

Applied Visual Analytics for Economic Decision-Making

Anya Savikhin*

Ross Maciejewski†

David S. Ebert‡

*Purdue University Department of Economics

†‡Purdue University Regional Visualization and Analytics Center (PURVAC)

ABSTRACT

This paper introduces the application of visual analytics techniques as a novel approach for improving economic decision making. Particularly, we focus on two known problems where subjects' behavior consistently deviates from the optimal, the *Winner's* and *Loser's Curse*. According to economists, subjects fail to recognize the profit-maximizing decision strategy in both the Winner's and Loser's curse because they are unable to properly consider all the available information. As such, we have created a visual analytics tool to aid subjects in decision making under the *Acquiring a Company* framework common in many economic experiments. We demonstrate the added value of visual analytics in the decision making process through a series of user studies comparing standard visualization methods with interactive visual analytics techniques. Our work presents not only a basis for development and evaluation of economic visual analytic research, but also empirical evidence demonstrating the added value of applying visual analytics to general decision making tasks.

1 THE WINNER'S AND LOSER'S CURSE

In order to best demonstrate the applicability of visual analytics to decision making problems, we have chosen to analyze a familiar situation, bidding at an auction. When a person bids at an auction, economists assume that the item under bid is of equal value to all participants. This type of auction is known as a *Common Value Auction*. Here, each participant has some estimate of the item's worth with a degree of uncertainty, and each places a bid accordingly. The winner of the auction is the individual who bids the highest. Unfortunately, if we assume that the average bid of all the participants represents the actual worth of the item, then the winner of the auction has now clearly overpaid. This phenomenon is known as the *Winner's Curse* [2, 10, 14]. As the number of bidders increases, so does the severity of the amount overpaid and the more likely it is that some of them have overestimated the auctioned item's value. In technical terms, the winner's expected estimate is the value of the first order statistic, which increases as the number of bidders increases. Since most auctions involve a degree of common value and uncertainty, this deviation from the optimal bid is an important phenomenon to study. In this case, companies must assess the value of the market in that area and make a bid accordingly.

The sister problem to this is known as the *Loser's Curse* [7], which appears when the parameters of the experiment are changed in such a way that the naive bid is now below the profit-maximizing bid, and the profit-maximizing bid results in winning the company. This is because the lower bound for the range of possible company values is increased, reducing the bid range in which high negative profits would result (as in the Winner's Curse setup). Subjects

who do not win the company have lower profits than those who bid higher and win the company, but subjects consistently underbid in this setup and do not find the profit-maximizing solution.

In order to better understand subject motivations, the Winner's and Loser's Curse have also been studied in a simplified framework which is referred to as the Acquiring a Company problem, first formulated by Samuelson and Bazerman [14]. The Acquiring a Company problem is simplified because subjects no longer interact with or bid against one another; instead, the subject bids on a company, the value of which is randomly determined by the computer but unknown to the subject. The subject is also told that the company is worth more to him than it is to the seller (in this case, the computer is the seller). Even though subjects no longer compete against one another to win the company, both the Winner's and Loser's Curse are still present.

In order to compare our results with previous work, we have chosen to employ the Acquiring a Company framework in our visual analytics application and evaluation work. The motivation for this work is to help subjects overcome this failure by helping them better consider the relevant information, thereby improving their decision making abilities. We present an extensible visual analytics framework for economic decision making, and demonstrate the added value of such techniques through the results of a user study. Our work shows the benefit of visual analytics in economic decision making, improving subject performance in the Acquiring a Company framework, and aiding in overcoming the Winner's and Loser's Curse phenomena.

2 APPLYING VISUAL ANALYTICS TO ECONOMIC DECISION-MAKING

We present this study at a crucial time when many researchers are calling for new techniques and systems to help non-expert users of varying abilities in complex decision-making and analysis tasks [9]. According to the NIH/NSF Visualization Research Challenges report, visualization is essential to the solution of complex problems in every field and systems must allow ordinary people to experiment with "what if" situations [9]. Our research addresses these points through the application of a visual analytics tool that can aid in complex decision making tasks in real world applications with uncertainty.

2.1 Related Literature in Economics

Many studies have attempted to address the Winner's and Loser's Curse problems through the Acquiring a Company framework over the past 20 years. The goal of these studies was to find a way for subjects to avoid the Winner's or Loser's Curse by improving their decision-making abilities. Some of these studies allowed for role reversal or modified the game, while others attempted to make the game "easier" for subjects, or to prepare subjects better. In one study, a training package was designed to help subjects, with training on various conditional probability problems, but this was only marginally effective in improving learning [8]. In another study, training was employed where subjects were given a lesson in probability and then performed four similar tasks where the Acquiring

*e-mail: anya@purdue.edu

†e-mail:rmacieje@purdue.edu

‡e-mail:ebertd@ecn.purdue.edu

Table 1: Previous Results

	Baseline Learning Rate	Type of Enhancement	Subsequent Learning Rate
Ball et al, 1991	7%	Role Reversal	6%-9%
Tor and Bazerman, 2003	N/A	Training	8.5%
Bereby-Meyer et al, 2002	14%*	Extra Feedback	26%-50%*
Charness and Levin, 2006	<10%*	Reduced Complexity	30%-60%*
			*total percentage of optimal bids (not learning rate)

a Company task was done third in every order [17]. In the manipulations discussed above, it was found that the Winner’s Curse persisted for most subjects. It was found that the Loser’s Curse was equally persistent [7].

