

Information Sharing Using an Electronic Marketplace

David Montana, Alice Leung and Marshall Brinn
BBN Technologies
10 Moulton Street, Cambridge, MA 02138
{dmontana,aleung,mbrinn}@bbn.com

Abstract— The intelligence failures preceding the September 11 terrorist attacks highlight critical shortcomings with the sharing of information between government organizations. Fixing these problems involves overcoming technological and cultural barriers, fostering openness while not endangering privacy and security. We describe an approach to information sharing using an electronic marketplace, similar in some respects to eBay, that facilitates and incentivizes cooperation while minimizing exposure to abuses of privacy and security. To investigate the properties of this approach, we have devised a game, Find the Ring, that centers on an electronic marketplace for information sharing among the players. We analyze this game theoretically and describe the results of experiments run using software agents as players.

1. INTRODUCTION

In the aftermath of the September 11 attacks, it became clear that there existed enough information in the intelligence community to have anticipated and prevented the attacks. The problem was that the information was distributed among individuals in different organizations, and the pieces of information could not be integrated into a single coherent picture. This has prompted calls for greater sharing of information between organizations, including the passage of the Homeland Security Information Sharing Act. (A summary of the act is available at [7].)

There are two main obstacles to improving the sharing of information. The first is the inertia of separation, both cultural and technological. Traditionally, government agencies have resisted cooperating and have avoided sharing information with each other, although this attitude is rapidly changing. More difficult to overcome is the incompatibility of the different organizations' computer systems, complicating the possibility of sharing information by integration of these systems.

The second obstacle is the danger to security and privacy engendered by unchecked information sharing. While the most commonly voiced concern is the compromise of information sources, there are also other types of abuses that can occur when too much information is made too easily available. For example, consider the case of the spy Robert Hanssen, who used the database intended for sharing of information among FBI agents to easily and anonymously gather much of the FBI's most sensitive information for sale [8]. If he also had access to the CIA's and NSA's data, he could have done even worse damage. An important doctrine of U.S. national secu-

rity (that applies equally well to protection of civil liberties) is restricting access to sensitive information to those who have a "need to know", hence minimizing losses when a leak does occur. Too much sharing of information can undermine this.

We propose an approach to information sharing, described in greater detail in Section 3, that circumvents these obstacles. An electronic double-auction marketplace allows individuals to buy and sell information (and informational services) as needed. A broker matches buyers and sellers without revealing the information being sold until the sale is complete. (Keeping the offers secret is necessary not just to enhance the security and privacy of the information but also because information, unlike a physical object, loses its value once its content is revealed.) We drive the marketplace with a large monetary reward for the individuals that assemble all the pieces of information about a security threat to form a coherent picture.

This approach to information sharing provides some important advantages over the standard concept of open databases. First, it facilitates and incentivizes cooperation. There is no need to worry about database formats and incompatibilities, since the marketplace is inherently about humans, and not machines, interacting. Furthermore, because all parties are rewarded for working together, it is likely that people will overcome whatever institutional boundaries might separate them. A second big advantage of our approach is that it minimizes exposure to abuses of security and privacy. Forcing people to pay for whatever information they receive should restrict the flow of information to what is truly needed and encourage auditing of transactions.

An actual implementation of such an electronic market would be a major undertaking, which requires investigation of a variety of issues beforehand. Included among these issues are computer security, matching of buyers and sellers, and creation of audit trails. In this paper, we focus on economic and game-theoretic issues centering on how to ensure that the electronic market efficiently generates information sharing. To this end, we introduce the game "Find the Ring", described in Section 4, which allows us to investigate an information sharing market without other complicating factors. This permits us to present a preliminary analysis, both theoretical (Section 5) and experimental (Section 6), of an information sharing market.

2. BACKGROUND

Information markets are artificial economic markets whose primary purpose is to gather information. The idea for such markets started with the realization that markets are good predictors of future events. For example, the orange juice futures market is a good predictor of future weather in Florida [16]. The next step was not just gathering information as a side effect but instead establishing markets specifically for this purpose [5], [14], providing an alternative to other information aggregation methods like polling or the delphi method [1].

Until now, such information markets have been analogous to futures markets, with investors betting on possible outcomes by buying “shares” in the outcomes. One of the earliest and longest-running of such information markets is the Iowa Electronic Markets, whose results have provided accurate predictions of elections for over a dozen years [3]. Other examples of such markets are the Hollywood Stock Exchange for predicting Academy Award winners and the Foresight Exchange for general predictions about the future [13]. Hewlett-Packard has even introduced such a market for predicting future sales [4].

There have been a variety of investigations of issues in such information “futures” markets and how to handle them. For example, some research has addressed problems such as information traps [12] and lack of equilibrium [2]. One major issue is how to handle the case of a “thin” market (i.e., when there are only a small number of investors in the market for a particular “security”) to allow this technique to be applied when there are only a small number of people with knowledge of the subject area. One approach, described in [4], is to use a mathematical analysis beyond just the price of the security. A second major issue is the limitation of the futures market approach to predicting answers to questions with only a small number of possible answers rather than generating open-ended responses. For example, an information futures market can answer the question “will terrorists attack this year?” with a probabilistic “yes” or “no”, but not with a set of names of terrorists or their plans to fly planes into buildings. One approach to allowing information futures markets to handle a greater range of possible outcomes is combinatorial information markets, where investors can bet on not just possible outcomes but also conditional probabilities of different classes of outcomes [6]. However, these are still limited in what conclusions they can generate.

Despite these efforts, the information futures market approach is inherently not suited to true information sharing. Sharing information only via effects on a futures market does not generally impart enough information to the person who needs it and provides too much information to those who do not. Instead, information sharing requires a means to bring together people with complementary information and have them deal directly with each other. We therefore turn to a better model than futures markets for this purpose. Open electronic marketplaces, such as those run by eBay (var-

ious items), monster.com (employment opportunities) and match.com (romantic attachments), provide an efficient way to bring together people with complementary needs. These online versions of traditional goods and services markets use electronic search to perform fast and efficient matching of potential buyers and sellers. We apply the same basic approach, albeit with some important domain-specific differences, to the distribution of information rather than jobs or romantic partners.

Another thread of work that we draw upon is the use of artificial agents to investigate the behavior of economic markets. While most experimental work on artificial or simulated markets uses human test subjects, some work has successfully used software-based agents as an alternative. Groundbreaking work in this area was that on the Santa Fe artificial stock market, which investigated the dynamics of markets using artificial traders [10]. More recent research using agent-based economic modeling is [11]. A comprehensive set of guidelines for constructing agent-based markets is provided in [9]. The use of software agents rather than, or in addition to, human traders provides some advantages. It allows investigation of a wider range of conditions and strategies. Furthermore, it permits more extensive experiments to investigate how the system scales as the number of participants grows.

3. INFORMATION SHARING MARKET OVERVIEW

We now describe our vision of how an electronic market for information sharing should work. This description is geared towards the sharing of intelligence data, because this is the driving motivation, but it is applicable to other types of information as well. We start with an overview of the general framework and then discuss some of the possible variations in the details of the implementation.

