of the four-person sessions with announcements, bids were about 14.5 tokens away from the burning-out CPNE. NA is insignificant, implying that whether or not there was an announcement made no difference to the distance from the burning-out CPNE in the last period. The insignificance of the interaction between PER and NA indicates that whether or not there was an announcement did not affect the speed of convergence to the CPNE either.

This result is in hindsight not particularly surprising. For example, suppose it is announced that the winning bids in stage one were 39 and 40. Then in the next period of play, we might expect the losers to bid very close to 39 and 40. The winners might maintain their bids or even reduce them. However, there is no compelling reason why announcements should induce the players to bid close to $B = 50$ in the next period. Notice that in a one-stage auction, the losers might bid more than 40 in the next period if they were informed that the winning bid was 40 in the previous period. In our game, this would be a very risky strategy, since bidding too high in stage one could leave one with too few resources to win the prize in stage two.

In contrast, 8P is negative and highly significant ($p = 0.000$), indicating that more players push bids significantly closer to the CPNE. The sum of $\gamma_{000} + \gamma_{002}$, which represents an estimate of the distance from the CPNE in the last period of the eight-person sessions, is insignificant, indicating that bids were very close to the burning-out CPNE in the eight-person case.

The coefficient on PER is not significant, implying that in the four-person games, there is no significant movement towards or away from the CPNE. However, the interaction between PER and 8P is negative and highly significant ($p = 0.009$), indicating that in the eight-person sessions the period-to-period movement towards the CPNE was significantly higher than in the four-person case. The sum of $\gamma_{100} + \gamma_{102}$, which represents that movement, is significant ($p = 0.001$)

Table 5
Distance from burning-out CPNE results

Independent variables	Estimate	t value	$Pr > \|t\|$
Intercept [γ_{000}]	14.544341	6.053	0.000
No announcement (NA) [γ_{001}]	0.425206	0.155	0.881
8 Participants (8P) [γ_{002}]	−15.001736	−5.464	0.000
Adjusted period (PER) [γ_{100}]	−0.159603	−0.474	0.646
PER × NA [γ_{101}]	−0.227421	−0.615	0.553
PER × 8P [γ_{102}]	−1.254216	−3.347	0.009
Other hypothesis tests			
$\gamma_{000} + \gamma_{002}$	−0.457395	−0.195	0.850
$\gamma_{100} + \gamma_{102}$	−1.413819	−4.577	0.001

Notes. Repeated measures three-level hierarchical linear model with random effect on intercept and adjusted period using full maximum likelihood. Equation estimated: Eq. (8).

Table 6
Summary of stage-two behavior in student sessions

Announcement	Stage-one winning bids	Both spend rest of endowment	One spends rest of endowment	None spend rest of endowment	Total
Yes	Tie	16	1	0	17
Yes	Difference = 1	9	5	1	15
Yes	Difference > 1	5	8	3	16
No	Two chosen randomly	6	1	1	8
No	One chosen randomly	5	1	0	6
No	No random draw	23	9	2	34

and equal to about −1.41, indicating that from period to period, bids moved about 1.41 tokens closer to the burning-out CPNE in the eight-person case.

How did participants behave in stage two? Table 6 summarizes stage-two bids in the student sessions. In all of the pure-strategy Nash equilibria, both participants who reach stage two after bidding identical amounts as required by all the pure-strategy equilibria in stage one, should bid all of their remaining endowments in the second stage. In 16 out of the 17 cases in which the announcement indicated that the two players entering stage two were tied in stage one, both players did in fact bid all of their remaining endowments in stage two as predicted. The professors did so in four out of four tied cases.

There were cases, however, in which the announcement revealed that the two participants entering stage two bid different amounts in stage one, despite the fact that such behavior is not part of a pure-strategy Nash equilibrium. Since in these cases the participants in the stage-two subgame have unequal caps, there is no pure-strategy equilibrium for the subgame, but only an equilibrium in mixed strategies (Che and Gale, 1997). While we did not set out to test the mixed-strategy equilibria of the one-stage all-pay auction in Che and Gale (1997),[19] we nevertheless wish to make a few comments on the stage-two experimental evidence. The mixed-strategy equilibrium in Che and Gale (1997) has the property that when two players face different caps, the player with the larger cap puts a positive mass on the smaller cap and distributes the remainder uniformly on $(0, B_2)$, where B_2 is the smaller cap. The other player puts a positive mass on zero and distributes the remainder uniformly on $(0, B_2)$. An immediate implication is that it is not part of a mixed-strategy equilibrium for both players to bid their caps in stage two, given that they are different. However, we found that of the 15 instances in which the announced winning stage-one bids differed by one token, both players bid the rest of their endowments nine out of 15 times. Out of the 16 instances in which the announced winning stage-one bids differed by more than one token, both players bid the rest of their endowments five out of 16 times. These results seem inconsistent with the mixed-strategy equilibrium in the stage-two subgame.