In most of these studies, the guideline for “learning” was defined as bidding close to the optimal bid at some point and continuing this strategy for the duration of the study. The general result of most studies for the baseline case is a learning rate of around 7%. Table 1 summarizes the results of selected previous studies in order to give the reader an understanding of typical findings in this area of research.

2.2 Related Work in Visual Analytics

Visual analytics is the science of analytical reasoning assisted by interactive visual interfaces, which has already been applied and found to be effective in social sciences such as management, finance, marketing and organizational behavior to aid in decision making [1, 5, 16]. Previous work has employed visuals and found that subjects given a problem statement in the form of a visual performed better than subjects given the same information in the form of a written statement [5]. Also, subjects given a visual interface outperformed those using a text-based interface in decision making in low and high complexity tasks [16]. Research in the field of energy and management has found that a system that employs visualization is an effective tool in planning and decision-making for bio-energy production [1]. As such, we feel that applying visual analytics to economic decision making is a natural extension.

Since it has been shown that visualization can extend a person’s working memory [3], it is natural to assume that the application of visual analytics will improve subjects’ economic decision making skills. Furthermore, economists believe that the reason subjects succumb to the Winner’s and Loser’s Curse is that they have a limited focus of attention and therefore fail to consider all of the information needed to solve the problem correctly. The decision making failure can be explained by the theory of *bounded awareness*, which suggests that people are rational but sometimes fail to make decisions rationally because the limits of cognitive abilities do not allow them to consider all the information available [4]. A generalized version of this theory is the more well-known concept of *bounded rationality*, which is a school of thought in economics that explains the limits of rational behavior and has been researched by behavioral theorists and experimental economists [6, 13, 15]. We believe that the application of visual analytics can help subjects overcome bounded rationality because of improvement in the productivity of their cognitive effort. Because cognitive effort is costly, the effort expended by individuals may not be enough to solve the problem and discover the optimal solution. We propose that subjects use an information production function, where effort level (chosen by the subject) and the type of visual aid (chosen externally) are inputs. Subjects maximize their earnings in the Acquiring a Company game by choosing optimal levels of effort and their bid. An economic model is under development that will help explain this phenomenon theoretically.

2.3 Contribution

We have evaluated the effects of various applications through the implementation of an interactive visual analytics program and a

simple visual representation, as well as a baseline tabular representation of the data, under the hypothesis that such representations will enhance subjects’ abilities to overcome the bounded awareness issues. In particular, the interactive visual analytics program should be most effective because subjects use this to explore the link between decisions and outcomes *before*, rather than *after* they make their decisions as compared to the simple visual or tabular representations, in which they receive the information as feedback after the decision has been made. The interactive visual representation is the only way for subjects to receive information before they make their decision, because non-interactive representations (the simple visual and the table) do not have this capability. To test our hypothesis, we have performed a user study in which one group of subjects was given the data in tabular form, one group was provided a simple visual representation of the data, and another group was given the use of an interactive visual analytics program. Consistent with our hypothesis, we found that the group given the visual representation outperformed the group given the data in tabular form. Also, the group given the interactive visual analytics program outperformed both the group given the visual representation and the group given the data in tabular form, where the latter was statistically significant. This result is very promising for future work in the application of visual analytics to various economic decision making problems.

3 EXPERIMENTAL DESIGN AND SETUP

We considered two problems: a Winner’s Curse treatment and a Loser’s Curse treatment. Six treatments were run in total (Winner’s Curse-Interactive Visual Analytics (WC-IV), Winner’s Curse-Simple Visual (WC-V), Winner’s Curse Table (WC-T), Loser’s Curse-Interactive Visual Analytics (LC-IV), Loser’s Curse-Simple Visual (LC-V) and Loser’s Curse Table (LC-T)). The treatments and the experiment parameters are summarized in Table 2. The baseline user study provided subjects with feedback information on outcomes in tabular form. The second user study provided subjects with feedback information in a simple visual representation. Finally, the third user study allowed subjects to make use of an interactive visual analytics program at all times during their decision-making process. Our prediction was that subjects who participated in the interactive visual analytics study would bid closer to the profit maximizing decision than those participating in the simple visual or tabular studies. Furthermore, we also predicted that even those participating in the simple visual study would outperform those participating in the tabular study.

3.1 Experimental Design

Our design utilized parameters similar to previous research. Each subject, independent of any other subject, was given the role of the buyer. As a buyer, the subject had to make the decision of what to bid for a company. The company could take values between X and U , with all values in that range equally likely. In accordance with Samuelson and Bazerman [14], the company was worth 50% more to the buyer than to the seller for all treatments. During the experiment, the value of the company was determined randomly by the computer. If the bid was higher than or equal to the value of the company, the computer always sold the company to the buyer. If the bid was lower than the value of the company, the computer never sold the company to the buyer. The buyers were aware that the seller was informed about the value of the company before the

Table 2: Winner's/Loser's Curse treatments and parameters.