In an electronic information sharing market, the marketplace is actually a computer system, which provides participants each with their own highly restricted view of the market. The participants in the market are buyers and sellers of information and information services. (An example of an information service is checking a database to determine whether someone has entered the country.) Sellers submit their information, along with potentially a description of this information, to the marketplace. However, this submitted information is not revealed to any other participants yet. Buyers submit offers or requests for particular types of information. A broker, which is a software agent, accepts such queries from buyers and finds sellers whose information descriptions appear to match. (Depending on the type of information, this matching can be a difficult problem.) If the buyer and seller settle on a price (and there are various approaches to determining the price), the buyer receives a copy of the information from the seller. A large reward for piecing together snippets of information into a full picture is the ultimate incentive that drives the market. (Whoever is running the market must see value in information properly pieced together, and hence should be

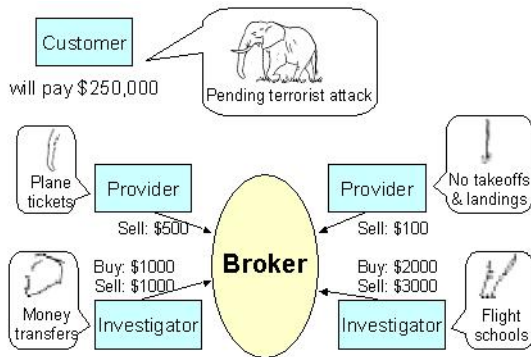


Figure 1. Like the blind men in the parable, the participants do not individually have sufficient information to see the full picture. By buying and selling information, they work to fit the pieces together.

willing to pay someone for doing this successfully.) The process is pictured in Figure 1.

An important part of an information sharing market is having the proper mechanisms to prevent participants from purchasing information to which they are not entitled. Restriction of certain types of information to certain participants can be automatically enforced, with access permission based partly on the clearance level of the participants and the classification level of the information. For example, a local police officer might be able to buy from another police department but not from the FBI; an ordinary citizen might not have permission to buy anything, but just to sell. Audit trails are another mechanism to ensure the legitimacy of transactions with respect to security and privacy. Need-to-know would be chronicled in this audit trail of transactions, and individuals could be held accountable for the information that they purchase.

There are two main reasons why sharing of intelligence data is best accomplished with an electronic market rather than an electronic bulletin board, which is a traditional way to allow sharing of information between people electronically. First, an electronic market incentivizes sharing rather than hoarding of information. It encourages people to sell information (including taking the time to properly advertise the information that they are selling) as well as to “create” information (by which we mean going out and finding information rather than making it up, the latter of which we need mechanisms to discourage). This incentive to share transcends organizational boundaries. The second reason is that it discourages wanton gathering of information, forcing those retrieving information to focus on what they truly need to pursue their lines of investigation. An audit trail of sales transactions makes people responsible to their organizations for how they spend their money and hence for any unjustified acquisition of information.

This discussion of the general market framework has left open a variety of questions regarding specific implementations, whose answers depend on the situation. These ques-

tions include:

1. What is the currency? In particular, are participants trading straight cash or shares of the reward?
2. How are potential buyers and sellers matched without revealing the seller’s information to the buyer prior to the sale?
3. How is the price of a transaction determined?
4. What rights to the information are transferred when information is sold and what rights are retained by the buyer? For example, can someone who bought a piece of information subsequently sell that information?
5. What data about general market trends and/or specific market transactions should be revealed to the participants?
6. How does the market prevent cheating, particularly with respect to the sale of bad information? Does the market require a reputation system [15]?

In our description in the next section of a simple information sharing market that we have actually implemented, we answer these questions among others for an example market.

4. A GAME: FIND THE RING

To investigate the properties of electronic information sharing markets, we needed an example of such a market. Therefore, we invented a game, Find the Ring, that has an information sharing market as a central component. In the game, players buy and sell clues (information) about suspects, attempting to identify a subset of the suspects called the *ring*. The objective for the players is to maximize their money at the end of the game.

The initial setup of the game is as follows:

- There are S suspects, R of whom are in the ring. The names of all the suspects are known, but the identities of the ring members are not.
- There are P players, each starting with $\$M$ as their initial money.
- There are three types of clues about whether or not the suspects are in the ring. *Equivalence* clues tell that two suspects are either both in the ring or both out of the ring. *Anti-equivalence* clues tell that, for a pair of suspects, one but not both are in the ring. *Negative* clues tell that a single suspect is not in the ring. Each player receives E equivalence clues, A anti-equivalence clues, and N negative clues to start the game. These initial clues should have the property that no individual player has enough clues to fully determine the members of the ring but all the players’ clues in aggregate are sufficient.

Play proceeds as follows:

- Play is divided into rounds. Each round has two phases, one for the players to make their offers followed by one for the broker to resolve all the offers.
- Players can make three different types of offers. A *sell offer* indicates that a player will sell a specified clue in his possession for a minimum of a specified price. A *buy offer* indicates that a player will buy any of a specified set of clues for at most a specified price. The set of clues of in-

terest are specified by selecting a subset of the clue types, a subset of the suspects acceptable as one of those referred to by the clues, and a subset of the suspects acceptable as the second suspect referenced. (This second subset of suspects is ignored for negative clues, where there is only one referenced suspect.) An *identification offer* is an attempt by a player to identify all the members of the ring.

- A buy offer and sell offer match if (a) the clue type and suspects of the clue designated by the sell offer match the specifications for clues given by the buy offer and (b) the price of the buy offer is at least as great as the price of the sell offer. The best buy offer for a given sell offer is the one with the highest price. Conversely, the best sell offer for a given buy offer is the one with the lowest price.
- A transaction can occur when there is a set of matching buy and sell offers. The actual price is set halfway between the two offer prices. Upon execution of a transaction, an amount of money equal to the price is transferred from the buyer to the seller, and a copy of the clue is transferred from the seller to the buyer. The seller retains a copy of the clue, and in future rounds both players can sell the clue.
- The first action of the broker each round is to resolve all identification offers. Each identification offer that is incorrect receives a penalty of $\$P$ for the player making the offer. A single reward of $\$R$ is dispensed for a correct identification. If there are multiple correct identifications, then only one is accepted with the winner chosen randomly. If there is a correct identification, then the game is over.
- If the ring has not been identified, the broker next selects a random ordering of the players and sequentially selects one buy offer from each player. For each buy offer, the broker finds the best matching sell offer (resolving ties randomly) and, if there is a matching offer, performs the resulting transaction. The broker then repeats this same procedure, except this time selecting sell offers and finding the best matching buy offers.
- If Q consecutive rounds pass without a transaction being completed, then the game is declared a *stalemate* and ended with no reward given.

The following information is made available to each player:

- Game parameters including: the number of players, the number of clues of each type initially provided, the number of ring members, and the size of the reward and penalty for correct and incorrect identification offers
- The player’s own state including: the set of all his clues (both those bought and those provided at the start), the set of his currently outstanding offers, how much money he has, and a list of all his completed transactions (including the offered and actual price)
- Some summary information about the full game state including: a “leader board” showing a measure of progress towards identifying the ring for different players and a “transaction board” showing the number of transactions and average price per transaction for each clue type last round

The Find-the-Ring market is a variety of a double auction,

and typically a double auction market functions continuously rather than being broken into rounds [17]. However, while some of our decisions on the rules of the game are arbitrary, we have a good reason for our use of rounds. All the information needed to identify the ring exists amongst the different players at the start of the game. Without the use of rounds, the game devolves into a contest of who can generate their offers quickest. In particular, when playing with a mix of humans and automated agents, the agents can generally identify the ring long before the humans can make their first offer. The use of rounds makes strategy, rather than speed, the primary determinant of quality of play. (In the long run, for an actual intelligence information sharing market, we envision the use of a more traditional continuous double auction because (a) there will be all human participants and (b) the information will not all exist at the start but will rather arrive into the market in a continuous fashion.)

We have created an online version of the game for an arbitrary number of people and automated agents to play. Contact the authors about how to access the game. Figure 2 shows a screenshot from this online version.

