Level-K and Iterated Reasoning in a Bertrand Price Competition Experiment

Brian O. Buck

Submitted to the Distinguished Majors Program Department of Economics University of Virginia April 2024 Advisor: Charles Holt

Level-K and Iterated Reasoning in a Bertrand Price Competition Experiment

Brian O. Buck

Abstract

This paper considers cognitive hierarchy theory in the context of a Bertrand price competition model. A laboratory experiment is run with a multi-round price competition Bertrand Game, varying the marginal cost and range of feasible prices between eleven groups of students. The Level-K Cognitive Hierarchy Model is used to evaluate pricing behavior, both in the first round and in subsequent rounds with random matching. The main finding is that that Level-K models fail to provide useful predictions in simple, multi-round Bertrand Price Competition Games, unless the range of feasible prices is restricted to a narrow range. More precisely, the "level" of iterated reasoning that is needed to explain behavior with a wide range of prices differs from the level that explains behavior with a narrow range. While behavior is dependent on the range of prices, this paper finds preliminary evidence that subjects given broader pricing options artificially narrow their choices of iterated reasoning to mimic behavior observed under a narrower range.

Table of Contents	4
Introduction	4
Level-K Theory	6
Theoretical Predictions	9
Bertrand Price Competition Nash Equilibrium	9
Prediction for a Level 0 Player	10
Experimental Design	11
Procedures	14
Data	15
Results	16
Baseline Treatment	
Price Floor Treatment	
High Marginal Cost Treatment	19
\$10 Upper Limit Treatment	21
\$50 Upper Limit Treatment	
Level-K Analysis Between Treatments	
Level-K Predictions in Multi-Round Bertrand Price Competition Games	
Comparing Results in the \$10 Upper Limit and the \$100 Upper Limit Treatments	
Conclusion	
References	33
Appendix	34
Appendix A. – Average Price per Round in Each Treatment by Session	34
Appendix B. – Level 0 Prediction Derivation for Bertrand Price Competition	37
Appendix C. – Bertrand Price Competition Lab Experiment Instructions	
Appendix D. – Holt-Laury Risk Aversion Test Instructions	40

Introduction

Laboratory experiments are often used to help economists isolate why individual behavior differs from theoretical predictions. A key economic model where behavior generally deviates from the equilibrium predictions is in Bertrand price competition games. Bertrand price competition was first recognized by French economist Antoine Agustin Cournot who gained recognition for deriving a model of firm competition based on quantities produced. In his work, Cournot alluded to a form of competition based on setting prices marginally below your competitor. This form of price competition was later formalized by Joseph Louis François Bertrand in his review of Cournot's *Recherches sur les Principes Mathématiques de la Théorie des Richesses* (Bertrand 1883). Bertrand's research of firm competition based on setting prices would later be known as Bertrand price competition, the subject of countless economic papers.

The simplest design of a Bertrand price competition game involves two individuals, producing homogeneous goods who compete to sell their goods in the market with identical production costs. Each firm hopes to acquire the highest achievable profit by setting their price high but must not set a price too high such that consumers would not purchase goods from that firm. Each firm sets their price without knowing the decision of their competitor. Once prices are selected, the seller with the lowest price sells their goods at their selected price. The competing firm who set a higher price will be excluded from the market, failing to sell any goods earning them zero profit. In a scenario of prefect information, the best strategy is for each firm to set their price just below the price of their competitor. This ensures that consumers buy their goods from that firm while making the highest possible profit. Yet, given that each firm knows the pricing decision of the other firm, each has an incentive to continually undercut the price their opponent sets as failing to have the lowest price will lead to no goods sold equaling no profit. As both

competitors continue to undercut one another in price, the Nash equilibrium predicts the price will converge to the marginal cost. Any price below marginal cost will cause the firm to lose money and any price above marginal cost, assuming that both firms are setting their price equal to marginal cost, will lead to no sales. This is the Bertrand Paradox where even with two firms competing in price, they will converge to the competitive prediction of setting price equal to marginal cost. Given a simple Bertrand price competition game, the theoretical prediction is that in a one-shot game, both competitors will set their price at marginal cost.

While the Bertrand Model provides insights into price competition in theory, often market interactions persist over many periods with price setting behaviors deviating from theoretical predictions. To better understand pricing decisions in Bertrand competition by financially motivated individuals, this paper uses laboratory experiments to applies the Level-K Cognitive Hierarchy model as a possible explanation for deviations in behavior from Nash predictions. This paper explores initial pricing decisions of financially motivated subjects in a laboratory setting, testing whether Level-K Cognitive Hierarchy Model explains iterated reasoning observed in laboratory experiments.

As described above, the simplest version of the Bertrand Price Competition Game consists of two individuals, producing homogenous goods, with identical, constant marginal costs, selling to a market consisting of buyers in one period. The experiment discussed in this paper maintains most of the properties of the simple Bertrand game, varying the marginal cost and pricing bounds in five treatments and introducing multiple periods to what is traditionally a one-shot game. A multi-round Bertrand price competition game allows subjects to learn from previous rounds to inform their next decision. Rather than relying solely on their intuition to predict their opponent's price, subjects use information on pricing strategy of their opponents observed in prior rounds a to select the profit maximizing price in the current period. This paper is primarily concerned with price selections made in the first round emulating a one-shot game. A future research topic using this data may include introducing a learning component to Level-K iterated reasoning model, but this paper does not directly address learning and iterated reasoning in multi-round games.

Laboratory experiments modeling Bertrand price competition find that subjects, almost universally, do not follow the Nash Equilibrium of setting price equal to their marginal cost in the first round. Instead, subjects select prices above the marginal cost and as they participate in additional rounds, tend to converge toward the Nash prediction after enough periods. This paper compares Level-K behavioral predictions in Bertrand price competition experiment to test the conditions under which a Level-K iterated reasoning provides useful predictions in one-shot and multi-shot treatments of Bertrand price competition game.

Level-K Theory

The theory focused on in this paper attempting to explain the behavioral differences between Nash Equilibrium and laboratory experiments is Level-K Cognitive Hierarchy Model, introduced by Dale Stahl (Stahl 1993). Level-K Cognitive Hierarchy Theory breaks down the iterated reasoning an individual undertakes when they determine the best pricing strategy given no information about their opponent. Instead of randomly choosing a price, Level-K theory posits that individuals, given an optimal strategy, estimate the sophistication that they expect their opponent to choose based on the optimal strategy. Given the expected strategy from their opponents, individuals will choose one level of sophistication above all other opponents. By selecting one level of sophistication that they suspect is above all other players, they ensure that they will win if their predictions about their competitors level of sophistication are true. In the case of Bertrand Price Competition, sophistication equates to the pricing decisions of an opponent. If a firm suspects that their opponents are sophisticated, they may want to set a lower price to ensure they capture the market. If a firm suspects their competitor is not sophisticated, they will set a higher price to make a larger profit without worrying that their opponent will undercut. Level-K helps breakdown and expand out the possible prices that your opponent could select in terms of "levels" that an individual might think through when determining their best pricing strategy.