Winner's Curse	Loser's Curse
$X = 4$	$X = 50$
$U = 100$	$U = 100$
$M = 1.5$	$M = 1.5$
Optimal Bid = 8	Optimal Bid = 100
Naive Bid = 52	Naive Bid = 75

transaction took place. The buyers were also aware that they would not know the value of the company before they placed their bids. Each buyer was given 350 tokens at the beginning of the experiment with which to bid. Each buyer was given 30 periods in which to bid on the company, and informed that the company takes a new value in each period, independent from the previous period or from any other period. Tokens were converted to dollars and subjects were paid in cash based on how well they performed at the end of the experiment. After each period, subjects were shown the information screen, which was either the visual or a simple table.

We also introduced using the *Surprise-Quiz* methodology as described by Merlo and Schotter to quantify learning in the Acquiring a Company game [12]. According to Merlo and Schotter, if the experiment requires subjects to learn the properties of a static equilibrium, the surprise quiz methodology can be used. In the Surprise-Quiz, subjects were given the opportunity to use the visual (or the table information screen, in the other treatment) to find the "optimal bid". Subjects were told that an optimal bid exists, that they had one chance to submit a bid they thought was close to the optimum, and that they would be paid based on how close their bid was to the optimal bid. It was explained to subjects that the optimal bid is one that would maximize earnings over time, regardless of the company value. The subjects were given as much time as they required (up to 1 hour, although no subject spent this much time) to come up with their decisions. The bid submitted in the bonus round is the measure of learning in the experiment and we used this bonus round bid to compare decision making across treatments.

3.2 Theoretical Predictions

The assumption in economic theory is that a risk-neutral individual will maximize his or her expected net earnings over time. However, most subjects fail to take into account the fact that the value of the company is known to the seller, and bid the expected value of the company instead, which results in the Winner's or Loser's Curse. Profit-maximizing subjects will take conditional expectations into account and solve the maximization problem:

$$\max_B \left(\frac{B-X}{100-X} \right) \left(1.5 \left(X + \frac{B-X}{2} \right) - B \right) \quad (1)$$

where B is the bid. Differentiating this equation with respect to B yields:

$$\frac{1}{100-X} \left(1.5 \left(X + \frac{B-X}{2} \right) - B - .25(B-X) \right) \quad (2)$$

By setting Equation 2 equal to zero and solving for B, we find that $B = 2X$.

The optimal bid in the Winner's Curse treatment would be 8, while the optimal bid in the Loser's Curse treatment would be 100. The naive bid of 52 for the Winner's Curse treatment would result in an expected profit loss, while the naive bid of 75 for the Loser's Curse treatment would result in lower expected profit than any higher bid. This is because raising the lower bound for the range of company values reduces the decision space where outcomes could be negative, which is very large in the Winner's Curse treatment.

To understand this result intuitively, suppose that you are a buyer bidding on a company in the Winner's Curse treatment. The naive

approach would be to bid the expected value of the company (52), under the assumption that your bid will be accepted by the seller and since the company is worth 50% more to the buyer, the buyer will benefit with higher than expected earnings. However, you must take into account the fact that the seller knows the value of the company before accepting the offer. Therefore, at a bid of 52, the bid will only be accepted if the company is worth less than 52. In this case, if the company value is between 4 and 33, there will be possible negative profit for the buyer, and if the company value is between 34 and 50, there will be possible positive profit. This results in an expected profit loss at a bid of 52; therefore, this bid is not optimal. In fact, any bid higher than 8 will result in expected profit loss. Bidding 8, on the other hand, will result in acceptance only if the value of the company is between 4 and 8. However, the result in this case is that of expected profit gain.

Suppose now that you are a buyer bidding on a company in the Loser's Curse treatment. The naive approach would be to bid the expected value of the company (75), without accepting that asymmetric information plays a role here. This time, increasing the bid at the margin would result in a larger gain than the bid of 75. For example, we know that at the bid of 75, expected profit is 9.75, where the buyer just breaks even (gets 0) if the company value is 50, gains positive profit if the company value is between 50 and 75, and does not acquire the company if the company value is greater than 75. Increasing the bid to 76 results in an expected profit of 9.99, with small negative profit of -1 if the company value happens to be 50, positive profit between and including company values 51 and 76, and no profit for company values above 76. Continuing to increase the bid, finally, at the optimal bid of 100, one finds that expected profit is 12.75, with some small negative profit at company values below 67, and large positive profit at company values above and including 67. Since company values are uniformly distributed and the bidder bids for 30 periods, it is rational to bid the value which will result in highest expected profit - in this case, this is 100.