In the treatments where the successful bids were not announced, a participant moving on to stage two would only know his own stage-one bid and whether there had been zero, one, or two random draws. Since such draws were used only in the event of a tie for one or both of the two winning positions, the following inferences could be drawn. If there were two draws, three or more players must have been tied, requiring two draws to choose the two players who would advance to stage two. Thus, in this case, the two advancing players could determine that they must have bid identical amounts in stage one and thus have identical caps in stage two. This is consistent with all of the game's pure-strategy equilibria, each of which requires the advancing

[19] Rapoport and Amaldoss (2004) experimentally test a mixed-strategy equilibrium in a discrete all-pay auction.

players to bid the rest of their endowments in stage two. This actually occurred in six of the eight cases in which there were two draws. Behavior in the lab when there were fewer than two draws is summarized in the last two rows of Table 6. Note that stage-two behavior in this treatment cannot generally be tested based on the mixed-strategy equilibrium in Che and Gale (1997). This is because the size of a player's cap in stage two remains private information when there are no announcements, while in Che and Gale (1997), a player's cap is common knowledge. With no knowledge of the other person's cap, both players bid their caps in 28 out of the 40 such cases.

5. A cognitive hierarchy (CH) analysis[20]

In this section, we analyze the outcome of our game using a recently developed model of decision-making by boundedly rational agents due to Camerer et al. (2004).[21] As in Camerer et al. (2004), we assume that players think in steps. This captures the empirical fact that human beings have limited thinking capacities (e.g., Stahl, 1998). Also, as shown by Camerer et al. (2004), this idea explains some experimental data very well.

Suppose all players solve the game in k steps, where $k = 0, 1, \ldots, J$. If $k = 0$, a player does no thinking and hence does not behave strategically. As in Camerer et al. (2004), we assume that such players make their decisions by randomizing uniformly on some support. In particular, we assume that in stage one, such players randomize uniformly on $[0, b]$, where $0 < b < B$.[22] Each player doing k steps of thinking assumes that all other players do fewer than k steps of thinking. Hence each player is overconfident, believing that he/she is smarter than everyone else.[23]

As in Camerer et al. (2002, 2004), we assume that the actual distribution of thinking steps follows a Poisson distribution with a mean number, τ, of thinking steps. However, we do not use the Camerer et al. (2002, 2004) assumption that a k-step thinker believes that other players do 0 to $k-1$ steps of thinking according to a normalized Poisson distribution. Instead, we follow Nagel (1995), Stahl and Wilson (1995), and Ho et al. (1998) by assuming that a k-step thinker believes that all other players do exactly $k-1$ steps of thinking.[24] Players hold false beliefs. However, best response functions are correct, given those false beliefs.

Consider a 0-step thinker in stage one. He/she will randomize uniformly on $[0, b]$. Now consider a 1-step thinker who believes that all other players are 0-step thinkers. Then when he/she bids e in stage one his/her probability of successfully moving on to stage two, given $F = 2$, is

$$\rho_1 = \left(\frac{e}{b}\right)^{N-1} + (N-1)\frac{b-e}{b}\left(\frac{e}{b}\right)^{N-2}.$$

The first term is the probability that each of the other $N-1$ players bid less than e, and the second term is the probability that, out of $N-1$ players, $N-2$ players bid less than e and one

---

[20] We thank an associate editor for suggesting this line of research.

[21] This work builds on Nagel (1995), Stahl and Wilson (1995), Ho et al. (1998), and Costa-Gomes et al. (2001).

[22] Shortly, the need for $b < B$ will be obvious.

[23] Camerer et al. (2004) make a number of arguments in support of this overconfidence assumption. See also Binmore (1988) and Selten (1998).

[24] Using the normalized Poisson distribution of beliefs in Camerer et al. (2002, 2004) produces results identical to the ones derived here with a minor exception noted in footnote 29 below. The analysis under this assumption is available at http://www.uoguelph.ca/~jamegash/CH_normalized_poisson.pdf, or from the authors upon request. Camerer et al. (2002, fn.13) note that the $k-1$ assumption used here is "... easy to work with theoretically because the sequence of predicted choices can be computed by working up the hierarchy without using any information about the true distribution ..." We adopt it because it leads to empirically indistinguishable predictions, and enormously simplifies the exposition.

player bids more than e.[25] Clearly, it could be argued that there is more than one step of thinking in computing the probability above. So while we refer to this as 1-step thinking, we do so only in the sense that a higher-step thinker goes through more thinking steps than a 1-step thinker, or that a 1-step thinker best responds to the 0-step thinkers. Indeed, a 1-step thinker does even more strategic thinking by looking ahead to the outcome of the game in stage two. We assume that a 1-step thinker believes that his opponent, if he makes it to stage two, is a 0-step thinker who randomizes uniformly on the support $[0, \hat{B}]$,[26] where $\hat{B} = B - e$ is his opponent's cap in stage two. To find $\hat{B}$ as stage one commences, a 1-step thinker must compute the conditional density function that his/her eventual opponent in stage two will have emerged as the winner from the $(N-2)$ other 0-step thinkers by bidding $\tilde{e}$. This is the density of $\tilde{e}$, conditional on success in stage one. This conditional density function is the density function of the largest order statistic of the $(N-1)$ random variables. This gives $f(\tilde{e}|s) = (N-1)\tilde{e}^{N-2}/b^{N-1}$.