5. THEORETICAL ANALYSIS

We start with an examination of how to piece together clues to identify the ring members. This will allow us to assess the value of a particular clue to a particular player at a particular time. We can then analyze a variety of buy and sell strategies, with a preliminary determination of which strategies make sense in which situations. We conclude with a theoretical investigation of some cases where the market fails, i.e. where the information to identify the ring exists but the incentives work to discourage piecing them together, and how we can modify the market to overcome these failures. Both the pricing strategies for buying and selling and the failure modes analysis have direct relevance to possible real-world information sharing markets. Hence, we see the benefit of using this “toy” market as a tool for investigating general market behavior.

Piecing Together the Clues

The key to piecing together clues is viewing the property of being in the ring as defining an equivalence class and the property of not being in the ring defining a second, complementary equivalence class. The effect of equivalence and anti-equivalence clues is to build up subsets of these equivalence classes using the following transitive relationships: (i) $a \sim b$ and $b \sim c$ imply $a \sim c$, (ii) $a \sim b$ and $b \not\sim c$ imply $a \not\sim c$, (iii) $a \not\sim b$ and $b \not\sim c$ imply $a \sim c$.

Define a *linked pair* as two sets, such that (i) any two elements in the same set are in the same equivalence class (ii) any two elements in different sets are in different equivalence classes, and (iii) the sets are of maximal size insofar as there is no element outside the linked pair for which there exists a known relationship to an element inside the linked pair. The

Your name: dave Round number: 2 Ring size: 3 Game started by: dave
 Your money: \$4688 Seconds left in round: 229 Number of players: 10 Game started at: 09:32:50

Actions

Finished for this round: Clue: Clue Type: Suspect 1: Suspect 2: Price:

 Possible Ring Members:

Clues

| Equivalences | | Antiequivalences | | Negatives |
|--------------|-------------|------------------|-------------|------------|
| Suspect 1 | Suspect 2 | Suspect 1 | Suspect 2 | Suspect |
| character47 | character45 | character47 | character21 | character6 |
| character8 | character50 | character16 | character29 | |
| character22 | character6 | character15 | character21 | |
| character4 | character25 | | | |

Game Summary

| Leader Board | | | | Transactions: Round 1 | | |
|--------------|-------|-----------------|--------|-----------------------|--|--|
| Player | Score | Clue Type | Number | Avg Price | | |
| player8 | 6 | Equivalence | 1 | 39 | | |
| player4 | 4 | Antiequivalence | 12 | 146 | | |
| dave | 3 | Negative | 0 | 0 | | |
| player3 | 3 | | | | | |

Offers

| Sell Offers (black-active red-inactive) | | | | Buy Offers (black-active red-inactive) | | | |
|---|--------------------------|-------|---|--|-------------|-------------|-------|
| Clue Type | Suspects | Price | | Clue Type | Suspect 1 | Suspect 2 | Price |
| Equivalence | character47, character45 | 30 | Cancel Deactivate | | character1 | character1 | |
| Equivalence | character8, character50 | 30 | Cancel Deactivate | | character10 | character10 | |
| Equivalence | character4, character25 | 30 | Cancel Deactivate | | character11 | character11 | |
| Antiequivalence | character47, character21 | 30 | Cancel Deactivate | | character12 | character12 | |

Transactions

| Sell Transactions | | | | Buy Transactions | | | |
|-------------------|--|-------|----------|------------------|--|-------|----------|
| Buyer | Clue | Price | Time | Seller | Clue | Price | Time |
| player4 | Antiequivalence - character47, character21 | 114 | 09:35:26 | player9 | Antiequivalence - character15, character21 | 254 | 09:35:26 |
| player5 | Equivalence - character47, character45 | 39 | 09:35:26 | player4 | Antiequivalence - character16, character29 | 242 | 09:35:26 |
| player7 | Antiequivalence - character47, character21 | 31 | 09:35:26 | | | | |

Figure 2. A screenshot from Find the Ring shows the information and actions available to a player.




Figure 3. The clues $A \sim B$, $C \not\sim D$, and $D \not\sim E$ for suspects A-F produce the pictured linked pair decomposition.

set of all elements (i.e., suspects) is divided by the equivalence and anti-equivalence relations (i.e., clues) into a disjoint set of linked pairs. Figure 3 shows an example of this. Each new equivalence or anti-equivalence clue that is not redundant (i.e., specifying a relation between two elements already in the same linked pair) acts to join two linked pairs into one via transitivity.

There are two ways to establish which subset in a linked pair has elements in the ring and which has elements not in the ring. One way is to use a negative clue that specifies that an element in the linked pair is not in the ring, hence indicating that the subset containing this element is not in the ring. The second method is counting, comparing the size of each subset to the maximum number allowed in (or not in) the ring, and excluding the set that has too many elements. As a simple example, if we know that there are exactly R ring members

and one subset of a linked pair has more than R members, then we can conclude that the too-large subset must not be in the ring.

We can get more sophisticated with counting. If each subset of a linked pair has at least k elements, then that linked pair accounts for at least k ring members. So, from the remainder of the suspects, there can be at most $R - k$ ring members. Applying this same reasoning to multiple linked pairs P_i for $i = 1, \dots, n$, where P_i accounts for at least k_i ring members, the remainder of the linked pairs can contain at most $R - \sum_{i=1}^n k_i$ ring members. In the example linked pair decomposition shown in Figure 4, knowing that there are only three ring members allows us to use this counting argument to conclude exactly which suspects are in the ring. Since linked pairs P_2 and P_3 account for two of the three ring members, then we know that C must be the ring member from linked pair P_1 . Similar reasoning identifies F and I as ring members.

In fact, if less than half the suspects are in the ring, i.e. $\frac{R}{S} < \frac{1}{2}$, one can identify a ring of size R with as little as a total of $2R$ equivalence and anti-equivalence clues. This requires that the $2R$ clues be non-redundant and result in a linked pair decomposition such that each linked pair has one subset containing exactly one more member than the other subset. This




Figure 4. If there are only three ring members, we can use counting to conclude that C, F and I are the ring members.

allows us to measure how far along a player is in finding the ring, with one “point” for each of the $2R$ clues needed. Given that the two subsets of linked pair P_i have sizes n_{i1} and n_{i2} with $n_{i1} \leq n_{i2}$, then the measurement of progress is $\sum_i [n_{i1} + \min(n_{i1}, n_{i2} - 1)]$. For the situation pictured in Figure 3, this “score” is 2, while for that pictured in Figure 4, the score is 6 (which indicates that the three ring members can be identified).

We can restate this analysis and expand it to include the case when there are also negative clues as follows. Define a *balanced* linked pair as one that has both subsets the same size and has no negative clue about any of its suspects. An *unbalanced* linked pair is one that is not balanced. The *minority* subset of an unbalanced linked pair is the one with no associated negative clue, or if neither has a negative clue, then it is the smaller subset. (The minority subset of a balanced linked pair is selected arbitrarily.) The *size* of a linked pair is the size of its minority subset and represents the minimum number of ring members contained in this linked pair. Then, the ring can be identified if all the linked pairs are unbalanced and the sum of their sizes equals the number of ring members. The score measuring progress is twice the sum of the sizes of the linked pairs minus the number of balanced linked pairs.

Given this measure of progress, we can formulate a strategy for which clues to attempt to buy. If we make offers for each of the following three classes of clues, then we are guaranteed to improve our score and progress towards a solution, assuming that sufficient clues to determine the ring do exist (which our game generator described below does guarantee) and are for sale at the right price:

1. any clue that has one suspect from a balanced linked pair and the other suspect (if not negative) from an unbalanced linked pair
2. any equivalence or negative clue that has one suspect from the minority subset of an unbalanced linked pair and the other (if not negative) from the majority subset of an unbalanced linked pair
3. any anti-equivalence clue between two suspects both in the majority subset of an unbalanced linked pair

[Actually, to be complete, there is a fourth such class: any anti-equivalence clues between two suspects both in the minority subset of an unbalanced linked pair. However, such clues are unlikely to occur and extremely unlikely to be the only clue allowing a player to advance. Therefore, it is best not to waste a buy offer on this class.]