It will be useful to precisely describe Level-K iterated reasoning in the context of the experiments discussed in this paper. It is helpful to start from the perspective of a seller selecting a price for their goods sold with no prior information about the price setting behavior of their opponent. If a firm suspect that their opponent has no sophistication where they choose a price at random, given the ability to choose a price between and including \$0 and \$100 with a uniform probability distribution, they would be expected, on average, to choose a price of \$50.¹ Level-K describes individuals acting randomly to be "Level 0" players (L0). If a seller suspects that they are playing against all L0s, their best response, given the ability to only choose integer values, would be to select a price of \$49. Level-K labels an individual choosing a best response to a L0 player as "Level 1" individuals (L1). If a firm suspects that you are playing against an opponent at L1, your best response would be to choose a price of \$48, meaning they are a Level 2 individual (L2). This continues until the Nash Equilibrium and in the case of a Bertrand Competition with the parameters above, the maximum level of sophistication is a player playing the Nash Equilibrium of \$0 at level 50.

¹ See Appendix B. for derivation.

Another feature of Level-K iterated reasoning model is the assumption by each player believes they are playing the "most sophisticated" strategy (Nagel 1995). Given that each player has no information on the behavior of their opponents, each must make a decision about the maximum level of sophistication they believe their competitors may play. For instance, if individual A believes that individuals B and C are both playing a Level 1 strategy, individual A's best response would be to play a Level 2 Strategy. Even if an additional individual were to play a Level 0 strategy, this would not change the best response of individual A as they would need to select an L2 strategy to beat individuals B and C. This means that each individual, regardless of the success of their strategy, believes, *ex ante*, that they are the most sophisticated player.

Level-K can be particularly useful when modeling behavior in games where "levels" are easily extracted from an optimal strategy. A popular game that makes use of Level-K is one in which a group of individuals place a bid to win an amount of money. The winning bid is the one closest to half of the average of the bids submitted. Given that bids may be placed between \$0 and \$100, the Nash prediction is that each subject will bid \$0. This contradicts laboratory experiment findings. Instead of expecting all opponents to calculate the Nash prediction and price accordingly, each player makes a judgement about the sophistication of their fellow players. For subjects who think that all other players will act randomly, the best response is to bid \$25, half of the expected value of randomization between \$0 and \$100. Those bidding \$25 would be labeled as Level 1 players, as they abstract one level below Level 0 players. A Level 2 player would anticipate Level 1 players optimal bid of \$25. Given that information, a Level 2 player would submit a bid half the average of a Level 1 player, or \$12.50. In this game, the intuitive nature of bid benchmarks to select the optimal bid based on a prediction of your opponent's sophistication makes Level-K theory an ideal model for explaining the deviations in laboratory data from the theoretical predictions of setting price at the Nash Equilibrium of \$0.

In our experiments, it is evident from the data that subjects update their beliefs between rounds, adding new information on the complexity of their opponent's strategy. In a game where individuals are matched with the same person every round, each round provides a new behavioral anchor price along with information on their opponent's level of sophistication. Both the new anchor, depending on whether a subject submitted the lowest price, and learning the sophistication of your opponents will impact the price chosen in subsequent rounds. When introducing random matching between rounds, individuals still use information provided in previous rounds to inform their pricing decisions. It is evident from the experimental results that even when matched with new individuals each round, prices begin in round 1 above the Nash Equilibrium and start to fall, converging close to the Nash prediction by round 10. This suggests that although matched with new subjects each round, participants use the information from previous rounds to determine the price set in subsequent rounds.

Theoretical Predictions

In this section, I define the Nash optimal pricing strategy in a simple Bertrand Price game as well as the derivation for price set by a Level 0 player in Bertrand Price Competition.

Bertrand Price Competition Nash Equilibrium

The simple Bertrand price competition game used in this paper's laboratory experiments consisted of sellers producing homogenous goods with identical marginal costs. These firms cover the entire market and compete on price in a one-shot game. The Nash Prediction, as described above, is where firms set their price equal to marginal cost. Suppose Firm A and Firm B are two

companies in a duopoly.² In order to sell their goods, each firm must set a price below their competitor. If Firm A were to set a price above their marginal cost, Firm B's best response would be to price just below firm A's price, capturing the entire market and giving Firm A no profit. Given that firm A knows that firm B will undercut them if they set any price above their marginal cost, Firm A will choose to set their price at marginal cost. Firm B, using the same reasoning, will choose to set their price equal to marginal cost. Both firms will lose money if they select a price below their marginal cost meaning that neither firm can achieve a higher profit by deviating from setting their price equal to marginal cost. This means that pricing equal to marginal cost is a Nash Equilibrium.

The Nash Predictions in the laboratory experiments differ from the predictions described in the paragraph above due to the discrete integer prices used in the experimental design. Using integer pricing options creates a new Nash Equilibrium to MC + 1 as well as a Nash Equilibrium at MC. While not relevant in the purely theoretical Nash Equilibrium prediction with continue pricing choices, the experiment resolves pricing ties by dividing the sales equally among firms. If firms set their price to MC + 1, each could split the small profit. Neither firm has an incentive to set their price equal to marginal cost, as that would result in no profits. Additionally, neither has an incentive to increase their price beyond MC + 1 as they make a profit of zero while their opponent services the entire market.

Prediction for a Level 0 Player

The construction of a "level" in a Level-K cognitive hierarchy model depends on level of best response from the expected decision of a player choosing randomly from the available options.

 $^{^{2}}$ This game theoretical approach will work for markets with N \geq 2 firms. I use a duopoly to explain the reasoning for simplicity.

To derive the expected price set by L0, all other players believe that L0 employs a random, mixed strategy that is uniformly distributed between $U[\overline{x}, \underline{x}]$ (Tong and Freeman, 2021). Appendix B. provides the full derivation of L0's strategy in which the expected price to set becomes $x_0 = \frac{\overline{x}-\underline{x}}{2}$.³ A Level 1 (L1) player would respond to L0's strategy by setting their price equal to one integer value below the L0 strategy leading L1 to set their price equal to $x_1 = \frac{\overline{x}-\underline{x}}{2} - 1$.⁴ A generalization to determine a level given as price is $x_L = \frac{\overline{x}-\underline{x}}{2} - L$ s.t. $x_L \ge \underline{x}$.

Experimental Design

To study the explanatory power of Level-K iterated reasoning models in Bertrand price competition games, experiments were performed to study subjects behavior in simple, multi-round Bertrand competition games. Subjects participated in one of five treatments, three consisting of 10 rounds and two consisting of one single-shot Bertrand game followed by 10 round Bertrand game. In all treatments, subjects were informed that they would be making decisions as a seller in a market. Participants were told they would be randomly matched with a new subject in each round. In order to sell their goods, each subject selected a price to sell their goods that is less than their opponent, without knowing the price set by their competitor in advance. After selecting prices, subjects were notified whether their selected price was lower than their opponent as well as their total earning from the round. Each treatment contained the setup as described above, adjusting parameters such as the marginal cost, imposing a Price Floor and restricting the upper bound of the price.