Recall that a 1-step thinker believes that his/her opponent, if he/she makes it to stage two, is a 0-step thinker who randomizes uniformly on the support $[0, B-\tilde{e}]$. From the standpoint of stage one, a 1-step thinker computes his/her success probability when he/she bids x in stage two as $\rho_2(\tilde{e}) = x/(B-\tilde{e})$. Since a player who is burnt out in stage two cannot randomize over any support, the only belief by a 1-step thinker and higher-step thinkers consistent with the belief that 0-step thinkers randomize in both stages is $b < B$. In what follows, we set $b = 49.99$. Hence the expected payoff of a 1-step thinker in stage two with valuation V_i is[27]

$$\Pi_{2i} = V_i \int_0^b \rho_2(\tilde{e}) f(\tilde{e}|s)\,\mathrm{d}\tilde{e} - x = x\left(\frac{V_i(N-1)}{b^{N-1}} \int_0^{49.99} \frac{\tilde{e}^{N-2}}{50-\tilde{e}}\,\mathrm{d}\tilde{e} - 1\right). \tag{9}$$

By setting $N \in \{4, 8\}$ and $b = 49.99$, we use the math software, Maple V, to compute the integral in (9). We find that the term in brackets is positive, given the values of $V_i \geqslant 100$ and $N \in \{4, 8\}$ used in our experiments and $b = 49.99$. Hence the optimal bid for a 1-step thinker in stage two is $\hat{x} = B - e$. This gives

$$\hat{\Pi}_{2i} = (B-e)\left(\frac{V_i(N-1)}{b^{N-1}} \int_0^{49.99} \frac{\tilde{e}^{N-2}}{B-\tilde{e}}\,\mathrm{d}\tilde{e} - 1\right). \tag{10}$$

Therefore, in stage one, a 1-step thinker chooses e to maximize $\Pi_1 = \rho_1 \hat{\Pi}_2 - e$. Let the optimal value be $\hat{e} = \hat{e}(N, V_i)$. We find that $0 < \hat{e} < B$ (i.e., is an interior solution) and is increasing in N and V_i. We obtain these results for $B = 50$, $V_i \geqslant 100$ and $N \in \{4, 8\}$. We check that second-order conditions for a local maximum hold.[28] We also look at graphical plots of Π_1, where e ranges over the interval $[0, B]$. We find that $\hat{e}$ is a unique global maximum. The results are summarized in Table 7. In general, $\hat{e}$ increases with V_i, and should thus differ among 1-step

[25] Note that a uniform distribution has no mass points, so the probability of a tie is zero.

[26] This is consistent with Camerer et al.'s remark (2004, p. 892) that "... extending the model to extensive-form games is easy by assuming that 0-step thinkers randomize independently at each information set, and higher-level types choose best responses at information sets using backward induction."

[27] For notational convenience, we suppress the i subscripts on the bids.

[28] There are three solutions to the first-order condition for an interior solution. Only one solution satisfies the second-order condition for a local maximum. Of the two remaining solutions, one solution is a minimum and the other solution violates the budget constraint (i.e., $e > B$).

Table 7
Values for $\hat{e}$, given $B = 50$, $N = 4$ or 8, and various values of V

V	$\hat{e}$ $(N = 4)$	$\hat{e}$ $(N = 8)$
100	28.5562	39.5947
120	28.6188	39.6166
150	28.6808	39.6382
170	28.7097	39.6484
200	28.7422	39.6597
230	28.7661	39.6681
270	28.7896	39.6764
300	28.8031	39.6811
400	28.8333	39.6918
500	28.8514	39.6981
600	28.8634	39.7024

thinkers with different valuations. However, the predicted differences are so small as to be inconsequential in our experiments where only integer bids were allowed. When $N = 4$, $\hat{e}$ rounds off to 29, while for $N = 8$, it rounds off to 40 for all valuations used in the experimental sessions.

What happens when N becomes very large? Given $f(\tilde{e}|s) = (N-1)\tilde{e}^{N-2}/b^{N-1}$, the expected highest bid among the $(N-1)$ randomizers in stage one, from the standpoint of a 1-step thinker, is $E(\tilde{e}|s) = [(N-1)/N]b$. Notice that $E(\tilde{e}|s) \to b$ as $N \to \infty$. Also, the expected second-highest bid among the $(N-1)$ randomizers can be shown to be $\hat{E}(\tilde{e}|s) = [(N-2)/N]b$. Again, $\hat{E}(\tilde{e}|s) \to b$ as $N \to \infty$. Hence, when N is large, a 1-step thinker requires a bid very close to b in stage one to be successful. Therefore, $\hat{e}(N, V) \to b$ as $N \to \infty$. Since $b = 49.99 \approx B$, it follows that the 1-step thinkers almost burn out when N is very large.

Now consider a 2-step thinker. Note that a 2-step thinker knows that a 1-step thinker bids $\hat{e}$ in stage one and $B - \hat{e}$ in stage two. Since a 2-step thinker believes that all other players are 1-step thinkers, his optimal response is also to bid $\hat{e}$ in stage one and $B - \hat{e}$ in stage two.[29] Similarly, all higher-step thinkers will bid $\hat{e}$ in stage one. We summarize our analysis in Proposition 3:

Proposition 3. *Consider a two-stage contest in which the contest in each stage is an all-pay auction where the ith contestant has valuation, V_i, contestants have a common budget constraint, B, and behave according to the cognitive hierarchy* (*CH*) *step model of thinking, $i = 1, 2, \ldots, N$. Then* (i) *there exists a pure-strategy CH outcome in which* 1*-step thinkers and higher-step thinkers with valuation V_i bid $\hat{e} = \hat{e}(N, V_i) \leqslant B$ in stage one and $\hat{x} = B - \hat{e}$ in stage two, where $\hat{e}$ is increasing in N and V_i, and* (ii) $\hat{e}(N, V_i) \to B$ *as* $N \to \infty$.