These specifications can be directly translated into buy offers. As an example, consider the case that there are 26 suspects A-

Z and we so far have three clues $A \not\sim B$, $C \not\sim D$ and $D \sim E$. Then, we should make the offers

| Offer | Clue Types | Suspect1 | Suspect2 |
|-------|------------|----------|----------|
| 1 | A,E,N | A,B | C-Z |
| 2 | E,N | C | D-Z |
| 3 | A | D-Z | D-Z |

with the pricing discussed in the next section. Note that, clearly, if there are no balanced linked pairs, then there will be no offer 1. Similarly, if the sum of the sizes of the linked pairs equals the size of the ring, then there is no need for offers 2 or 3, since any matching clue will result in an increase of the sum of the sizes and hence cannot exist.

Buy and Sell Pricing Strategies

Having characterized what types of clues to buy, we now analyze how to determine the prices to offer for both buying and selling. For buy offers, there are multiple factors determining what a player should be willing to pay:

- the size of the reward (A sure losing strategy is to spend more gathering clues than is offered as a reward.)
- the amount of money the player possesses (A player cannot spend more than he has, so he must plan accordingly.)
- the effect of the acquisition of the specified clue on the player’s probability of winning the reward (If the player is far behind and unlikely to obtain the reward anyway, it is not worth paying much for a clue. Instead, he should focus on making money by selling clues, only buying at true bargain prices.)
- the expected additional expenses to continue buying enough clues to identify the ring (Unless this clue is the last clue, then one purchase means nothing without further buys requiring additional spending.)
- what prices other buyers are offering (During the phase when sell offers are matched with the best buy offer, out-bidding other potential buyers can result in additional purchases. This requires gauging of other players’ strategies.)
- what prices sellers are offering (Since a buy offer needs a price at least as high as a corresponding sell offer to match at all, it is important to anticipate what prices sellers will be offering.)

Of course, many of these factors are not known for certain, so a player needs to take his best guess when formulating a pricing strategy.

For sell prices, a player should consider the following factors:

- what prices buyers are offering (To match a buy offer, the sell offer price must be below the buy offer price.)
- what prices sellers are offering (To make multiple sales in the same round, a seller must underbid any competitors.)
- the amount of competition (If a clue is likely to be unique in the sense that there are not other clues providing the same information, then a good strategy is to hold out and exploit the “monopoly” by forcing buyers to pay a high price. However, if there do exist alternative clues with the

same information, a good strategy is to bid lower in an attempt to underbid competitors.)

- the effect on the race to identify the ring (Any clue you sell to another player will advance that other player closer to identifying the ring. Particularly if the clue is unique, selling the clue decreases your chances by improving your competitor's position. As discussed below, this can potentially lead to withholding of information from the market.)
- the need to finance buying (If a player has little money left, selling a clue can be a way to provide enough money to buy a clue, and this can affect how to set the price.)

Based on all these factors to consider, we can see that pricing strategy is complicated, particularly because of the need to estimate competitors' strategies and adapt one's own strategy accordingly. It is difficult even for a human player to determine a reasonable strategy. Still, we have formulated a set of pricing strategies for automated agents that takes into consideration many of the above factors and provides modifiable parameters as a means of exploring variations as a means of adaptation. (We discuss how we vary these parameters in Section 6.)

Standard Buy Pricing - The buy offer price is determined by computing four different prices BP_1 - BP_4 and taking the minimum. Before giving the formulas for these prices, we first define a few quantities used in these formulas:

- *mysteps*: the expected value of the total number of clues that this player will have to buy to identify the ring (This has to be an expected value because it is not known a priori which equivalence and negative clues initially given to the player will serve as grounding clues.)
- *avgsteps*: the expected value of the total number of clues that an average player will have to buy to identify the ring
- *steps*: the larger of *mysteps* and *avgsteps*
- *myprob*: the probability that this player will win the reward given that all players continue to purchase clues, as computed using a Monte Carlo simulation
- *compprob*: the largest probability of any of the competing players winning the reward
- *buyfactor*: a parameter used to adjust the "aggressiveness" of the buy pricing, which will generally vary between agents to provide a range of pricing strategies
- *maxfact*: the maximum value that *buyfactor* can take
- *reward*: the payout for finding the ring
- *initdoll*: the starting money for the player
- *currdoll*: the amount of money the player has

Then, the four prices from which to choose the minimum are:

- $BP_1 = (myprob \cdot buyfactor \cdot reward) / steps$
- $BP_2 = 1 + (1.1 \cdot compprob \cdot maxfact \cdot reward) / steps$
- $BP_3 = (buyfactor \cdot initdoll) / mysteps$
- $BP_4 = currdoll$

We now explain the significance of each of these four prices. BP_1 addresses the constraint that the player should on average be rewarded for pursuing the ring. To see this, observe that $myprob \cdot reward$ is the expected payoff, and hence $(myprob \cdot reward) / mysteps$ is the average amount to pay

for each clue to provide a zero expected return. (Using *steps* rather than *mysteps* takes advantage of any better-than-average luck to increase expected returns.) BP_2 addresses the constraint that a player should not bid higher than what it takes to outbid all competitors, nomatter what the clue is worth is worth to him. BP_3 addresses the constraint that a player needs to stay within his budget for purchasing clues nomatter how big the reward. If he uses too much money on early purchases, then there will not be enough left for later ones. BP_4 handles the constraint that a player cannot pay more than he has.

Standard Sell Pricing - The price of a sell offer is determined by computing two different prices SP_1 and SP_2 and taking the maximum. As before, we first define some quantities that we will use in the formulas for the prices:

- *sellfactor* and *sellfactor2*: parameters analogous to *buyfactor*
- *buyprice*: the price that the current leader would pay for the clue, making the assumption that the leader has *buyfactor* equal 1.0 and has *steps* equal *avgsteps*.
- *probdiff*: the difference in the probability that the leader will win the reward between the case when the leader buys the clue and when he does not
- *uniqueprob*: the a priori probability of a particular clue being irreplaceable, i.e. the game being unsolvable without it
- $SP_1 = sellfactor \cdot buyprice$
- $SP_2 = sellfactor2 \cdot probdiff \cdot reward$

We now explain the significance of these two prices. SP_1 sets the price somewhere between zero and the maximum other players are willing to pay (in a worst-case scenario). The parameter *sellfactor* provides the tradeoff between trying to underbid the competition and trying to increase the price. SP_2 addresses the constraint of not selling for a price less than the decrease in the player's expected reward. The quantity $probdiff \cdot reward$ is the expected decrease in the worst case that the clue is sold to the player's primary competitor for the reward and there would have been no other player willing to sell the competitor an equivalent clue. The parameter *sellfactor2* allows an adjustment for the facts that the worst-case assumption often does not hold and that the averaging of the buy and sell prices means that the actual price will generally be more than the offer price.

Stalemate Backoff Pricing - In the case that we detect a potential stalemate, based on one or more rounds with no transactions occurring, we need to reconsider our pricing strategy. All the pricing analysis above is based on the assumption that the game is a constant-sum game with a reward always given. However, in the case of a stalemate, there is no reward given, and hence we need to modify our analysis. Given that we generally have multiple rounds to detect a potential stalemate and react, we can use the standard pricing strategies until we see signs of a potential stalemate. At that point, we need to re-examine the expected payoffs and modify the prices accordingly. For a sell offer, the expected harm a player does to

his expected reward by selling a clue is zero if the reward is not given, so it is worth reducing the price beyond this point to at least try and get something. Similarly, for a buy offer, a player should be willing to raise the price as long as it does not go beyond the point where what he expects to pay for clues *from this point on* is no greater than the expected return from the reward.