³ See Appendix B.

⁴ Since subjects are selecting integer prices, their best response must an integer value.

TABLE 1: Overview of Experimental Treatments

Treatments	Base	eline Treatm	nent	Price	Floor Treat	ment	High Mar	ginal Cost T	reatment	\$10 Reserv	ation Price	\$50 Reserv	ation Price
Session	Session 1	Session 2	Session 3	Session 1	Session 5	Session 6	Session 7	Session 8	Session 9	Session 1*	Session 2*	Session 1*	Session 2*
Rounds Played	10	10	10	10	10	10	10	10	10	1	10	1	10
Subjects	12	12	12	12	12	12	12	12	12	12	12	12	12
Firm Output Units	1	1	1	1	1	1	1	1	1	10	10	2	2
Marginal Cost	\$0	\$0	\$0	\$0	\$0	\$0	\$10	\$10	\$10	\$0	\$0	\$0	\$0
Maximum Price	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$100	\$10	\$10	\$50	\$50
Minimum Price	\$0	\$0	\$0	\$10	\$10	\$10	\$0	\$0	\$0	\$0	0	\$0	\$0

*In these sessions, the same subjects were used for Session 1 and Session 2.

The first or Baseline treatment, named Baseline, consisted of the simple Bertrand price competition game described above between two, randomly matched individuals, where subjects select prices at integer values between \$0 and \$100 with no marginal cost. In this treatment, the Nash Equilibrium prediction is that both sellers will set their price close to the marginal cost, in this case, at \$1. Given the pricing choices, \$0 is a Nash Equilibrium, but is dominated by the Nash Equilibrium at \$1 as no profits can be made at the \$0 Nash Equilibrium.

The second treatment, described as the Price Floor treatment, consisted of an identical set up to treatment one, except subjects were restricted to select a price between \$10 and \$100 instead of between \$0 and \$100, as if a Price Floor were imposed at \$10. Marginal Cost is held constant at \$0. In this treatment, the Nash Equilibrium predicts that subjects will set their price equal to \$10 in each round. In most conceptions of Bertrand Equilibrium, sellers make no economic profit for they are forced to price at marginal cost to make a sale. In the Price Floor treatment, the addition of a floor gives sellers the ability to make profit even by acting optimally and pricing at the Nash Equilibrium of \$10. The Price Floor treatment provides is the only treatment where a unique Nash Equilibrium exists. In other treatments, multiple Nash Equilibria exist. Although not behaviorally relevant, it is important to note. The third treatment, titled the High Marginal Cost treatment, takes the same design as the Baseline treatment, except imposing a marginal cost equal to \$10. The discrete choice set available to subjects creates a new Nash Equilibrium at \$11. Given the tie rule explained in the experimental design, when both firms set their price equal to \$11, they split the profit equally. If either deviated to marginal cost or higher than \$11, both scenarios would lead to no profit thus giving us a Nash Equilibrium of \$11.

The fourth and fifth treatments, titled \$10 Upper Bound and \$50 Upper Bound, respectively, take an identical setup to the Baseline treatment but restrict the sellers from setting a price above the session's upper bound. In the \$10 Upper Bound treatment, subjects continue to choose a prices equal to integer values but are restricted to only choosing prices between \$0 and \$10. The \$20 Upper Bound treatment employs the same set up, increasing the range to be between \$0 and \$20. To create identical subject payouts to the Baseline, Price Floor, and High Marginal Cost treatments, firms in the \$10 and \$50 Upper Limit treatments sell 10 and 2 units in each round, respectively. By increasing the firm output, the \$10 and \$50 Upper Bound treatments equalize the payout to treatments where subjects can set a price between \$0 and \$100. Since subjects in the \$10 Upper Bound treatment can only select integer prices between \$0 and \$10, by increasing their sales quantity in the round, setting a price of \$5 where they sell 10 units creates an equivalent payout to someone in one of the \$100 Upper Bound treatments selecting a price of \$50, selling one unit.

The Baseline, Price Floor, and High Marginal Cost treatments all allow subjects to select prices between \$0 and \$100. When comparing these treatments to the \$10 and \$50 Upper Bound treatments, the Baseline, Price Floor, and High Marginal Cost treatments will be referred to as the \$100 Upper Bound treatments.

Procedures

To study the applicability of Level-K to predict behavior in Bertrand Price Competition, a total of 11 laboratory experiments were conducted, nine in October of 2013 and two in April of 2024. Three sessions were conducted for the Baseline, Price Floor and High Marginal Cost treatments, while only one session was conducted for the \$10 Reservation Price and \$20 Reservation Price treatments. In total, 132 student subjects were recruited from the University of Virginia to participate in these experiments. Subjects were financially motivated and paid a portion of their lab earnings based on the their behavior.

Each of the 11 sessions began with a Holt-Laury Risk Aversion test to elicit risk preferences.⁵ Subjects then played a simple Bertrand Price Competition Game. In the Baseline, Price Floor, and High Marginal Cost treatments, subjects played 10 rounds in which they were randomly matched with a new competitor in each round. In treatments the \$10 and \$50 Upper Bound treatments, subjects played a single-round Bertrand Price Competition game before moving on to play 10 rounds where subjects were again randomly matched with a new competitor in each round. Subjects in all treatments were not informed in multi-round treatments when the final round would occur. At the conclusion of the experiment, subjects were paid earnings based on their decisions in the Holt-Laury Risk Aversion Test and in all rounds of the Bertrand Price Competition Game.

Across all sessions, subjects were asked to choose a price to sell a number of goods in a market. They were informed that they would be randomly matched with a new competitor in each round, where the individual setting the lowest price would clear the market. In the event of a tie in

⁵ See Appendix D.

which both subjects select the same sale price, both will sell half of their goods, splitting the profit equally. Across all treatments, subjects were paid 10% of the lab dollar earnings at the end of each session.

Subjects did not participate in a practice round before making their initial decisions in the Bertrand price competition game. Detailed instructions were provided explaining the structure of the game.⁶ Since the experiment aims to elicits pricing decisions without prior information on competitor behavior, providing a practice round may have influenced individuals' decisions in the first round. A drawback to not including a practice round is subjects who did fully understand the instructions may not choose the same price/strategy if they were fully informed on the experimental procedures. It is not possible to discern from the data which pricing decisions were due to a lack of understanding or intentional iterated reasoning, but is should be noted that some outlier behavior in the first round may be attributed to uninformed decisions.