We wish to emphasize that the predicted CH bid $\hat{e}$ is increasing in the number of players, N, as indicated in Table 7. Note that this is not the case in all-pay auctions with non-boundedly rational players (Baye et al., 1996; Che and Gale, 1997). This result is however consistent with Anderson et al. (1998) who, using a quantal response equilibrium model in which players choose best response functions stochastically, show that aggregate expenditure is increasing in the number of players in a one-stage all-pay auction. Moldovanu and Sela (2001, 2006) also obtain a similar result in an all-pay auction where a player's ability is private information and is assumed by

[29] If τ is sufficiently high, for example $\tau \geqslant 0.4$ for $N = 8$ we can show that, using the normalized Poisson distribution of beliefs in Camerer et al. (2004), a 2-step thinker will bid $\hat{e} + 1$ and all other higher-step thinker will bid likewise. The proof is available at http://www.uoguelph.ca/~jamegash/CH_normalized_poisson.pdf.

his opponents to be a random variable that is drawn from some distribution. It is interesting to note that the number of players has no effect on individual or aggregate bids in all-pay auctions with complete information, mutually consistent beliefs, and unboundedly rational players. This suggests that some exogenous randomness either in the bidding behavior of the players as in the present paper and Anderson et al. (1998), or in some player-specific parameter (e.g., ability or valuation) as in Moldovanu and Sela (2001, 2006) is required to obtain a relationship between the number of players and bids in all-pay auctions. Similarly, using the CH model, Camerer et al. (2004) also find a group size effect in the predicted outcome of the "stag hunt" game consistent with experimental evidence, while Nash equilibrium makes no such prediction. An important observation is that our claim that players almost burn out when N is sufficiently large requires that a sufficiently high proportion of players be non-0-step thinkers. It is the 1- and higher-step thinkers who burn out when N is high, while 0-step thinkers always place their bids randomly.[30]

5.1. Estimation and model comparison

Since our game has a multiplicity of strategies available to the players, estimating τ to fit the data by maximum likelihood would be inefficient (Camerer et al., 2002, pp. 23–26). Therefore, we follow Camerer et al. (2002, 2004) in choosing τ such that the predicted mean bid is close to the actual mean bid in the data. Given a Poisson distribution of thinking types, the probability mass function of k-step thinkers is $g_k = \tau^k \exp(-\tau)/k!$. Restricting bids to integer amounts and given the Poisson CH model, we predict that a proportion, $g_0 = 1/\exp(\tau)$, of the players will randomize on $[0, 49]$ and the rest will bid $\hat{e}$. For $N = 8$, the predicted mean is $\bar{e}_8 = 24.5\exp(-\tau) + 40(1 - \exp(-\tau))$. For $N = 4$, the predicted mean is $\bar{e}_4 = 24.5\exp(-\tau) + 29(1 - \exp(-\tau))$.

To fit the data, we focus our analysis on the first two periods of each treatment. This allows us to abstract from possible learning effects that may take place as players receive feedback between periods. Note that our CH analysis does not take account of learning. The development of a dynamic CH-learning model is beyond the scope of this paper.

The τ estimates for the first period are presented in Table 8. In two out of ten cases in our sample, the mean bid was greater than $\hat{e}$. However, as τ approaches infinity, the mean bid predicted by the CH model approaches $\hat{e}$ from below. Since there is no τ that predicts the observed mean bids in these cases, we did not provide a τ estimate.[31] In ten out of 12 cases, the estimated $\tau \in (0, 1)$.[32] These results imply that the mean number of thinking steps in period one was very low. Most players were apparently 0-step thinkers. This is not surprising in a dynamic game like ours where the players have to figure out the equilibrium via backward induction.[33] The average of the sample means for $N = 4$ and $N = 8$ were 25.04 and 24.77 respectively. This is also not

[30] Note that the normalized Poisson belief assumption in Camerer et al. (2004) is equivalent to the $k - 1$ assumption as $\tau \to \infty$. Camerer et al. (2004) shows that as $\tau \to \infty$, the prediction of the Poisson CH model converges to one of the Nash equilibria, when a Nash equilibrium is reached by iterated deletions of weakly dominated strategies. In our model, almost everyone bids $\hat{e}$, if τ is very large. It is interesting to note that we obtain convergence to a Nash equilibrium although, unlike Camerer et al. (2004), our Nash equilibria are not reached by such iterated deletions.

[31] If τ were really close to infinite in these cases, implying belief consistency, we would expect the standard deviations of the bids to be close to zero. However, they were not. Two similar cases arise in period two. In the $N = 8$ case (session 11), the standard deviation is quite low and the distribution of bids appears very close to a Nash equilibrium.