However, it is important to be a good negotiator. A player should try to make concessions as small as possible while forcing the other person to make them as large as possible. There is a tradeoff: too little backoff from the original price and a player might miss a chance for a sale, too large a backoff and the other player gets most of the benefits. Therefore, we have implemented the following procedure for modifying the price as a gradual backoff procedure. Call P_1 the original (buy or sell) offer price and P_2 the worst-case price beyond which it never makes sense to go. For sell offers, P_2 is zero. For buy offers, we say that P_2 is the minimum of $myprob \cdot reward / maxStepsLeft$ and $buyfactor \cdot currdoll / maxStepsLeft$, where $maxStepsLeft$ is the maximum number of clues that the player will need to buy. Then, after $rounds$ rounds without any transaction, we set the backoff price to $P_1 + (backoffparam \cdot rounds / maxrounds) \cdot (P_2 - P_1)$, where $backoffparam$ is a parameter, generally between 0 and 1, and $maxrounds$ is the number of rounds without a transaction before stalemate is declared.

Instances of Market Failure

The properties of the game and the market vary greatly depending on the parameters defined during setup, including the number of players, the starting money, the size of the reward, the number of clues of each type, the number of ring members, etc. Here, we discuss two cases (in terms of setup parameters) where the market fails, i.e. information is not shared due to market incentives, and we suggest some simple changes to the game and market that can correct these failures. We then relate these failure modes to potential real-world market involving intelligence data, and we thus illustrate how a theoretical analysis of a simple game provides insights into a potential real information sharing market.

Low-Capital Stalemate—This type of market failure occurs when the players start the game with a small amount of money in comparison with the reward, i.e. when the the reward-to-starting-money ratio is large. The fundamental problem is that a buyer cannot pay more for a clue than he possesses, while a seller does not have incentive to sell for a price too far below the decrease in his expected reward caused by selling the clue. When the buyer’s ceiling is less than the seller’s floor, even after a “backoff” procedure like that described above, the sale will not occur. Since the money a buyer possesses is tied to the starting money and the seller’s change in expected reward is tied to the size of the reward, this situation becomes more and more common as the reward-to-starting-money ratio increases, with the flow of informa-

tion slowing accordingly.

The problem is easy to see in the following simple situation. There are only two players in game, each one clue away from a solution. Each player is understandably reluctant to sell a clue and therefore allow the other player to identify the ring and claim the reward. A player is only willing to sell a clue for a sufficiently high price that it is worth giving up his own aspirations of claiming the reward. If the reward is \$50,000, then selling a clue to the other player for \$25,000 is essentially splitting the reward, and hence makes sense from both the buyer’s and seller’s viewpoints. However, selling a clue for only a small fraction of \$25K does not make sense from the seller’s viewpoint. If neither player has more than, say, \$100 to spend on purchase of clues, then the game is in a stalemate situation. Even though each player would clearly rather split the reward than watch it go unclaimed, there is no current mechanism to allow splitting the reward if each player only has a small amount of money.

The obvious way to correct this failure is to either raise the starting money or lower the reward. This solution works fine for a simulated game like Find the Ring, but is impractical for the real world. In a real-world information sharing market, the starting money is largely fixed by what the participants have available, while the reward can only be lowered so much before people are no longer incentivized to wholeheartedly participate in the market. When the simple solution is not practical, an alternative is to invite “venture capitalists” to provide cash to market participants in exchange for a share of their rewards gained. If this is also impractical, yet another solution is to allow exchanges of information not only for actual money but also for shares of future rewards. Essentially, the market would use an analog to stock or stock options as an alternative currency. The problem with the stock-as-currency approach is that it provides a no-cost method for gathering information, providing perhaps too free a flow of information for the purposes of security and privacy. This approach requires further analysis and investigation.

Monopoly—It is not always the case that prices naturally adjust themselves to a “good” setting, i.e. high enough to encourage the gathering of information and its sale on the market but low enough to ensure that information makes its way to the people who need it. An example of this is “information monopolies”. In Find the Ring, a player can have a clue that everybody else needs, e.g. if the ring has size three and there are only a total of three anti-equivalence clues. The players originally dealt these clues know that they can hold out for a high price, high enough to discourage certain players from buying them. Similarly, in the real world instance of intelligence data, the U.S. Citizen and Immigration Services (USCIS) is the guardian of all immigration data and hence has a monopoly on this data. Therefore, its agents could raise prices on this data artificially high. When such a monopoly exists, it may make sense to “regulate” it by placing in artificial maximum on the price of sell offers for this data.

An interesting case is when the seller of an irreplaceable clue is also a buyer pursuing the ring. Then, as discussed above, that player’s incentive is to set the price high as a way to hinder others in finding the ring, hence enhancing his own chances of finding the ring first. In fact, this is what happens, as we just discussed, with a low-capital stalemate. In a situation where the players have sufficient money, there should be a way to negotiate way out of a possible stalemate with a backoff procedure. However, this does slow down progress towards a solution and potentially stops progress if the players are unwilling to back off enough.

6. EXPERIMENTAL RESULTS

We start with a description of the methodology we used for running our experiments. We then provide the results of the experiments we executed and an analysis of these results.

Experimental Methodology

Each experiment uses a single game setup, where a game setup consists of values for the number of players, number of suspects, number of ring members, size of the reward, amount of starting money, and number of each type of clue per player. An experiment tests different combinations of strategies for the given game setup, measuring the performance of the different strategies and searching for a good set of strategies for that game setup. Here, a strategy consists of a particular set of values for the parameters *buyfactor*, *sellfactor*, *sellfactor2*, and *backoffparm* defined in Section 5.

An experiment consists of multiple runs, each run measuring the performance of a particular combination of strategies in competition with each other. We specify the possible strategies as a set of possible values for each parameter plus a selection probability associated with each value. To make statistically significant comparisons, each run uses a Monte Carlo technique, executing 10,000 independent games and accumulating the average return for each strategy. (The need for roughly 10,000 trials to ensure reliable results is something that we discovered empirically, although we do provide a rationale below in our discussion of Experiment 1a.)

To enable this approach, we have developed a game generator that creates random initial configurations of clues for a game given a particular game setup. The game generator does check that the aggregate set of clues distributed among all the players does allow the ring to be identified, which means that any time the game ends in stalemate it is because the right clues were not being shared. Additionally, the game generator selects values for each of the strategy parameters randomly from the set of possible values for this parameter with no correlation between the different parameter selections.

There are two types of information that we are interested in obtaining from a run. The first is the tendency of the game to end in stalemate, i.e. *stuck*, rather than to progress to the identification of the ring. We therefore report both the percent of

games that ended stuck as well as the percent of games that were temporarily stuck (i.e., had at least one round with no transaction occurring). Secondly, we want to know the average amount of money made (defined as the ending money minus starting money) by a player executing each strategy. We report the results on average return for each possible value of each parameter averaged over all players using that value. Note that this ignores any possible effects of correlations between parameters, e.g. that a particular value of *buyfactor* may work well with particular value of *sellfactor*. However, reporting per parameter value provides the information we need while keeping the analysis simple.