Data

Experiments were run between October 17th, 2013, and November 1st, 2013, as well as between April 18th, 2024, and April 24th, 2024, in the VeconLab at the University of Virginia. Data was collected on each subject's chosen prices for all rounds as well as the random competitor in the round. In addition, cumulative earnings were calculated given each individual's price and whether they set the lowest price. Additionally, each subject was given a Holt-Laurey Risk Aversion Test to measure the risk aversion of each individual. On average, subjects selected 6 safe choices out of the 10 available options. This Holt-Laury Risk Aversion score suggests that the average individual was risk averse (Holt and Laury, 2002).

⁶ See Appendix C.

On average, subjects earned \$77.95 in lab dollars across all treatments, translating into actual earnings of \$7.95 from the experiment accompanied by a guaranteed \$10 for participating in the experiment. The minimum amount a subject earned across the experiments was \$14.50 in lab dollars while the maximum amount earned was \$187.00 lab dollars.

Results

On average, subjects in all treatments placed bids in the first round above the Nash prediction. Although players were randomly matched in each round, the average bid price begins to decline towards the Nash prediction in all treatments. In the Baseline, Price Floor, and High Marginal Cost treatments, the average price set in the first round was \$39.31 where, on average, subjects decrease their price by \$2.81 per round.

In the \$10 and \$50 Upper Bound treatment, results follow a similar trend to the Baseline, Price Floor, and High Marginal Cost treatments where prices begin above the Nash prediction and decline in each subsequent round. Given that the \$10 Upper Bound treatment only allows subjects to choose a price between \$0 and \$10, the comparison between the Baseline, Price Floor, High Marginal Cost treatments comes in Level-K predictions. The average level of subjects across the Baseline, Price Floor, and High Marginal Cost treatments, given subjects ability to select a price between \$0 and \$100, is Level 11⁷. By contrast, subjects in the \$10 Upper Bound treatment, on average, subjects were playing as Level 1.

Using a permutation test to measure the statistical difference between the Price Floor and High Marginal Cost treatments to the Baseline treatment, the null hypothesis can be rejected that pricing decisions in the first round were pulled from different distributions. Comparing the

⁷ Level 11 predicts that individuals reasoned 11 best responses away from a random (L0) individual.

Baseline and the Price Floor treatments, the average difference in price set was \$3.17 higher in the Price Floor treatment. Although the average price is higher in the Price Floor treatment, the two-tailed permutation test leads to a P-Value comparing the Baseline and the Price Floor treatments of 0.45 supporting no statistically significant difference between behavior in the two treatments. Comparing the Baseline to the High Marginal Cost treatments, subjects, on average, set a price \$0.56 lower in the High Marginal Cost treatment. Again, the High Marginal Cost treatment is not statistically significant different from the Baseline treatment, with a two-sample permutation test providing a two-tailed p-value of 0.73, rejecting the null hypothesis that samples were pulled from different distributions.

Baseline Treatment

The Baseline treatment yielded an average starting price of \$38.08 in round 1, declining by an average of \$2.96 per period. The decline in price was more pronounced in the first five rounds, falling by \$4.03 on average in rounds 1 through 5, while only averaging a drop of \$2.11 in the final five rounds. In the final round, all Baseline sessions ended nearly \$10 above the Nash prediction of \$1. Given the 101 possible pricing options, it is reasonable to suspect that it would take more than 10 rounds to arrive closer to the Nash prediction.



**N=36

Price Floor Treatment

The Price Floor Treatment exhibits similar behavior to the Baseline treatment where prices begin above the Nash prediction, declining sharply in the first five rounds and flattening out in the final periods. The average price set in the first round was \$44.14 with an average decline in subsequent rounds of \$3.34. Unlike the other four treatments, the Price Floor Treatment provides a unique Nash Equilibrium at \$10. By the tenth round, 39% of subjects across all sessions set a price equal to the Nash prediction. The increased number of subjects pricing at Nash Equilibrium is likely due to the sizeable profit available at Nash Equilibrium, which does not occur in the other

treatments. These results are similar to those found in the Baseline treatment, so it is difficult to say whether the increased number of subjects setting prices equal to the Nash prediction can be attributed to the higher Nash prediction or trends see in the Baseline treatment.



**N=36

High Marginal Cost Treatment

Of the three treatments allowing subjects to set prices between \$0 and \$100, the High Marginal Cost treatment produced the lowest round 1 price at \$35.72. On average, prices decreased by \$2.14 per round again, declining sharply in the beginning and diminishing in the final rounds. In session 3, a large spike in the average price occurs in rounds 8 and 9. This deviation occurs due

to two subjects setting their prices at \$100 and \$80 in round 8 and one subject setting their prices at \$77 in round 9. These large increases in price skew the averages in session 3 while most other subjects pricing similarly to subjects in sessions 1 and 2. Across all High Marginal Cost treatments, only one subject set their price equal the Nash Prediction of \$11, with none selecting a price equal to the marginal cost.



**N=36

\$10 Upper Limit Treatment

The \$10 Upper Limit treatment used an alternative design to the ones used in the Baseline, Price Floor, and High Marginal Cost treatments. Subjects participated in one-shot Bertrand Price Competition Game, selecting any integer price between \$0 and \$10. Subjects were then directed to a multi-round game with an identical setup as the one-shot game. The gray square in the table below represents the average price set in the one-shot game, while the orange triangles represent the average prices set by individuals in the multi-round game. Consistent with treatments allowing subjects to price between \$0 and \$100, subjects began setting prices above the Nash prediction averaging to \$4.50 in the one-shot game, dropping to \$3.25 in the first round of the multi-round game. This sizable decrease in average price provides evidence that suggests that subjects learn about the optimal pricing decisions between rounds, solidifying the use of no practice rounds in the experimental design. In subsequent rounds, the average decrease in price amounted to \$0.21 per round. In round 5, an increase in the selected price is due to a subject who sets their price equal to \$7 while the average price excluding the outliner in round 5 was equal to \$2.18. By the 10th round, 75% of subjects selected a price equal to the Nash prediction of \$1.



**N =12

\$50 Upper Limit Treatment

The \$50 Upper Limit treatment utilizes the same design as the \$10 Upper Limit treatment, just modifying the maximum price subjects can set from \$10 to \$50, allowing any integer prices between \$0 and \$50. Subjects participated in a single-round Bertrand Competition Game, followed by a 10 round Bertrand Competition game with random matching in each round. The average price set in the one-shot game is \$25.67, almost identical to the average expected price of a subject selecting a L0 price. Deviating from our other treatments, the first three rounds in the multi-round session yielded higher prices than the one-shot price. This deviation could be due in part to small sample size (N=12) with noisy responses as well as learning that occurs in the first round. In the

one-shot game, 25% of subjects selected a price greater than \$45 while another 25% chose a price less than or equal to \$10. The large variation in the prices set in the one-shot game were not seen in other treatments. Running more sessions of the \$50 Upper Limit Treatment would help determine whether large variations are consistent with the treatment or whether the session was an outlier.