[32] We do not fit the CH model to the professor treatment because we believe the Nash equilibrium is more applicable to the behavior of the professors. They were not randomizing. Zero bids were almost all associated with low prize values, and non-zero bids were almost all close to 20.

[33] See Johnson et al. (2002).

Table 8
Period-one data and CH estimates of τ in various samples

Sample	Winning bids announced	Sample size	Sample mean in period one	Sample Std Dev period one	*Estimated* τ	Predicted mean from CH model	CH Std Dev
1	No	4	24.75	11.5	*0.06*	24.75	13.79
2	No	4	33.50	5.80	*N/A*	*N/A*	N/A
3	No	4	26.50	14.34	*0.59*	26.50	10.78
4	Yes	4	20.50	4.20	*0.00*	24.50	14.14
5	Yes	4	33.25	7.89	*N/A*	N/A	N/A
6	Yes	4	11.75	6.24	*0.00*	24.50	14.14
7	No	8	32.25	13.19	*0.69*	32.25	12.65
8	No	8	25.13	13.53	*0.04*	25.13	14.18
9	No	8	26.63	4.37	*0.15*	26.63	14.17
10	Yes	8	15.75	16.45	*0.00*	24.50	14.14
11	Yes	8	28.00	17.70	*0.26*	28.00	14.02
12	Yes	8	20.88	14.53	*0.00*	24.50	14.14

Notes. The predicted variances $= (1/\exp(\tau))(\sigma_0^2 + (\bar{e}_0)^2) + (1 - 1/\exp(\tau))(\sigma_1^2 + (\hat{e})^2) - (\bar{e})^2$, where $\sigma_0^2 = (49-0)^2/12$ is the variance of the bid of a 0-step thinker, $\sigma_1^2 = 0$ is the variance of the bids of all higher-step thinkers, $\bar{e}_0$ is the predicted mean bid of 0-step thinkers, and $\bar{e}$ is the predicted mean bid for all players. N/A indicates that the mean bid fell above $\hat{e}$, which is inconsistent with any τ.

Table 9
Period-two data and CH estimates of τ in various samples

Sample	Winning bids announced	Sample size	Sample mean in period two	Sample Std Dev period two	*Estimated* τ	Predicted mean from CH model	CH Std Dev
1	No	4	20.00	13.88	*0.00*	24.50	14.14
2	No	4	27.75	13.55	*1.28*	27.75	7.72
3	No	4	21.00	14.31	*0.00*	24.50	14.14
4	Yes	4	22.75	1.89	*0.00*	24.50	14.14
5	Yes	4	36.00	4.69	*N/A*	N/A	N/A
6	Yes	4	19.75	3.40	*0.00*	24.50	14.14
7	No	8	34.88	18.15	*1.12*	34.88	10.92
8	No	8	31.50	14.8	*0.60*	31.50	13.01
9	No	8	32.00	13.46	*0.66*	32.00	12.77
10	Yes	8	21.13	19.35	*0.00*	24.50	14.14
11	Yes	8	47.63	1.41	*N/A*	N/A	N/A
12	Yes	8	33.88	14.44	*0.93*	33.88	11.68

See notes to Table 8.

surprising because most of the players were randomizing in this period, and hence the average bid should be very close to 24.5 regardless of sample size. Camerer et al. (2004) also obtained very low estimates of τ in beauty contest games. As in Camerer et al. (2004, Table II), the predicted standard deviations were not very far away from the actual standard deviations, although the τs were not chosen to match the standard deviations.

The τ estimates for the second period are presented in Table 9. The average of sample means for $N = 4$ is 24.54, while it is 33.52 for $N = 8$. In four of six cases, the estimated τ for $N = 4$ is 0. In four out of six cases, the estimated τ for $N = 8$ is at least 0.60. It would seem that greater competition when $N = 8$ than when $N = 4$ motivates players to think harder, resulting in a more sophisticated understanding of the beliefs and expected behavior in the former case after one period of play. The higher level of τ, together with the higher $\hat{e}$ in the eight-person treatments, lifts the bids in these treatments above those in the four-person treatments.

Although the CH model makes a qualitative prediction that the bids of strategic players, $\hat{e}$, will increase with V, Table 7 indicates that given the integer bidding rule, predicted $\hat{e}$ rounds off to 29 for all valuations when $N = 4$, and to 40 when $N = 8$. Bids of non-strategic 0-step thinkers are chosen randomly, and should thus also be unrelated to V. Thus, although it has already been shown that valuation has a substantial effect on whether to bid at all, it is interesting to examine whether bid level is unrelated to valuation once a participant decides to place a non-zero bid. We ran separate OLS regressions of bid level on prize value for the four-person and eight-person treatments for each of the first two periods. The results, reported in Table 10, show that active bid level is positively related to valuation in period one ($p = 0.078$ for $N = 4$, $p = 0.003$ for $N = 8$), albeit with marginal significance for the four-person case. This relationship becomes insignificant by period two. It would seem that players are influenced by their valuations when choosing their initial bids. However, once they learn how their bids fare among the others placed in period one, they mark up or down their period-one bids in period two without taking their valuations into account.