An important issue is how to select the game setups for the experiments and the possible strategies for the runs within each experiment. We are limited in the number of experiments and runs we can perform, and therefore need to select these carefully. We gear our experiments towards proving (or disproving) two main hypotheses (and one corollary) about what should happen in a game under certain conditions, both of which we have discussed in Section 5. Within a single experiment, we vary the different strategies across runs with the goal of finding the best strategies or sets of strategies within the context of the game setup. We start with a random mix of strategies and use the disparities between the returns as a way of performing a “gradient” search for a good set of values. We describe this process in the context of specific experiments in our discussions of Experiments 1a and 2a below. Note that this approach for finding a good mix of strategies is simpler than the evolutionary approach suggested in [9] but is still effective.

Experiments and Results

Hypothesis 1 - When sales are competitive, sell offer prices naturally tend to be low and the market proceeds smoothly to a solution. However, when there is no competition in sales, sell offer prices naturally tend to be higher and the market has a better chance of getting stuck in a stalemate.

This hypothesis is basically a restatement of the theoretical analysis of Section 5. Our task here is to provide an experimental validation (or invalidation) of this theoretical observation. To do this, we need to use different game setups leading to different amount of sales competition and then observe the effects.

While an exact measure of the sales competition for a clue (i.e., “uniqueness” of a clue) is in general complicated, we will here utilize a rough approximation that is sufficient for our purposes. This approximate measure of uniqueness is based on the observation that the most common way for players to identify the ring is to accumulate (i) R anti-equivalence clues, one containing each of the R ring members, and (ii) R equivalence or negative clues, each matched with, i.e. containing a suspect in common with, one of the anti-equivalence clues. Then, the measure of uniqueness for anti-equivalence clues is A/R , where A is the total number of anti-equivalence

clues in the game, since this is the expected number of anti-equivalence clues associated with any particular ring member. A smaller value means a clue is more likely unique, with a value of 1 meaning all such clues are unique. The measure of uniqueness for equivalence and negative clues is $(2E + N)/S$, where E is the total number of equivalence clues, N the total number of negative clues, and S the number of suspects. For the experiments described below, we will use the anti-equivalence competition factor as the primary measure because: (a) it is the smaller of the two in all these cases and (b) it is easier to substitute an anti-equivalence clue for an equivalence clue than vice versa when deviating from the most common configuration.

To investigate the effect of sales competition on strategies and the progress of the market, we will run a variety of experiments with different game setups representing a range of different amounts of sales competition. In each experiment, we will adjust the parameters of the buy and sell strategies until they seem stable and near optimal and see how the different game setups lead to different strategies. By a *stable* set of strategies, we mean that no one strategy significantly outperforms another, hence leading to an incentive to shift strategies. One strategy parameter we will not adjust along with the others is `backoffparm`, which controls the stalemate backoff. Instead, we will leave stalemate backoff turned off (i.e., the parameter value set to 0), leaving the investigation of the selection of backoff strategies to the experiments associated with Hypothesis 2 below. The one exception to this is that we will rerun Experiment 1b with stalemate backoff turned on (`backoffparm` = 0.5) to show how great a difference this can make on pricing strategies in a market without sales competition.

In this sequence of experiments, the only aspects of the game setup we will vary are the number of players (P) and the number of ring members (R). We will keep the following values fixed:

- number of suspects = 50
- reward = \$5000
- starting money per player = \$5000
- equivalence clues per player = 10
- anti-equivalence clues per player = 1
- negative clues per player = 0

The competition factor for anti-equivalence clues is P/R and for equivalence clues is $\frac{2P}{5}$.

Experiment 1a ($P = 6, R = 3$) - In this experiment we investigate a game setup with six players and three ring members, providing a sales competition factor of $6/3 = 2$ for anti-equivalence clues. This is an example of a competitive environment.

The results from the sequence of runs of this experiment are shown in Table 1. These show how we were able to converge on a stable set of strategies, i.e. one in which no single strategy outperforms the others. In the first run, there were three

different values of `buyfactor` (each selected with equal probability): 1.3, 1.0 and 0.7. These resulted in average returns of \$933, \$859 and \$667 respectively. These clearly indicate a better return the larger the value of `buyfactor`. We therefore followed the “gradient” and used larger values of `buyfactor` for the next run. Through this process of “gradient search”, we converged on a stable set of strategies after a few runs. The results from the final run are included also in Table 2 along with the results from the other experiments for Hypothesis 1.

We can now see why we needed 10,000 games for each run. The standard deviation of the average return for a strategy is very roughly \$2,500 for a single game. Differences on the order of \$25-\$50 are significant when looking for a gradient to follow. By averaging over 10,000 games, we obtain a standard deviation of roughly $2500/\sqrt{10,000} = 25$, which is enough to conclude that gradients are statistically significant.

The results show that, in this competitive sales environment, there is strong pressure on buyers to bid high to obtain as many clues as possible and win the race to the reward. There is also strong pressure on sellers to bid low to make as many sales as possible. The fact that the broker averages the buy and sell offer prices to find the sales price is what allows buyers to bid above their break-even point (with `buyfactor` well above 1) and sellers to offer to give away their clues. The high buy offer prices and low sell offer prices keep the information flowing as desired, with the result that the market does not get stuck in a stalemate, even temporarily, even once in the 10,000 runs.

Experiment 1b ($P = 3, R = 3$) - This is an example of a non-competitive sales environment. With three players and three ring members, the sales competition factor for anti-equivalence clues is $3/3 = 1$. We actually ran this experiment twice, once the standard way without stalemate backoff (`backoffparm` = 0) and the other time with stalemate backoff (`backoffparm` = 0.5). The results of both are shown in Table 2.

The first thing to notice about the results is that the players and the market are behaving very differently than in the competitive case. With optimal values of `sellfactor` well above 0.0, there is clearly advantage for sellers to ask for a higher price from buyers without fear of being undercut. With optimal values of `sellfactor2` so high, there is also a big incentive to protect one’s own position relative to the reward by asking for higher prices on sales. Even in the case of no backoff, sellers hold out to the extent that over 18% of the games end in stalemate.

As we can see from the second version of this experiment, with stalemate backoff on, the only thing keeping sale offer prices as low as they are is the threat of stalemate. With the ability to negotiate out of stalemate, sell asking prices rise much higher still. The sell prices are so high that over 91% of the games reach the point of temporary stalemate even though

| | | buyfactor | | | sellfactor | | | sellfactor2 | | | stuck | temp |
|-------|--------|-----------|-------|-------|------------|-------|-------|-------------|-------|-------|-------|-------|
| Run 1 | Value | 1.3 | 1.0 | 0.7 | 0.8 | 0.5 | 0.2 | 0.8 | 0.5 | 0.2 | 1.64% | 2.48% |
| | Return | \$933 | \$859 | \$667 | \$729 | \$816 | \$911 | \$802 | \$799 | \$857 | | |
| Run 2 | Value | 1.6 | 1.3 | 1.0 | 0.6 | 0.3 | 0.0 | 0.8 | 0.5 | 0.2 | 0.43% | 0.57% |
| | Return | \$870 | \$825 | \$794 | \$738 | \$862 | \$889 | \$762 | \$834 | \$892 | | |
| Run 3 | Value | 1.8 | 1.5 | 1.2 | 0.2 | 0.1 | 0.0 | 0.4 | 0.2 | 0.0 | 0.00% | 0.01% |
| | Return | \$867 | \$835 | \$796 | \$769 | \$828 | \$904 | \$758 | \$815 | \$926 | | |
| Run 4 | Value | 2.1 | 1.8 | 1.5 | 0.0 | | | 0.0 | | | 0.00% | 0.00% |
| | Return | \$834 | \$831 | \$833 | \$833 | | | \$833 | | | | |