Consistent with all treatments, prices decline in each subsequent round, falling by \$2.07 on average. In the first five rounds, the average decline in price was \$2.86 dropping to a \$1.44 decline per period in the final five rounds. In the final round, no subject selected a price equal to the Nash prediction. The lowest price set by subjects was \$5. This is a similar result to the Baseline Treatment where the 10th round prices ended higher than the Nash prediction with no subject setting their price equal to the Nash Equilibrium.



**N=12

Level-K Analysis Between Treatments

The lack of statistically significant distributional differences between the Baseline, Price Floor, and High Marginal Cost Treatments provides backing to assume that pricing decisions are not statistically different between the three treatments. Aggregating the data from the Baseline, Price Floor, and High Marginal Cost treatments into a \$100 Upper Limit treatment provides a comparison of levels in the \$10 and \$50 Upper Limit treatments. Rearranging the generalized Bertrand price prediction outlined in the theoretical given an individual's level, I use the pricing data to determine each individual's level in round 1 for each treatment.⁸

Camerer, in his seminal work on Level-K, cites that most subjects select strategies that are 1.5 levels away from a L0 (Camerer et al. 2004). In the \$100 Upper Limit treatments, a small minority of subjects selected a strategy indicative of a L1 or L2 player.⁹ The majority of players set their price at the L0 prediction of \$50 (9% of subjects). Interestingly, the second highest count of levels was at Level 10 followed by a large mass around Levels 31 and 32. While the vast majority of subjects set their price at or below \$50, 15% of individuals selected a price above \$50, with two individuals setting a price at \$100. Some of the pricing decisions that seem to deviate significantly from expectations may be attributed to a lack of understanding the instructions of the game, but some may occur from subjects acting randomly as described by a L0, setting their price randomly across the range of \$0 to \$100. These findings deviate from Camerer's suggestion that most individuals select a strategy 1.5 levels away from an L0, suggesting that Level K iterated reasoning model provides little predictive power for Bertrand Price Competition when subjects have a broad array of pricing options.¹⁰

⁸ Given $x_L = \frac{\overline{x} - \underline{x}}{2} - L$ s.t. $x_L \ge \underline{x}$, we rearrange such that given any x_L , $L = \frac{\overline{x} - \underline{x}}{2} - x_L$. ⁹ Nearly 5% of subjects selected a L1 price of \$49 with another 5% selecting an L2 price of \$48.

¹⁰ Camerer mentions that there are many games for which Level-K predictions are useful, but predictions fit "poorly in others" (Camerer et al. 2004).



**N=108

An important consideration before positing that Level-K provides little predictive power in Bertrand competition is the level of noise present when subjects are asked to select prices across 101 options. While Level-K theory may provide useful insights under narrow pricing options or obvious benchmark levels, other models of iterated reasoning may provide more insight into individual behavior. For this reason, it may be useful to use a model of noisy introspection as an alternative to Level-K theory (Goeree and Holt 2000). Using the \$10 Upper Limit Treatment as a comparison to the \$100 Upper Limit treatments helps support the proposition that noise may be the reason Level-K fails to provide useful predictions with numerous pricing options. In the \$10 Upper Limit Treatment, subjects were restricted to choosing integer prices between \$0 and \$10. An L0 player would, on average, be expected to set a price equal to \$5. Displayed in the graphic below, a majority of subjects set their price as if they were playing as an L0. This is consistent with findings in the \$100 Upper Limit Treatments. The second highest prices set are indicative of L1s playing the best response to an L0. It should also be noted that an equal number of subjects to L1s played what might be characterized as a L-1 strategy, pricing one level above an L0 prediction. It is hard to definitively say whether the price selection was due to failure to understand the game, random behavior, or a strategy based on iterated reasoning. Given more sessions of the \$10 Upper Limit Treatment, it may be easier to uncover whether L-1 strategy is due to noise or some level of iterated reasoning.



**N = 12

Given the narrow range of pricing options in the \$10 Upper Limit Treatment, results contain less noise than in the \$100 Upper Limit Treatments. In addition, the results from the \$10 Upper Limit Treatment align closer with Camerer's finding of most subjects selecting a strategy 1.5 levels away from L0, with an average level of 0.5 in the \$10 Upper Limit Treatment. With more observations, it is not unreasonable to suspect that the average could move closer to Camerer's prediction of 1.5 levels of iterated reasoning.

Level-K Predictions in Multi-Round Bertrand Price Competition Games

Level-K analysis is traditionally used to explain single-shot games as it relies on an individual positing the expected level of sophistication based on no outside information. Since learning is not a factor in the prediction of sophistication, it is questionable whether Level-K can be extrapolated to explain the behavior in multi-shot games. In her seminal work on iterated reasoning models, Rosemary Nagel observes that Level-K does not adequately explain behavior in periods beyond the first round. The laboratory results in this paper provide evidence that confirms Nagel's findings.

All treatments randomly matched individuals with new participants between rounds. Subjects were informed that they would be matched with a new subject in each round. Yet, across all treatments, subjects begin by selecting a price above Nash predictions and, on average, decreasing their price as the experiment progresses. These results, which remain consistent across sessions, suggest that individuals learn between rounds, using that information in their decisionmaking process in future rounds.

Level-K theory, as discussed in this paper, does not account for learning that occurs after the first round. Given that pricing decisions change substantially throughout the experiment, it is evident that Level-K, in its standard conception, does not explain subject behavior beyond the first round. Adjusting the Level-K model to account for learning between rounds may provide insight into how subjects change their pricing decisions given information on the sophistication of their players. Combining Level-K theory with a noisy introspection model may provide interesting results that have the power to explain multi-round Bertrand price competition results.

Comparing Results in the \$10 Upper Limit and the \$100 Upper Limit Treatments

In the \$100 Upper Limit treatments, a considerable mass of levels pop up around multiples of 10 (i.e. L0, L10, L30). One potential theory is that subjects select round numbers due in part to the round number bias (cite). A potentially compelling theory is that subjects, given 101 pricing choices, simplify their iterated reasoning by choosing prices in multiples of \$10 or \$5, anchoring off of a L0 prediction of \$50. Instead of playing a best response to an L0, which would induce a selection of \$49, subjects scale their options up selecting the next multiple of 10 below the maximum level of sophistication they expect from their opponent. In this way, they reduce the number of options by selecting multiples of 10. This brings the number of effective levels from L0 – L50 to L0 – L5. If true, subjects simplify the available options before iteratively reasoning their best response.