In subsequent periods, as indicated in the statistical results reported earlier, active bids moved progressively higher in the eight-person treatments relative to the four-person treatments. Although this may be partially explained by both higher τs and a higher $\hat{e}$ in the eight-person treatments, the CH model cannot fully explain the mean bid levels that result. From period three through eight, the mean bid for active players[34] was greater than $\hat{e} = 40$ in the eight-person treatments fully 94.4% of the time, validating our decision not to apply the CH model and estimate τ in the these periods. For the four-person treatments the corresponding percentage of mean active bids above $\hat{e} = 29$ was 61.1%. Some players bid zero or one in these latter periods. As previously discussed, such bids were significantly more common when low valuations were drawn. Thus, they likely represent situations were the expected payoff was not high enough to compensate for the risk of bidding as opposed to randomly-chosen bids by 0-step thinkers.

Standard deviations of active bids generally fell as the game progressed, particularly in the eight-person treatments as exemplified in the multi-period results displayed in Figs. 1 to 5 and the final period mean bids and standard deviations for active players reported in Table 3. In five

Table 10
OLS regressions of bid level of active (non-zero) bids on prize value

Period and number of players	Independent variables	Estimate	t value	$Pr > \|t\|$
Period 1, $N = 4$	Intercept	14.179	2.281	0.033
	Prize value	0.0539	1.848	0.078
Period 2, $N = 4$	Intercept	24.709	4.529	0.000
	Prize value	0.0099	0.399	0.694
Period 1, $N = 8$	Intercept	12.086	2.317	0.026
	Prize value	0.0371	3.150	0.003
Period 2, $N = 8$	Intercept	35.646	6.634	0.000
	Prize value	0.0042	0.346	0.731

[34] As explained in a previous footnote, we defined active players as those bidding more than one. If active players are defined more stringently as those who bid more than zero, the percentage of mean bids above $\hat{e}$ becomes 88.9% in the eight-person treatments and remains at 61.1% in the four-person treatments.

of the six eight-person cases, the standard deviation was less than one. In two of the three eight-person sessions with announcements, all active players burned out by bidding exactly 50. Thus, in many of our samples, randomizing behavior seems much reduced and convergence towards a Nash equilibrium with consistency of beliefs seems to be occurring by the end of the game. Although the CH model predicts burning out as the number of players becomes very large, it does not do so for $N = 8$, where $\hat{e}$ is predicted to be 40.

6. Burning-out in the eight-player treatment: a formal CPNE explanation

As argued above, we believe that the CH model is applicable to period one and perhaps period two of each treatment, but has less applicability in later periods. Since burning out in the lab was observed in later periods, we shall examine the difference in behavior for $N = 8$ versus $N = 4$ by returning to the CPNE model. The effect of the number of players on the likelihood of burning out does not support the CPNE predictions, since these equilibria are independent of the number of players. Nonetheless, the number of players could affect the likelihood of burning out.

Consider any Nash equilibrium that is not a CPNE. These are the non-burning out equilibria. Suppose p is a player's subjective probability that an opponent in stage one will deviate to a higher bid.[35] We assume that a player holds this subjective belief at any non-burning-out equilibrium. Let d be the number of deviators. Then, of the $N - 1$ other players, the probability that d players will deviate is

$$\text{prob}(d) = \frac{(N-1)!}{d!(N-1-d)!} p^d (1-p)^{N-1-d}.$$

If $d \geqslant 2$, the probability that a non-deviating player will advance to stage two is zero, given $F = 2$. If $d < 2$, the probability that a non-deviating player will advance to stage two is $(F - d)/(N - d) = (2 - d)/(N - d)$. Therefore, the probability that a non-deviating player will advance to stage two is

$$\text{prob}(\text{adv}) = \sum_{d=0}^{1} \frac{2-d}{N-d} \frac{(N-1)!}{d!(N-1-d)!} p^d (1-p)^{N-1-d}.$$

Now suppose that each player believes that there is a 50% chance that a randomly chosen opponent will deviate, i.e. $p = 0.5$. Then $\text{prob}(\text{adv} \mid N = 4,\ p = 0.5) = 0.1875$ and $\text{prob}(\text{adv} \mid N = 8,\ p = 0.5) = 0.00977$. So, when $N = 8$, a player has a very small chance (i.e., 0.977%) of advancing to the next stage if he/she does not deviate to a higher bid. In contrast, when $N = 4$, a player has a much higher chance (i.e., 18.75%) of advancing if he/she does not deviate. Since these probabilities are the same at any non-burning-out equilibrium and a player does not know by how much other deviators have increased their bids, it is reasonable that if a player decides to deviate, he/she should deviate to the burning-out equilibrium. Since, if a player does not deviate, the probability of advancing when $N = 8$ is almost 1/20 of the probability of advancing when $N = 4$, it is reasonable to argue that a player is more likely to deviate from any non-burning out equilibrium, when $N = 8$ than when $N = 4$.[36] The higher is N, the higher is the probability that two or more of the other players will deviate to a higher bid. With only two slots and eight

[35] Notice that it does not make sense to deviate to a lower bid.

[36] This ratio is smaller than 1/20 for $0.5 < p < 1$. Holding N fixed at 4 or 8, the probability of advancing is decreasing in p. If one were to assume that p is an increasing function of N, this would strengthen our result since $p(8) > p(4)$.

contestants, a player feels that he/she has to bid more to get to stage two when he/she has to beat six out of seven other players rather than two out of three players. Aggressive bidding may stem from a player's belief that deviations to higher bids are more likely with more than with fewer players.