Table 1. The results from the sequence of runs of Experiment 1a illustrate how the “gradient search” proceeds to a stable set of strategies.

| | | buyfactor | | | sellfactor | | | sellfactor2 | | | stuck | temp |
|----------------------|--------|-----------|------|------|------------|------|------|-------------|------|------|--------|--------|
| Expt 1a | Value | 2.1 | 1.8 | 1.5 | 0.0 | | | 0.0 | | | 0.00% | 0.00% |
| | Return | 834 | 831 | 833 | 833 | | | 833 | | | | |
| Expt 1b | Value | 1.6 | 1.4 | 1.2 | 0.8 | 0.6 | 0.4 | 0.8 | 0.7 | 0.6 | 18.70% | 22.56% |
| | Return | 1349 | 1352 | 1362 | 1357 | 1356 | 1351 | 1339 | 1374 | 1351 | | |
| Expt 1b (backoff) | Value | 2.0 | 1.8 | 1.6 | 1.1 | 0.9 | 0.7 | 1.9 | 1.7 | 1.5 | 2.52% | 91.39% |
| | Return | 1644 | 1623 | 1606 | 1620 | 1616 | 1637 | 1636 | 1600 | 1637 | | |
| Expt 1c | Value | 1.8 | 1.5 | 1.2 | 0.0 | | | 0.0 | | | 0.00% | 0.00% |
| | Return | 989 | 1015 | 994 | 1000 | | | 1000 | | | | |
| Expt 1d | Value | 1.7 | 1.5 | 1.3 | 0.5 | 0.3 | 0.1 | 0.9 | 0.7 | 0.5 | 6.52% | 8.11% |
| | Return | 1177 | 1172 | 1155 | 1143 | 1177 | 1184 | 1174 | 1168 | 1162 | | |
| Expt 1e | Value | 1.9 | 1.6 | 1.3 | 0.6 | 0.3 | 0.0 | 1.1 | 0.8 | 0.5 | 9.30% | 11.97% |
| | Return | 910 | 890 | 920 | 903 | 903 | 914 | 925 | 889 | 906 | | |

Table 2. The results from the final run of each experiment associated with Hypothesis 1 show a stable set of strategies for each game setup.

only 4% end that way. This confirms experimentally our analysis of Section 5 about how players who have a chance to win the reward will only sell if the price is sufficiently high to make it worthwhile to give the buyer, who is likely a competitor for the reward, an advantage.

Experiment 1c ($P = 5, R = 3$) - In the next three experiments, we explore the transition from a competitive to a non-competitive sales environment by examining anti-equivalence competition factors between 2 and 1. The game setup for this experiment has five players and three ring members, leading to a competition factor of $5/3$. The results of the experiment, given in Table 2, are very similar to the results from Experiment 1a, including sales offer prices of 0 and 0.0% games stuck in stalemate. Hence, we can conclude that this is still in the competitive realm.

Experiment 1d ($P = 4, R = 3$) - With four players and three ring members, the anti-equivalence competition factor is $4/3$. Based on the results shown in Table 2, we can conclude that this game setup is in the non-competitive realm. We see the same basic characteristics as in Experiment 1b, non-zero sales prices and a non-zero fraction of the games stuck in stalemate. Note that a far lower fraction of the games end in stalemate (6.5% vs. 18.7%), which is what we might expect given that

the competition factor is greater than for Experiment 1b.

Experiment 1e ($P = 5, R = 4$) - The previous four experiments all had game setups with three ring members, so here we vary this to show that the difference in market behavior is actually a function of the competition factor and not the number of players. In Experiment 1c, we showed a game with five players that was competitive. However, from the results in Table 2, we can see that this game setup is clearly in the non-competitive realm. With a competition factor of $5/4$, we would expect the fraction of games stuck to lie between those for Experiments 1b and 1d, and that is indeed the case.

Hypothesis 2 - When players have sufficient money, they will naturally tend to avoid stalemates through suitable initial pricing and negotiation (i.e., backoff). When they do not have sufficient money, the game will tend to get stuck in stalemates.

This hypothesis restates the ideas about stalemate and its avoidance discussed in Section 5 and the stalemate backoff analysis of Section 5. We attempt to validate this statement using a sequence of four experiments.

The first two experiments investigate the tendency of a game

to get into potential stalemate situations as a function of the reward-to-starting-money ratio. These two experiments both use the same game setup except with the starting money drastically reduced in the second experiment. In these first two experiments, the stalemate backoff is turned off in order to concentrate on the effects of the initial sell pricing, in particular the effects of `sellfactor2`. (Recall that `sellfactor2` is the parameter that controls the use of the sell offer price as a means to enhance a player's own chances of winning the reward.)

In the second two experiments, we investigate the effectiveness of the stalemate backoff procedure at avoiding stalemates. Here, we use the same two game setups as for the first two experiments. The difference is that this time we allow `backoffparm` to be non-zero and investigate the effect of different distributions of values for this parameter.

Experiment 2a - Despite the new name, this is exactly the same as Experiment 1b without stalemate backoff. Recall that the game setup has three players and 3 ring members. Because there is so little sales competition, this game setup has a strong tendency towards stalemate. As shown in Table 3 (as repeated from Table 2), with a stable set of strategies there is a 18.7% chance of having the game end in stalemate.

Experiment 2b - This has the same game setup as Experiment 2a except now with starting money equal to \$100 instead of \$5000. We keep the distributions of values of `buyfactor` and `sellfactor` the same as for Experiment 2a and just vary the distribution of the values of `sellfactor2`.

Because of the high reward-to-starting-money ratio, the effect of varying `sellfactor2` is essentially on-off. If `sellfactor2` equals 0, then the player is not letting the fact that selling a clue aids a competitor ever stop him from selling a clue. If `sellfactor2` \geq 0, then the player is not selling to anyone except when he is so far out of the running that he has no chance of winning. Therefore, instead of using different values of `sellfactor2` in the different runs, we instead vary the fraction of players using either a value of 0.0 or a value of 1.0.

Table 3 shows the results. We can see that with a stable mix of strategies, a large majority (7/8) of the players use `sellfactor2` equal 1, what in the language of the Prisoner's Dilemma would be called "defecting". Only a small minority (1/8) use `sellfactor2` equal 0, what would be called "cooperating". This leads to a very high probability (92.9%) of games ending in stalemate. This is a clear indication, before even considering stalemate backoff and just considering sell strategies, that too large a reward-to-starting-money ratio makes the situation much worse in terms of the tendency of games to end in stalemate. In the next two experiments, we will see that adding a stalemate backoff just amplifies this difference.

As a final note, we show a rough theoretical analysis that would allow us to predict approximately the mix of cooper-

ation and defection that we observed in this experiment. Let us assume that the only way that a player has a chance to win the reward is if both other players are cooperating. If one of the three players is defecting and the other two cooperating, then the defecting player always wins the reward. Hence, the expected return for a defecting player is $5000 * (prob)^2$, where `prob` is the probability that a particular player is cooperating. If all three players are cooperating, then they each have an equal chance of winning the reward. Hence, the expected earnings from the reward for a cooperating player is $(5000/3) * (prob)^2$. However, the cooperating player also makes money by selling clues. Let us assume that the sales price is about $\$100/4 = \25 , since each player needs on average a little less than four clues to find the ring and can afford to spend all his money. Then, a cooperating player's total return is $2(25) + (5000/3) * (prob)^2$. The expected total returns of defecting and cooperating players are equal when `prob` is approximately 1/8.

Experiment 2c - This experiment uses the same exact game setup as Experiment 2a. The difference is that this time we allow `backoffparm` to vary, and the primary emphasis of the different runs is to view the effect of different distributions of this parameter. The results are shown in Table 4.