To test this hypothesis, we normalize the \$10 Upper Limit Treatment by scaling up each pricing choice by a multiple of 10.¹¹ Given the difference in sample size, the percentage of subjects in each level is used to make a fair comparison between the two treatments. In the table below,

¹¹ If a subject selected a price of \$5 in the \$10 Upper Limit Treatment, we multiply that number by 10, normalizing the price in the \$10 Upper Limit Treatment. This equates levels of iterated reasoning such that we can compare individuals across the \$10 Upper Limit and the \$100 Upper Limit Treatments.

similarities between the levels in the two treatments suggest that subjects in the \$100 Upper Limit Treatment simplify their rationality down from the 50 possible levels to 5 by choosing prices in multiples of 10. To further strengthen this result, a 2 tailed permutation test allows us to reject the null hypothesis that the \$10 Upper Limit and \$100 Upper Limit Treatments are pulled from different distributions.¹²



While these experiments are not definitive proof that subjects simplify their iterated reasoning to rounds numbers or benchmarks, it is certainly suggestive that subjects, given a broad range of pricing options, will simplify their levels of reasoning similarly to treatments with narrow options. Attempting to further understand this phenomenon, the \$50 Upper Bound Treatment was

¹² This test was performed after scaling up the \$10 Upper Limit prices.

conducted to see whether individuals behave similarly, where normalized levels would lead to similar masses at normalized L0, L10, etc. After running one session with 12 subjects, the noisy nature of the data prevented any useful analysis from the treatment. Given that subjects can select 51 different prices ranging from \$0 to \$50, more sessions need to be run to average out the noise present in one session.

In addition to running more \$50 Upper Limit Treatments, a sixth treatment with an upper limit of \$20 may help uncover the similarities and differences when altering the maximum settable price. Running this treatment may help explain when Level K can helpful given the numeracy of available pricing options. Additionally, as stated in an earlier section, the noisy nature of the data may be one of the leading reasons why Level-K fails to provide a useful prediction when there are a broad range of available prices. A model of noisy introspection may be better able to explain patterns observed in the dataset.

Conclusion

Level-K theory, under the right conditions, provides useful insights into pricing decisions made by subjects in one-shot Bertrand price competition games. After running five treatments varying the range of pricing choices, Level-K theory provides higher predictive power in Bertrand price competition games with narrow ranges of pricing options. As the range of pricing options increases, Level-K predictions are unable to explain laboratory results. While strict Level-K theory does not always provide adequate predictions, experimental evidence suggests that in Bertrand games with broad ranges of pricing choices, subjects simplify the range of options or "levels" to intervals before engaging in iterated reasoning producing results that mimic behavior observed in session with narrow ranges of pricing options. More experimental evidence may help uncover the range and scope for which subjects simplify their iterated reasoning given numerous pricing options.

This paper finds additional evidence that Level-K fails to provide adequate predictions in multi-round Bertrand price competition games. In the first round, subjects across all treatments set their initial price above Nash predictions and, on average, decrease their price in subsequent rounds which converge towards Nash predictions. Despite random matching between rounds, subjects seem to alter their pricing strategy throughout the course of the experiment suggesting that learning influences their price selections. Using the Level-K framework in conjunction with a learning model may provide useful insight into subject behavior beyond the initial round. Level-K is a useful framework to explain the iterated reasoning that occurs in simple Bertrand competition games and lays the groundwork to expand the model to handle complexities such as multiple rounds and numerous pricing options. Running additional treatments in future work will help adapt Level-K theory to model laboratory results.

References

- Bertrand, J. (1883) "Book review of theorie mathematique de la richesse sociale and of recherches sur les principles mathematiques de la theorie des richesses", Journal de Savants 67: 499–508.
- Camerer, C. F., T.-H. Ho, and J.-K. Chong. "A Cognitive Hierarchy Model of Games." *The Quarterly Journal of Economics* 119, no. 3 (August 1, 2004): 861–98. https://doi.org/10.1162/0033553041502225.
- Goeree, Jacob K., and Charles A. Holt. "A Model of Noisy Introspection." *Games and Economic Behavior* 46, no. 2 (February 1, 2004): 365–82. https://doi.org/10.1016/S0899-8256(03)00145-3.
- Holt, Charles A., and Susan K. Laury. "Risk Aversion and Incentive Effects." *The American Economic Review* 92, no. 5 (2002): 1644–55. http://www.jstor.org/stable/3083270.
- Nagel, Rosemarie. "Unraveling in Guessing Games: An Experimental Study." *The American Economic Review* 85, no. 5 (1995): 1313–26. http://www.jstor.org/stable/2950991.
- Stahl, D. O. (1993). Evolution of Smart Players. Games and Economic Behavior, 5(4), 604-617.
- Tong, Hanh T., and David J. Freeman. "Anchors of Strategic Reasoning in the Traveler's Dilemma." *Journal of Economic Behavior & Organization* 191 (November 2021): 28–38. https://doi.org/10.1016/j.jebo.2021.08.026.

Appendix

Appendix A. – Average Price per Round in Each Treatment by Session¹³

A. Average Price By Treatment in the Baseline, Price Floor, and High Marginal Cost

Treatments

Treatment	Treatment Averages				
Round:	Baseline	Price Floor	High Marginal Cost		
1	\$38.08	\$44.14	\$35.72		
2	\$33.03	\$38.53	\$31.81		
3	\$29.19	\$31.83	\$27.36		
4	\$24.19	\$28.08	\$24.75		
5	\$21.97	\$23.72	\$21.25		
6	\$19.92	\$20.56	\$19.67		
7	\$17.56	\$20.11	\$18.53		
8	\$15.14	\$18.75	\$21.33		
9	\$13.75	\$16.14	\$19.39		
10	\$11.44	\$14.11	\$16.44		
First Half					
Average	\$29.29	\$33.26	\$28.18		
Second Holf					
Average	\$15.56	\$17.93	\$19.07		
Total					
Average	\$22.43	\$25.60	\$23.63		

¹³ Internal Session Codes: Baseline Treatment – BR10, BR12, and BR14 | Price Floor Treatment – BR11, BR13, and BR15 | High Marginal Cost Treatment – BR16, BR17, and BR18 | \$10 Upper Limit Treatment – BR 19 and BR23 | \$50 Upper Limit Treatment – BR24 and BR25.