3.3 Visual Analytics Application

Our contribution is helping subjects improve decision making through the use of an interactive visual analytics program. A simple visual representation, which basically increases the quality of information feedback, is also used for comparison. Both the interactive visual analytics program and the simple visual representation use 2D graphics to show the profit the subject would have received at every possible company value at his/her bid amount, after the outcome of the period is determined. According to Cleveland's graphics elements hierarchy, position along a common scale is one of the best ways to display this type of data [18]. The interactive visual analytics tool also allows the subject to explore the dataset on his/her own time before he/she makes a decision. First, the subject is able to move the proposed bid up and down to see how the possible profit changes. Second, the subject is able to 'sample' numerous possible company values with the click of a button, and see what he/she would have earned at the proposed bid. The simple visual representation and the baseline tabular representation programs track the amount of time spent on each period, as well as the subjects' bid, company value, and profit. The interactive visual analytics tool tracks the above but also goes further and tracks the subjects' changes in the proposed bid (up or down) and clicks of the 'sample company value' button.

Figures 1 and 2 show screenshots of the simple visual analytics representation at naive and optimal bids of the WC-V and LC-V treatments. Figure 3 shows the interactive visual analytics program for the Winner's Curse treatment (WC-IV) during the time the subject is deliberating a decision, and after the subject has placed a bid and the outcome has been revealed. The Loser's Curse interactive visual analytics treatment (LC-IV) is similar but has a different possible earnings distribution which is identical to LC-V. For com-

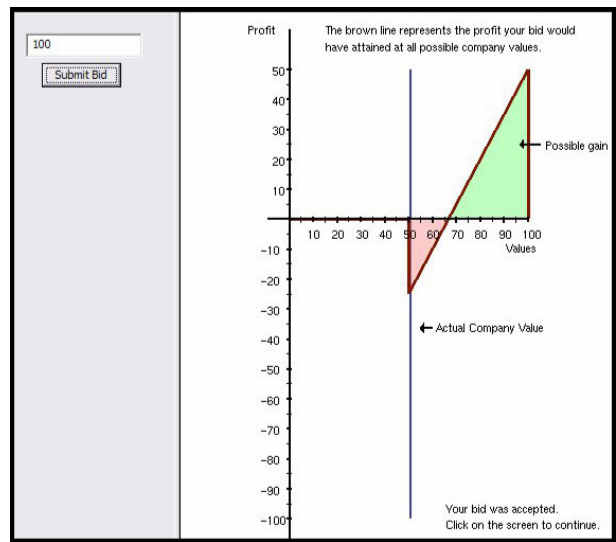
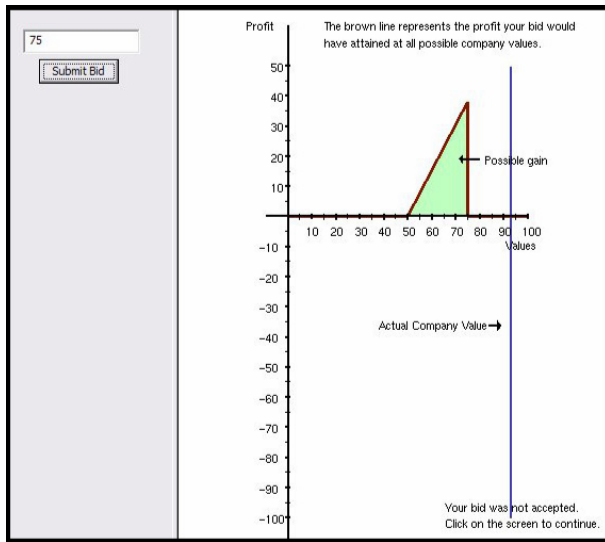


Figure 1: Visual analytics: Loser's Curse treatment





Figure 2: Visual analytics: Winner's Curse treatment

parison, Figure 4 is a screenshot of the information screen shown to subjects who were not given the visual display and were instead given the information in table format as the baseline case (WC-T and LC-T).

Both the simple visual display and the tabular display allow subjects to input a bid and look at all possible outcomes of that bid. The largest difference between the display shown in Figures 1 and 2 and the display shown in Figure 4 is that the display in Figures 1 and 2 employ visualization to disseminate the information to subjects. Both information screens show the relationship between profit and company value. The simple visual display is thought to be more effective because a visual can allow the user to see the relationship more easily and quickly. It can also save time, an important economic resource. One might believe that the information screen which is more effective is the one with the most information. It is important to be aware that the table includes the same amount, if not more, information than the visual display. The table depiction includes columns for the bid entered, the value of the company,

and the possible profit in each case, for all values of the company. The simple visual display is slightly more complex, but does not contain any extra information, except that color is used to draw attention. Green areas on the graph represent areas of possible positive profit, while red areas on the graph represent areas of possible negative profit. The actual company value also appears on the graph as a vertical blue line. In comparison, the interactive visual analytics tool is more complex than the two non-interactive treatments described above. The interactive visual analytics tool has all the capabilities of the simple visual representation, and also allows subjects to discover possible outcomes *before* making their decision by clicking "sample a company value" or moving their proposed bid up or down.