7. Conclusion

We have examined a two-stage sequential elimination game with a continuum of Pareto-rankable equilibria. The CPNE refinement rules out all equilibria except a Pareto-dominated burning out equilibrium in which players spend their entire budget in stage one. Our first finding is that some players withdraw from the game by bidding zero, while others bid substantial amounts. The probability of withdrawal is inversely related to valuation as our model predicts.

Our second finding is that a Pareto-dominant equilibrium is never attained in any of our sessions. Moreover, our experimental results show that bids are higher and that burning out is more likely to occur with more players. This somewhat parallels a recent experimental finding in Amaldoss and Rapoport (2005), where in a quite different context, bids are higher when there are more players in the first stage of a two-stage game. However, in Amaldoss and Rapoport (2005), such behavior is not consistent with any equilibrium, and first-stage bids are higher than predicted by Nash both with smaller and larger numbers of players in the first stage. The puzzle in Amaldoss and Rapoport (2005) is that bids are higher than predicted in both of their treatments. In contrast, our puzzle is that bids are lower than predicted by CPNE in our four-person treatment. Amaldoss and Rapoport (2005) propose that people get non-pecuniary utility, increasing with the number of competing players, from winning an auction, and show that this can explain much of the behavior they observe. In our context, unrealistically high levels of such utility would be required to explain the differences in our four- and eight-person results.

Third, we show that a cognitive hierarchy model predicts burning out when the number of players and the mean number of thinking steps are both large. However, it does not explicitly predict burning out in either of our treatments. Nonetheless, the idea emanating from the CH analysis that more players competing for the same number of spots means a higher probability of randomizers choosing very high bids, which in turn causes the bids of strategic players to be higher, is likely a driving force behind our burning-out results. In addition, our results suggest that with more competition for spots, people learn more quickly to think in a more sophisticated manner. Camerer et al. (2004) argue that high stakes encourage more sophisticated thinking. Apparently, more competition does as well.

Fourth, we show that CPNE is not a good predictor of behavior when four people compete for two second-stage spots, but does predict well when eight people compete for the two available spots. Allowing for joint deviations, we provide a formal analysis as to why CPNE is likely to have more predictive power in the eight-person case. With more players, there is a higher probability that two will deviate to a higher bid, leading to the breakdown of any equilibrium that is not coalition-proof, and convergence towards the unique CPNE equilibrium.

More competing players imply: a higher probability that two randomly-chosen high bids will be placed by 0-step thinkers; more strategic-thinking motivated by more competition leading to higher τs; and a higher probability that two competitors might deviate to a higher bid. These factors all suggest that more competition leads to higher bids, and that burning out is thus a competitive phenomenon.

Future research should focus on learning across periods and its interaction with the number of players. For example, the experience-weighted-attraction learning model proposed by Camerer and Ho (1999) might be employed to examine how bids adjust with experience. It would be easier to fit such a model to data from an experimental setting in which players had common valuations. It would be especially interesting to run experimental sessions for more than eight periods to determine whether four-player sessions would exhibit burning out with time. Similarly, it would be interesting to examine whether players who attain a burning-out equilibrium would stay put or ultimately attempt to attain a more efficient outcome. It would also be informative to examine an experimental treatment in which eight players compete for four spots in stage two and to investigate burning out in games with more than two stages.

Acknowledgments

We thank an associate editor and an anonymous referee for very helpful comments. We also thank Asha Sadanand, Gregory Besharov, Phil Curry, Claudia Keser, Talat Genc, Gordon Myers, Marco Runkel and participants in the Economic Science Association meetings in Tucson in October 2003 and the Canadian Economics Association meetings in Toronto in June 2004 for helpful comments. Cadsby would also like to acknowledge with gratitude funding from the Social Sciences and Humanities Research Council of Canada, grant #410-2001-1590.

References

Amaldoss, W., Rapoport, A., 2005. Excessive expenditure in two-stage contests: Theory and experimental evidence. Unpublished manuscript. Duke University.

Amegashie, J.A., 2004. Burning out in sequential-elimination contests. Unpublished manuscript. http://www.uoguelph.ca/~jamegash/Burning_out.pdf.

Anderson, S.P., Goeree, J.K., Holt, C.A., 1998. Rent seeking with bounded rationality. J. Polit. Economy 106, 828–853.

Anderson, S.P., Goeree, J.K., Holt, C.A., 2001. Minimum effort coordination games: Stochastic potential and logit equilibrium. Games Econ. Behav. 34, 177–199.

Baye, M.R., Kovenock, D., de Vries, C.G., 1996. The all-pay auction with complete information. Econ. Theory 8, 291–305.

Bernheim, B.D., Peleg, B., Whinston, M.D., 1987. Coalition-proof Nash equilibria I: Concepts. J. Econ. Theory 42, 1–12.

Berninghaus, S.K., Erhart, K.-M., Keser, C., 2002. Conventions and local interaction structures: Experimental evidence. Games Econ. Behav. 34, 177–205.

Binmore, K., 1988. Modelling rational players: Part II. Econ. Philosophy 4, 9–55.