Treatment	Baseline			Price Floor			High Marginal Cost		
	Session	Session	Session	Session	Session	Session	Session	Session	Session
Round:	1	2	3	1	2	3	1	2	3
1	\$36.58	\$37.58	\$40.08	\$43.33	\$48.75	\$40.33	\$33.83	\$38.83	\$34.50
2	\$31.33	\$33.00	\$34.75	\$35.25	\$45.17	\$35.17	\$30.58	\$36.42	\$28.42
3	\$26.83	\$30.92	\$29.83	\$31.00	\$38.83	\$25.67	\$24.00	\$34.67	\$23.42
4	\$21.42	\$23.67	\$27.50	\$27.92	\$35.33	\$21.00	\$21.67	\$30.50	\$22.08
5	\$19.92	\$21.58	\$24.42	\$24.92	\$27.92	\$18.33	\$19.25	\$25.25	\$19.25
6	\$17.92	\$18.17	\$23.67	\$22.33	\$23.17	\$16.17	\$18.17	\$23.08	\$17.75
7	\$16.08	\$14.58	\$22.00	\$20.83	\$21.42	\$18.08	\$17.33	\$20.50	\$17.75
8	\$14.17	\$11.92	\$19.33	\$18.58	\$19.92	\$17.75	\$16.25	\$18.58	\$29.17
9	\$12.50	\$10.75	\$18.00	\$16.58	\$18.33	\$13.50	\$15.50	\$17.67	\$25.00
10	\$9.08	\$9.83	\$15.42	\$13.58	\$16.58	\$12.17	\$14.83	\$17.00	\$17.50
First Half Average	\$27.22	\$29.35	\$31.32	\$32.48	\$39.20	\$28.10	\$25.87	\$33.13	\$25.53
Second Half Average	\$13.95	\$13.05	\$19.68	\$18.38	\$19.88	\$15.53	\$16.42	\$19.37	\$21.43
Total Average	\$20.58	\$21.20	\$25.50	\$25.43	\$29.54	\$21.82	\$21.14	\$26.25	\$23.48

B. Average Price Set by Session for Baseline, Price Floor, and High Marginal Cost

Treatments

	\$10 Upper	\$50 Upper
Treatment	Limit	Limit
Round:	Session 1	Session 1
Single-Shot	\$4.50	\$25.67
1	\$3.25	\$31.38
2	\$3.08	\$30.17
3	\$2.75	\$25.77
4	\$2.33	\$22.48
5	\$2.58	\$19.93
6	\$2.00	\$18.00
7	\$2.00	\$16.88
8	\$1.58	\$16.74
9	\$1.42	\$14.76
10	\$1.33	\$12.71
First Half		
Average	\$2.80	\$25.95
Second Half		
Average	\$1.67	\$15.82
Total		
Average	\$2.23	\$20.88

C. Average Price in \$10 and \$50 Upper Limit Treatments

**Solid line between rounds 5 and 6 signifies the difference between the first half average and the second price average, presented at the bottom of the table.

Appendix B. – Level 0 Prediction Derivation for Bertrand Price Competition

$$E[\pi_i] = E[\pi_i | x_i < x_j] + E[\pi_i | x_i > x_j]$$
(5)

$$E[\pi_i] = \left(1 - \frac{x_i}{\overline{x} - \underline{x}}\right)x_i + \left(\frac{x_i}{\overline{x} - \underline{x}}\right)(0) \tag{6}$$

$$E[\pi_i] = x_i - \frac{x_i^2}{\overline{x} - \underline{x}} \tag{7}$$

$$0 = 1 - \frac{2x_i}{\overline{x} - \underline{x}} \tag{9}$$

(8)

$$x_i = \frac{\overline{x} - \underline{x}}{2} \tag{10}$$

Appendix C. - Bertrand Price Competition Lab Experiment Instructions

Instructions Example:

Instructions (ID =), Page 1 of 6

• Rounds and Matchings: The experiment sets up markets that are open for a number of rounds. Note: In each round, you will be matched with another person selected at random from the other participants. There will be a new random rematching each round.

• Interdependence: The decisions that you and the other person make will determine your earnings.

• Price Decisions: Both you and the other person are sellers in the same market, and you will begin by choosing a price. You cannot see the other's price while choosing yours, and vice versa.

• Sales Quantity: Consumers will purchase from the seller with the lowest price.

Continue with Instructions

• Production Capacity: Each seller in the market has the capacity to produce and sell up to 1 unit in each round.

• Price: Your actual sales quantity will depend on the prices chosen. Your price decision must be between (and including) \$0.00 and \$100.00.

• Demand: Buyers will demand a total of 1 unit at any price less than or equal to \$100.00. Purchases will be made from the low-price seller. In the event of a price tie, each seller will have a sales quantity of 0.5 units.

• Production Cost: Your cost is \$0.00 for operating your capacity unit. So there is no cost for producing a unit.

• Sales Revenue: If you sell a unit, your sales revenue is your price. Since your sales are affected by the other's price, you will not know your sales revenue until market results are available at the end of the period.

Continue

- Earnings: Your profit or earnings for a round is the price for a unit that you sell (there is no cost).
- Cumulative Earnings: The program will keep track of your total (cumulative) earnings. Positive earnings in a round will be added, and negative earnings will be subtracted.
- Matchings: At the beginning of each round, there will be a new random pairing of all participants, so the person who you are matched with in one round may not be the same person you will be matched with in the subsequent round. Matchings are random, and you are no more likely to be matched with one person than with another.

• Price Choice: All people will begin a round by choosing a number or "price" between and including \$0.00 and \$100.00.

- Capacity: Each seller in the market has the capacity to produce and sell up to 1 unit in each round.
- Demand: Consumers are willing to purchase up to 1 unit at any price <= 100. The seller with the lowest price obtains the sale (with equal division in the event of a tie).
- Cost: Your total cost increases by \$0.00 for each unit you sell.
- Earnings: Your earnings are your total revenue (price times sales quantity) minus your total cost. Positive earnings are added to your cumulative earnings, and losses are subtracted.

• Special Earnings Announcement: Your cash earnings will be 10% of your total earnings at the end of the experiment.

Finished with Instructions

Subject Decision Screen Example:

Submit Decision for Round 1, ID: 1

Choose a price between \$0.00 and \$100.00.

The demand quanity is 1 unit for P <= \$100.00, with the purchase to be made from the seller with the lowest price (the sales unit is split equally in the event of a tie).

Your cost will be **\$0.00** for each unit that you sell.

Round	Your	Other's	Your	Total	Total	Your	Cumulative
	Price	Price	Quantity	Cost	Revenue	Earnings	Earnings
1	please choose \checkmark	Submit	*	*	*	*	\$0.00

Subject Round Result Page Example:

Results for Round 1, ID: 1

You chose a price of **\$5.00**, and the other seller chose a price of **\$9.00**.

Seller ID	Ranked Price	Sales Quantity
1 (you)	\$5.00	1 unit
2	\$9.00	0 units

So your total sales revenue is **\$5.00**. Your cost is **\$0.00** for each unit, so your total cost is **\$0.00**. Your earnings for this round are: **\$5.00 - \$0.00 = \$5.00**.

To summarize, your earnings = **\$5.00** for round 1. Your cumulative earnings = **\$5.00**

Please remember that your cash earnings will be calculated as 10% of your cumulative earnings in the final round.

Begin Round 2

Appendix D. – Holt-Laury Risk Aversion Test Instructions Instructions:

Instructions (ID =), Page 1 of 6

You will be making choices between two lotteries, such as those represented as "Option A" and "Option B" below. The money prizes are determined by the computer equivalent of throwing a ten-sided die. Each outcome, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, is equally likely. If you choose Option A in the row shown below, you will have a 1 in 10 chance of earning \$4.00 and a 9 in 10 chance of earning \$3.20. Similarly, Option B offers a 1 in 10 chance of earning \$7.70 and a 9 in 10 chance of earning \$0.20.