One other difference between the interactive visual analytics tool and the simple visual display is that the former shows information to the user *before* a decision is made, while the latter shows information to the user *after the decision has been made*. However, since the user makes decisions 30 periods in a row, it is likely that

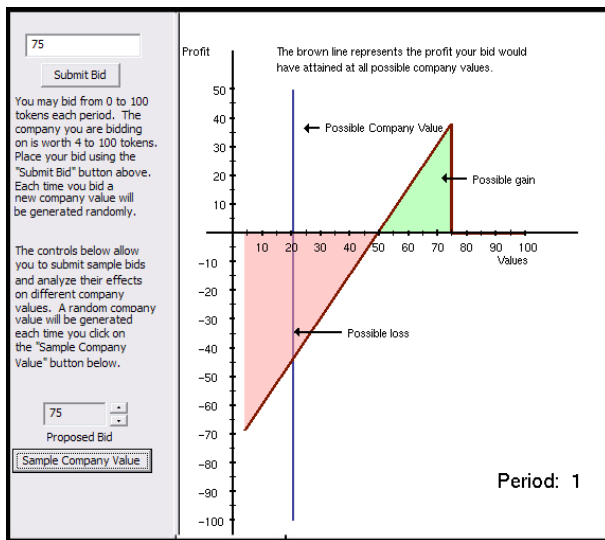


Figure 3: Interactive Visual Analytics - During Decision (left) and After (right) (WC)

the user remembers the previous period's outcome screen while he or she is making the next period's decision - if this is the case, then in all periods but the first period, we can suggest that the user has the simple visual display's information in mind *before the subjects makes the current period's decision*. In future work, it would be possible to test if this is the case by running a treatment where the user is shown the outcome from the previous period at the beginning of the current period, and instructed to think about his current decision at that time. The timing of showing this screen would be the same, with the only difference being the additional instructions. Because this change is so minor, we feel that the results would not be affected.

between the company value and the profit. Most importantly, the green positive profit and red negative profit areas were displayed in order to effectively convince the user of the importance of the size of these areas in making his/her decision. The relative size of the areas represents the probability of positive versus negative profit. Because the user chooses the bid value, and the visual looks different for different bid values, the user can change his/her bid values to observe the differences in positive and negative profit possibilities. The interactivity of the more advanced display is even more effective, because it allows users to view possible outcomes during the decision making process.

3.4 User Studies

Subjects in the WC-V and LC-V treatments were shown the simple visual display after each period, while subjects in the WC-IV and LC-IV treatments were also shown the interactive display during each period while they were making their decisions. Subjects in the WC-T and LC-T treatments were shown the tabular display after each period. At the end of the experiment, all subjects were also given the choice to use the visual analytics display again to make their decisions in the bonus round (subjects in the WC-T and LC-T treatments were given the choice of using the tabular tool as well.) The choices of subjects during the first 30 periods were recorded, and choices subjects made while using the display for the bonus round were also recorded. In the WC-IV and LC-IV treatments, all choices made during the deliberation process were recorded as well.

We recorded the amount of time spent on each decision as the time in seconds from the start of the period through the end. For the simple table and simple visual display this included the time spent deliberating (no visual or table during this time), plus the time spent observing the output screen for the future period and recording the result. For the interactive visual analytics tool, this included time spent deliberating, which in this case was supplemented with the interactive visual analytics tool, and time spent observing and recording the result (the result screen was the same as the most recent previous screen - the only difference being the addition of the actual company value line.) The mean amount of time spent on each period, with all periods and subjects was between approximately 8.5 seconds and 30 seconds. Within the Winner's Curse sessions, the most time was spent on the WC-T and WC-V treatments (21 seconds and 30 seconds, respectively) and the least time was spent

Your Bid	Firm Value	Possible Profit
60	50	15
60	55	23
60	57	26
60	58	27
60	59	29
60	60	30
60	61	0
60	62	0
60	63	0
60	65	0
60	70	0
60	75	0
60	80	0
60	85	0
60	90	0
60	95	0

Company's value was 57

Your bid was accepted.
Click on the screen to continue.

Figure 4: Tabular treatment for Winner's/Loser's Curse

The visual analytics display was chosen because it was believed that displaying the entirety of the information on the screen in a visual form would allow users to overcome bounded awareness issues and be more likely to focus on all the relevant information. The particular graph was chosen to clearly show the relationship

Table 3: Summary of subjects.

	Tabular Treatment	Simple Visual Treatment	Interactive Visual Treatment	Total
WC	11 subjects 100% passed training	7 subjects 90% passed training	13 subjects 100% passed training	31
LC	8 subjects 100% passed training	7 subjects 80% passed training	10 subjects 90% passed training	25
Total	19	14	23	56

on the WC-IV treatment (18.7 seconds.) Within the Loser’s Curse treatments, the most time was spent on the LC-T and LC-V treatments (19 seconds and 15 seconds, respectively), and the least time was spent on the LC-IV treatment (8.6 seconds.) We are in the process of investigating further the reason for these differences in time spent. In every treatment, there is a negative correlation between period number and time spent; that is, users spent less time deliberating each decision and outcome as periods progressed. This correlation is between -0.16 and -0.33, with no systematic difference across treatments.