Camerer, C., 2003. Behavioral Game Theory: Experiments in Strategic Interaction. Russell Sage Foundation/Princeton Univ. Press, New York/Princeton.

Camerer, C., Ho, T., 1999. Experience-weighted attraction learning in games: Estimates from weak-link games. In: Budescu, D., Erev, I., Zwick, R. (Eds.), Games and Human Behavior: Essays in Honor of Amnon Rapoport. Lawrence Erlbaum, Mahwah, NJ, pp. 31–52.

Camerer, C., Knez, M., 1994. Creating "expectational assets" in the laboratory: "Weakest link" coordination games. Strategic Manage. J. 15, 101–119.

Camerer, C.F., Ho, T.-H., Chong, J.-K., 2002. A cognitive hierarchy theory of one-shot games: Some preliminary results. http://www.hss.caltech.edu/~camerer/thinking2002.pdf.

Camerer, C.F., Ho, T.-H., Chong, J.-K., 2004. A cognitive hierarchy model of games. Quart. J. Econ. 119, 861–897.

Che, Y.K., Gale, I., 1997. Rent dissipation when rent seekers are budget constrained. Public Choice 92, 109–126.

Che, Y.K., Gale, I., 1998. Caps on political lobbying. Amer. Econ. Rev. 88, 643–651.

Clark, D.J., Riis, C., 1998. Competition over more than one prize. Amer. Econ. Rev. 88, 276–289.

Costa-Gomes, M., Crawford, V., Broseta, B., 2001. Cognition and behavior in normal-form games: An experimental study. Econometrica 69, 1193–1235.

Dufwenberg, M., Gneezy, U., 2002. Information disclosure in auctions: An experiment. J. Econ. Behav. Organ. 48, 431–444.

Garratt, R., Parco, J.E., Qin, C.-Z., Rapoport, A., 2005. Potential maximization and coalition formation. Int. Game Theory Rev. 7, 407–429.

Gavious, A., Moldovanu, B., Sela, A., 2002. Bid costs and endogenous bid caps. RAND J. Econ. 33, 709–722.

Gillette, A.B., Noe, T.H., Rebello, M.J., 2003. Corporate board composition, protocols, and voting behavior: Experimental evidence. J. Finance 58, 1997–2031.

Gillette, A.B., Noe, T.H. Rebello, M.J., 2004. Board structures around the world: An experimental investigation. Working paper. Tulane University.

Harsanyi, J.C., Selten, R., 1988. A General Theory of Equilibrium Selection in Games. MIT Press, Cambridge, MA.

Ho, T.-K., Camerer, C.F., Weigelt, K., 1998. Iterated dominance and iterated best response in p-beauty contests. Amer. Econ. Rev. 88, 947–969.

Johnson, E.J., Camerer, C., Sen, S., Rymon, T., 2002. Detecting failures of backward induction: Monitoring information search in sequential bargaining. J. Econ. Theory 104, 16–47.

Moldovanu, B., Sela, A., 2001. The optimal allocation of prizes in contests. Amer. Econ. Rev. 91, 542–558.

Moldovanu, B., Sela, A., 2006. Contest architecture. J. Econ. Theory 126, 70–97.

Nagel, R., 1995. Unravelling in guessing games: An experimental study. Amer. Econ. Rev. 85, 1313–1326.

Neugebauer, T., Selten, R., 2003. Individual behavior in first-price sealed-bid auctions: The importance of feedback in experimental markets. Discussion paper No. 3/2003. University of Bonn.

Ockenfels, A., Selten, R., 2005. Impulse balance equilibrium and feedback in first price auctions. Games Econ. Behav. 51, 155–170.

Parco, J.E., Rapoport, A., Amaldoss, W., 2005. Two-stage contests with budget constraints: An experimental study. J. Math. Psych. 49, 320–338.

Rapoport, A., Amaldoss, W., 2004. Mixed-strategy play in single-stage first-price all-pay auctions with symmetric players. J. Econ. Behav. Organ. 54, 585–607.

Raudenbush, S.W., Bryk, A.S., 2002. Hierarchical Linear Models: Applications and Data Analysis Methods, second ed. Sage, Thousand Oaks, CA.

Schelling, T., 1960. The Strategy of Conflict. Harvard Univ. Press, Cambridge, MA.

Selten, R., 1998. Features of experimentally observed bounded rationality. Europ. Econ. Rev. 42, 413–436.

Snijders, T., Bosker, R., 1999. Multilevel Analysis: An Introduction to Basic and Advanced Multilevel Modeling. Sage, London.

Stahl, D.O., 1998. Is step-j thinking an arbitrary modeling restriction or a fact of human nature? J. Econ. Behav. Organ. 37, 33–51.

Stahl, D.O., Wilson, P., 1995. On players' models of other players: Theory and experimental evidence. Games Econ. Behav. 10, 213–254.

Van Huyck, J.B., Battalio, R.C., Beil, R., 1990. Tacit cooperation games, strategic uncertainty, and coordination failure. Amer. Econ. Rev. 80, 234–248.

Van Huyck, J.B., Battalio, R.C., Beil, R., 1991. Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. Quart. J. Econ. 106, 885–909.

Van Huyck, J.B., Battalio, R.C., Rankin, F.W., 2001. Evidence on learning in coordination games. Unpublished manuscript. Texas A & M University.