Decision	Option A	Option B	Your Choice
1st	\$4.00 if the die is 1 \$3.20 if the die is 2 - 10	\$7.70 if the die is 1 \$0.20 if the die is 2 - 10	A : ○ or B : ○
1 A 1			

Continue with Instructions

Instructions (ID =), Page 2 of 6

Each row of the decision table contains a pair of choices between Option A and Option B.
You make your choice by clicking on the "A" or "B" buttons on the right. Only one option in each row can be selected, and you may change your decision as you wish.
Note: Note, try clicking on one of the radio buttons, then change by clicking on the other one.

Decision	Option A	Option B	Your Choice

1st \$4.00 if the die is 1 \$7.70 if the die is 1 \$3.20 if the die is 2 - 10 $(3 \times 3^{-1})^{-1}$ \$0.20 if the die is 2 - 10 $(3 \times 3^{-1})^{-1}$ **A**: \bigcirc or **B**: \bigcirc Continue

Instructions (ID =), Page 3 of 6

Even though you will make ten decisions, **only one** of these will end up being used. The selection of the one to be used depends on the "throw of the die" that is the determined by the computer's random number generator. No decision is any more likely to be used than any other, and you will not know in advance which one will be selected, so please think about each one carefully. This random selection of a decision fixes the row (i.e. the Decision) that will be used.

For example, suppose that you make all ten decisions and the throw of the die is 9, then your choice. A or B, for decision 9 below would be used and the other decisions would not be used

Decision	Option A	Option B	Your Choice
9th	\$4.00 if the die is 1 - 9 \$3.20 if the die is 10	\$7.70 if the die is 1 - 9 \$0.20 if the die is 10	A : ○ or B : ○

Continue

Instructions (ID =), Page 4 of 6

After the random die throw fixes the Decision row that will be used, we need to obtain a second random number that determines the earnings for the Option you chose for that row. In Decision 9 below, for example, a throw of 1, 2, 3, 4, 5, 6, 7, 8, or 9 will result in the higher payoff for the option you chose, and a throw of 10 will result in the lower payoff.

Decision	Option A	Option B	Your Choice
9th	\$4.00 if the die is 1 - 9 \$3.20 if the die is 10	\$7.70 if the die is 1 - 9 \$0.20 if the die is 10	A : ○ or B : ○

10th	\$4.00 if the die is 1 -10	\$7.70 if the die is 1 - 10	A : ○ or B : ○

For decision 10, the random die throw will not be needed, since the choice is between amounts of money that are fixed: \$4.00 for Option A and \$7.70 for Option B.

Continue

Instructions (ID =), Page 5 of 6

• Making Ten Decisions: After you finish these instructions, you will see a table with 10 decisions in 10 separate rows, and you choose by clicking on the buttons on the right, option A or option B, for each of the 10 rows. You may make these choices in any order and change them as much as you wish until you press the Submit button at the bottom.

• The Relevant Decision: One of the rows is then selected at random, and the Option (A or B) that you chose in that row will be used to determine your earnings. Note: Please think about each decision carefully, since each row is equally likely to end up being the one that is used to determine payoffs.

• Determining the Payoff: After one of the decisions has been randomly selected, the computer will generate another random number that corresponds to the throw of a ten sided die. The number is equally likely to be 1, 2, 3, ... 10. This random number determines your earnings for the Option (A or B) that you previously selected for the decision being used.

Instructions Summary (ID =)

• To summarize, you will indicate an option, A or B, for each of the rows by clicking on the "radio buttons" on the right side of the table.

• Then a random number fixes which row of the table (i.e. which decision) is relevant for your earnings.

• In that row, your decision fixed the choice for that row, Option A or Option B, and a final random number will determine the money payoff for the decision you made.

• Payoffs will be made in cash.

Finished with Instructions

Subjects Decision View:

Submit Decision for Round 1, ID: 1

Please select either A or B for each of the ten Decisions below. Remember: Each decision has an equal chance of being used to determine your earnings.

Real Money Payoffs: The choices that you make on this page will be used to determine your earnings, these are real money payoffs that will be paid to you in cash.

Decision	Option A	Option B	Your Choice
1st	\$4.00 if the die is 1 \$3.20 if the die is 2 - 10	\$7.70 if the die is 1 \$0.20 if the die is 2 - 10	A : ○ or B : ○
2nd	\$4.00 if the die is 1 - 2 \$3.20 if the die is 3 - 10	\$7.70 if the die is 1 - 2 \$0.20 if the die is 3 - 10	A : ○ or B : ○
3rd	\$4.00 if the die is 1 - 3 \$3.20 if the die is 4 - 10	\$7.70 if the die is 1 - 3 \$0.20 if the die is 4 - 10	A : ○ or B : ○
4th	\$4.00 if the die is 1 - 4 \$3.20 if the die is 5 - 10	\$7.70 if the die is 1 - 4 \$0.20 if the die is 5 - 10	A : ○ or B : ○
5th	\$4.00 if the die is 1 - 5 \$3.20 if the die is 6 - 10	\$7.70 if the die is 1 - 5 \$0.20 if the die is 6 - 10	A : ○ or B : ○
6th	\$4.00 if the die is 1 - 6 \$3.20 if the die is 7 - 10	\$7.70 if the die is 1 - 6 \$0.20 if the die is 7 - 10	A : ○ or B : ○
7th	\$4.00 if the die is 1 - 7 \$3.20 if the die is 8 - 10	\$7.70 if the die is 1 - 7 \$0.20 if the die is 8 - 10	A : ○ or B : ○
8th	\$4.00 if the die is 1 - 8 \$3.20 if the die is 9 - 10	\$7.70 if the die is 1 - 8 \$0.20 if the die is 9 - 10	A : ○ or B : ○
9th	\$4.00 if the die is 1 - 9 \$3.20 if the die is 10	\$7.70 if the die is 1 - 9 \$0.20 if the die is 10	A : ○ or B : ○
10th	\$4.00 if the die is 1 - 10	\$7.70 if the die is 1 - 10	A : ○ or B : ○
Press here after you have made ALL 10 decisions.			

Subject Results Page:

Final Results for for the Earnings-Relevant Round 1, ID: 1

The random number determined by the first die throw was: **4** so we use **Decision 4** . You chose **Option A** for this decision.

 Decision
 Option A
 Option B
 Your Choice

 4
 \$4.00 if the die is 1 - 4 \$3.20 if the die is 5 - 10
 \$7.70 if the die is 1 - 4 \$0.20 if the die is 5 - 10
 Option A

The second die throw is: 10

Therefore, your earnings for round 1 will be \$3.20

You have just finished the final round, Please note your cumulative earnings. Thanks for participating.

You may now close the browser window, unless you are instructed to return to the Login Menu for another experiment.

Return to Login Menu