From the post-experiment comments, which were submitted anonymously, it is clear that most subjects understood how to use both the simple visual display and the interactive visual analytics tool. Discussion of strategies included “find a value at which loss region is smaller so that it is more likely to gain profit”, “stay mostly around 30s because it seemed the loss and gain areas were close to equal” and “try to make the green [possible positive profit] area as large as possible”. Subjects using a table instead of the visual analytics display made comments such as “guess randomly” or “try to find the pattern of the company value.”

3.5 Experimental Setup

Inexperienced undergraduate students from Purdue University were recruited to participate in the study - 56 subjects in total participated. Subjects earned between \$5 and \$25 in the experiment, where their earnings depended on their success at learning the optimal solution to the problem. 23 subjects participated in the experiment using the interactive visual analytics display, 14 subjects participated in the experiment using the simple visual display, and 19 subjects participated in the experiment using the table display. The number of subjects participating in the experiment is summarized in Table 3.

3.6 Pre-Experiment Quiz & Training

Before the beginning of the experiment, subjects also received training on reading the information screen. For all treatments, this included a “Lemonade Stand” task which was not related to the experiment, where subjects were given an information screen which showed possible profit for all number of customers. An example of how to calculate possible profit was given, and then subjects were asked to answer 5 simple questions about the information screen to gauge how well subjects understood how to read the screen. The “Lemonade Stand” task was not a difficult or complex problem like the Acquiring a Company problem. Subjects who did not answer all 5 questions correctly were asked to leave and did not participate in the experiment. Subjects who participated in the WC-V, LC-V, WC-IV and LC-IV treatments were shown the graphical version of the possible profit, while subjects who participated in the WC-T and LC-T treatments were shown the table version of the possible profit. 100% of the subjects were able to answer the questions when the information screen was a table, while only 90% of the subjects were able to answer the questions when the information screen was a graph. This assured that only subjects who understood what the information screen was telling them were included in the study. After passing the training, subjects were explained the rules of the Acquiring a Company game and shown one period of the bidding process. The information screen was then explained to subjects on a display at the front of the room.

4 RESULTS

Subjects given the interactive visual analytics treatment learned the optimal solution more often than subjects who were only given the simple visual representation or the table information screen. Subjects given the interactive visual analytics treatment outperformed subjects given the table information screen for both the Winner’s and Loser’s Curse treatments, and this was statistically significant. Subjects participating in the Loser’s Curse simple visual representation treatment outperformed subjects in the Loser’s Curse tabular treatment, and this was statistically significant. The rest of the results also show the predicted effect, but these are not statistically significant. Also, when asked to make the optimal decision in the bonus round, subjects in the interactive visual analytics tool were more likely to utilize the information screen again than subjects in the simple visual treatment, and subjects in the simple visual treatment were more likely to utilize the information screen again than subjects in the table treatment (100% (23/23) versus 86% (9/14) versus only 25% (5/20)).

4.1 Comparisons of Learning

Learning in the game is measured based on how close the bonus round bids (or the surprise quiz) are to the optimal. These can be compared across treatments to understand the effectiveness of visual analytics at improving decision making. The bonus round results should be accurate because the bonus round allowed subjects to pause and think more carefully about their decisions. The bonus round also gave subjects the opportunity to explore the range of the visual display at their own pace and really “test out” different bids without any penalty. The bonus round result was included in the payment structure of each subject and subjects who bid closest to optimal in the bonus round were given a large monetary bonus as compared to the money earned in the first 30 rounds of the game.

With learning defined as bidding equal to or very close to the optimal value in the bonus round, learning rates can be reported. Here, one could define “close to the optimal” as within 10% of the total bid area, where the “bid area” is between 4 and 100 for WC treatments and 50 and 100 for LC treatments (since the company value can never fall below 4 or 50, respectively, it may not be reasonable to include those values in the total bid area - bidding these would be dominated by bidding the lowest possible company value, as this would increase possible earnings with zero risk - subjects seemed to understand this as no subject bid below these values in either the 30 period game or the bonus round.) It may have been difficult for subjects in the LC treatments to bid within just 5 values of the optimal bid (within 10% is within about 5 integers for this treatment, and subjects were only allowed to bid integer values), therefore expanding the ‘close to the optimal’ definition to ‘just under within 10 integers’ may be reasonable, which is approximately within 10% for the WC treatments. Learning was higher for subjects in the visual treatments versus subjects in the tabular treatment, and this is summarized in Table 4.

4.2 Statistical Analysis

We ran a two-sample Wilcoxon (Mann-Whitney) rank-sum test to check for differences between all possible pairings of groups within the WC treatment and all possible pairings of groups within the LC treatment bids in the bonus round, and these results are summarized

Table 4: Learning.

	Tabular Treatment	Simple Visual Treatment	Interactive Visual Treatment
WC - defined as within 10% of area	0% 0/9 subjects	14.3% 1/7 subjects	54% 7/13 subjects
LC - defined as within 10% of area	12.5% 1/8 subjects	28.5% 2/7 subjects	10% 1/10 subjects
LC - defined as within 10% of WC area	12.5% 1/8 subjects	28.5% 2/7 subjects	30% 3/10 subjects

Table 5: Test results summary.

	Test	Prediction	direction	Confirms With Hypothesis	p-value
Winner's Curse	Wilcoxon	WC-V \neq WC-T	WC-V < WC-T	Yes	0.928
		WC-V \neq WC-IV	WC-V > WC-IV	Yes	0.103
		WC-IV \neq WC-T	WC-IV < WC-T	Yes*	0.012*
Loser's Curse	Wilcoxon	LC-V \neq LC-T	LC-V > LC-T	Yes*	0.055*
		LC-V \neq LC-IV	LC-V > LC-IV	No	0.845
		LC-IV \neq LC-T	LC-IV > LC-T	Yes*	0.040*

in Table 5. The hypothesis would suggest that decision making for the groups receiving the interactive visual analytics display would be improved as compared to the groups who were shown the simple visual representation or the table, while groups receiving the visual representation would still outperform those receiving the table. Subjects using the interactive visual analytics tool or the simple visual display should choose to bid values closer to the optimal bid - lower values for the Winner's Curse treatment, and higher values for the Loser's Curse treatment, as compared to the group receiving the data in table format. The Wilcoxon test gives a measure and direction of the difference in medians between the two groups. The WC-IV group chose bids which were lower than the WC-T group as predicted by the hypothesis, and these results were statistically significant (p-value 0.012). The LC-IV and LC-V groups chose bids which were higher than the LC-T group as predicted by the hypothesis, and these results were statistically significant as well (p-values 0.040 and 0.055, respectively). However, no significant results were found when comparing the other pairs of groups. It is clear that the interactive visual analytics display was useful for both the Winner's and the Loser's Curse. The simple visual display was also useful for the Loser's Curse.

4.3 Reasoning

While it seems that in complicated situations an interactive visual analytics program can be very useful, we believe we may be able to explain why the simple visual representation was useful for the Loser's Curse treatment but not as useful for the Winner's Curse treatment. One possible reason is that the relative size of the gain region as compared to the loss region was much more pronounced in the Loser's Curse than in the Winner's Curse. Thus, the Loser's Curse had gain and loss regions which were easier for subjects to compare than those in the Winner's Curse graph. That is, it was clearer to users in the Loser's Curse treatment that the naive bid had a smaller relative possible positive profit region than the optimal bid than it was to users in the Winner's Curse treatment. This proposition is supported by previous research which suggests that graphic information is given more weight by subjects when visual representations show sharp contrast or greater variation in size [11]. We also strongly believe that increasing the number of subjects would increase the significance of the results, as the level of cognitive ability is heterogeneous across subjects.

4.4 Distances to the Optimal Bid

Figure 5 shows the distances of bids from the optimal in all treatments, with WC-IV and LC-IV groups in front of WC-V and LC-V groups, and WC-V and LC-V groups in front of WC-T and LC-T groups. For both the Winner's and Loser's Curse experiments,

our hypothesis was that subjects provided with the interactive visual program would bid closer to the optimal than subjects provided with a tabular representation or the simple visual representation. Also, subjects provided with the interactive visual analytics program would outperform subjects with the simple visual representation. Clearly we can see that subjects in the visual treatments bid closer to the optimal than subjects provided with the tabular representation, and this was statistically significant for all interactive treatments (WC-IV and LC-IV) as well as for the simple visual treatment in the Loser's Curse (LC-V). Each bar in Figure 5 represents a subject's distance to the optimal bid and we find that in general, WC-IV, LC-IV, WC-V and LC-V groups outperformed WC-T and LC-T groups in finding the optimal bid. The only contradiction was Subject 8 in LC-T, who found the optimal bid. In this subject's comments, he explained that he came to his answer by solving a mathematical formula maximizing his profit. He also stated that he did not use the information screen for his calculation. This is not the typical way that subjects look at this issue, as most subjects' comments included statements relating the information screen to their choices. Therefore, it is probable that Subject 8 in LC-T is an outlier and that, in a larger set of observations, he would have less weight.

It could be argued that the reason subjects using the interactive visual analytics tool and the simple visual display outperformed subjects using the simple table is because graphs are easier to understand and read than tables. However, results from the "Lemonade Stand" task would suggest otherwise. First, while every user was able to complete the task correctly in the case of the table version, only 90% of subjects were able to complete the task correctly in the graph version. Furthermore, while those subjects who did not correctly complete the "Lemonade Stand" task were asked to leave, even those who stayed had more trouble in the visual analytics treatments versus the simple table treatments. During the first few rounds of the experiment, no subjects had questions on interpreting the table, while some subjects did ask questions about interpreting the graph. The sessions with the table treatment took far less time in total than the sessions with the visual treatments for this reason. Even though it was more difficult to understand how to read the interactive visual or simple visual analytics information screens, users still outperformed those who had less trouble reading the table information screen.

The improvement in cognition may be due to two factors: informational advantage and computational advantage. The information presented in the simple visual display and the simple table was the same; thus, all improvement here is due to computational advantage - it was easier to interpret the information in visual form. However, the user can obtain more information in the interactive visual ana-




Figure 5: Distance of bids from the optimal value (Left) Loser's Curse (Right) Winner's Curse. Here each bar represents the distance a single subject's response to the surprise quiz was from the optimal bid. The shorter the bar, the closer the bid was to optimal.

lytics tool treatment than in the other treatments by moving the proposed bid to see various outcomes. Thus, in the interactive visual analytics treatment, increased information may also play a role. We do not distinguish between these two factors in this paper; however, this could be explored in future work.

The results suggest that visual analytics did help users to overcome some of the bounded rationality issues. In economic decision making problems, considering all the information at hand is extremely important. Because of cognitive limitations, users are not always able to consider all of the information in order to make rational decisions. Using visual analytics can improve cognitive abilities and thus overcome some of these problems.

5 CONCLUSIONS AND FUTURE WORK

To the best of our knowledge, this is the first study that has addressed how visual analytics can improve economic decision making. It has been widely reported that competitive decision making failures are common in Winner's and Loser's Curse problems in economics. We used visual analytics as a tool to help subjects consider all the information of the problem and make a decision that is closer to the optimal in the Acquiring a Company framework. The interactive visual analytics display was extremely successful: subjects using the visual analytics display made improved decisions as compared to subjects who had the same information in table form. The simple visual display was also successful. This suggests that showing all the information via a visual can be useful for overcoming bounded rationality issues that arise from the cognitive limitation of considering all the information at hand. This also suggests that the addition of an interactivity element is very important to the usefulness of visuals.

This research is the gateway for using visual analytics as a tool in other economic decision making problems. Visual analytics would be especially helpful in a situation where information feedback can aid subjects, as our results are convincing evidence that visual analytics can increase the quality of the feedback. The interactivity of the visual analytics program also increases participant involvement and understanding, which is necessary for successful decision making. A broader situation where the Winner's Curse is an issue is in common value auctions. One direction for future research is to create a visual analytics display that helps subjects make profit maximizing decisions in the common value auction framework. Common value auctions are often used in the fields of real estate, commercial construction, and book publication rights. Visual analytics, already proven to help subjects make decisions in the Acquiring a Company framework, would be extremely valuable in the common value auction framework as well. Visual analytics could then be ap-

plied by professionals in the workplace and would improve decision making and, therefore, economic efficiency.

REFERENCES

- [1] N. Ayoub, R. Martins, K. Wang, H. Seki, and Y. Naka. Two levels decision system for efficient planning and implementation of bioenergy production. *Energy Conversion and Management*, 48:709 – 723, March 2007.
- [2] E. Capen, R. Clapp, and W. Campbell. Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 23:641–653, 1971.
- [3] S. Card, J. Mackinlay, and B. Shneiderman. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers, San Francisco, 1999.
- [4] D. Chugh and M. Bazerman. Bounded awareness: What you fail to see can hurt you. 2004.
- [5] G. W. Dickson, G. DeSanctis, and D. McBride. Understanding the effectiveness of computer graphics for decision support: A cumulative experimental approach. *Communications of the ACM*, 29(1):40–47, January 1986.
- [6] G. Gigerenzer and R. Selten, editors. *Bounded Rationality: The Adaptive Toolbox*. MIT Press, 2001.
- [7] C. A. Holt and R. Sherman. The loser's curse. *The American Economic Review*, 84(3):642–652, June 1994.
- [8] L. C. Idson, D. Chugh, Y. Bereby-Meyer, S. Moran, B. Grosskopf, and M. Bazerman. Overcoming focusing failures in competitive environments. *Journal of Behavioral Decision Making*, 17:159–172, June 2004.
- [9] C. Johnson, R. Moorhead, T. Munzner, H. Pfister, P. Rheingans, and T. S. Yoo. NIH/NSF visualization research challenges report, 2005.
- [10] J. H. Kagel and D. Levin. *Common Value Auctions and the Winner's Curse*. Princeton University Press, Princeton, New Jersey, 2002.
- [11] N. H. Lurie and C. H. Mason. Visual representation: Implications for decision making. *Journal of Marketing*, 71(1):160 – 177, Jan 2007.
- [12] A. Merlo and A. Schotter. A surprise-quiz view of learning in economics experiments. *Games and Economic Behavior*, 28(1):25–54, July 1999.
- [13] A. Rubinstein. *Modeling Bounded Rationality*. MIT Press, 1998.
- [14] W. F. Samuelson and M. H. Bazerman. The winner's curse in bilateral negotiations. In V. L. Smith, editor, *Research in Experimental Economics*, volume 3, pages 105–137. JAI Press, 1985.
- [15] H. A. Simon. *Models of Bounded Rationality*. MIT Press, 1982.
- [16] C. Speier and M. G. Morris. The influence of query interface design on decision making performance. *MIS Quarterly*, 27(3):397–423, 2003.
- [17] A. Tor and M. H. Bazerman. Focusing failures in competitive environments: Explaining decision errors in the monty hall game, the acquiring a company game, and multiparty ultimatums. *Journal of Behavioral Decision Making*, 16(5):353–374, 2003.
- [18] L. Wilkinson. *The Grammar of Graphics*. Springer, New York, 1991.