The results show that, in order to optimize their own personal rewards, players tend to have a non-zero value of `backoffparm`, i.e. are willing to negotiate out of stalemate. This willingness to negotiate is enough to almost always allow stalemate to be averted; in fact, only 0.3% of games end in stalemate. There does exist a tradeoff between backing off not enough and ending up losing all the reward in stalemate and backing off too much and letting the other player reap a bigger portion of the reward. This experiment shows that the incentive to negotiate out of stalemate is so strong that stalemate rarely occurs.

Experiment 2d - This experiment is the same as Experiment 2c except with starting money of \$100. It is also the same as Experiment 2b except with `backoffparm` varying, not `sellfactor2`. The aim is to determine a stable mix of stalemate backoff strategies in the low-capital (i.e., high reward-to-starting-money) case.

In this case, since the sell offer price tends to be so many times greater than the maximum a player can offer as a buy price, partial backoff tends to be ineffective. Instead, the seller needs to back off essentially 100% for potential stalemate to be averted. Hence, we have the same all-or-nothing type of situation as in Experiment 2b, so we will again select two fixed values for the parameter and vary the probability of selecting them. The two fixed values for `backoffparm` are 0.0 and 1.25. A value of 1.25 for `backoffparm` represents a full backoff, since the maximum discount is that of $(4/5) * \text{backoffparm}$ offered after the fourth consecutive round of potential stalemate.

| | selfactor2 A | | | selfactor2 B | | | selfactor2 C | | | stuck | temp |
|---------|--------------|------|--------|--------------|------|--------|--------------|------|--------|--------|--------|
| | Value | Prob | Return | Value | Prob | Return | Value | Prob | Return | | |
| Expt 2a | 0.8 | 1/3 | 1339 | 0.7 | 1/3 | 1374 | 0.6 | 1/3 | 1351 | 18.70% | 22.56% |
| Expt 2b | 1.0 | 7/8 | 117 | 0.0 | 1/8 | 121 | | | | 92.92% | 92.92% |

Table 3. The results from the final runs of Experiments 2a and 2b showing a stable set of sell strategies in the high-capital and low-capital cases respectively.

| | backoffparm A | | | backoffparm B | | | backoffparm C | | | stuck | temp |
|---------|---------------|------|--------|---------------|-------|--------|---------------|------|--------|--------|--------|
| | Value | Prob | Return | Value | Prob | Return | Value | Prob | Return | | |
| Expt 2c | 0.7 | 1/3 | 1662 | 0.5 | 1/3 | 1668 | 0.3 | 1/3 | 1655 | 0.29% | 22.69% |
| Expt 2d | 1.25 | 1/13 | 232 | 0.0 | 12/13 | 234 | | | | 85.96% | 92.49% |

Table 4. The results from the final runs of Experiments 2c and 2d showing a stable set of stalemate backoff strategies in the high-capital and low-capital cases respectively.

The results are shown in Table 4. In a stable mix of strategies, a large majority (12/13) of the players will not back off (i.e., defect), hoping for another player to instead back off for them. A small minority (1/13) reap the small but assured payoff from backing off and making a sale. The result is that the games continue to usually end in stalemate, with 86% ending thus. This shows very clearly the detrimental effect of a high reward-to-starting-money ratio on the tendency of the game to end with the ring identified.

Hypothesis 3 - A change in the broker mechanism will change the specifics of the pricing strategies for a given scenario but will not change the basic conclusions of Hypotheses 1 and 2.

Recall that the standard broker mechanism we have selected for Find the Ring works as follows. First, let each buyer look for the lowest-price seller for one offer; if a seller is found, carry out the transaction at the average of the buy and sell offer prices. Next, let each seller look for the highest-price buyer; if a buyer is found, average the prices.

For the current set of experiments, we modify the broker in two ways. One change is that we now just let buyers look for sellers and not vice versa. A second change is that the transaction price is now that offered by the seller, not the average. This method of resolving buy and sell offers is similar to that of a typical goods and services market.

We have rerun Experiments 1a-1e with this new broker mechanism and called the new versions Experiments 3a-3e. The results are shown in Table 5. Examining the results, we find that indeed there are different pricing strategies adopted by the agents in response to this different broker mechanism. There are also some different numbers for the most important statistic, the stalemate probability, particularly between Experiments 1d and 3d. However, we still see similar qualitative behavior as the game setups transition from a competitive to a non-competitive sales environment: an increase in the sell offer prices and the probability of stalemate becoming non-

zero. The same game setups as before are competitive (1a/3a and 1c/3c) and non-competitive (1b/3b, 1d/3d and 1e/3e).

General Conclusions - We do need to be careful about drawing too-broad conclusions about information sharing markets in general from the results. The specific numbers we obtained are highly dependent on the rules of the game we have defined. Still, the combination of the theoretical analysis with the experimental results provide strong support for the two primary hypotheses, which are stated in very general terms. The knowledge about information sharing markets thus obtained can potentially be valuable in the future when designing real-world markets.

7. DISCUSSION

We now answer a few key questions about our work in greater detail than we were able to provide above.

Do the experiments with artificial software agents provide an accurate indication of what would happen with human players? We claim that the stable set of agent strategies approximate reasonably accurately what human strategies are once the humans learn how to play the game well. There are enough degrees of freedom to allow the strategies to adapt to the pressures of the marketplace. The tuning process does seem to find nearly optimal sets of strategies. Informal experiments with humans playing agents with strategies tuned to the particular game bears out that the agents play the game very well. So, the results for the tuned agents should approximate closely the results for experienced human players. Using agents provides an easy way to gather large enough quantities of data to produce statistically significant conclusions.

What the artificial agents approach does not capture is the effect on the market of inexperienced players. The tuning process for agents does not at all mimic the human learning process. Since we expect at least some of the players to be novices, we need to understand the effect of “irrational” actions on the market, and we can only do this by observing

| | | buyfactor | | | sellfactor | | | sellfactor2 | | | stuck | temp |
|---------|--------|-----------|------|------|------------|------|-------|-------------|------|------|--------|--------|
| Expt 3a | Value | 1.8 | 1.6 | 1.4 | 0.2 | 0.03 | 0.004 | 0.1 | 0.04 | 0.01 | 0.00% | 0.00% |
| | Return | 841 | 835 | 823 | 823 | 831 | 844 | 841 | 834 | 824 | | |
| Expt 3b | Value | 2.3 | 2.0 | 1.7 | 1.0 | 0.8 | 0.6 | 1.0 | 0.8 | 0.6 | 17.18% | 26.98% |
| | Return | 1382 | 1347 | 1411 | 1367 | 1419 | 1354 | 1370 | 1394 | 1376 | | |
| Expt 3c | Value | 1.8 | 1.6 | 1.4 | 0.3 | 0.13 | 0.05 | 0.2 | 0.05 | 0.01 | 0.00% | 0.00% |
| | Return | 1004 | 985 | 1010 | 994 | 1018 | 987 | 993 | 1024 | 981 | | |
| Expt 3d | Value | 2.3 | 2.0 | 1.7 | 0.6 | 0.4 | 0.2 | 0.7 | 0.5 | 0.3 | 0.57% | 1.69% |
| | Return | 1252 | 1232 | 1243 | 1243 | 1250 | 1233 | 1252 | 1226 | 1248 | | |
| Expt 3e | Value | 2.6 | 2.3 | 2.0 | 0.7 | 0.5 | 0.3 | 1.0 | 0.8 | 0.6 | 5.65% | 7.82% |
| | Return | 945 | 958 | 926 | 900 | 987 | 942 | 930 | 957 | 942 | | |

Table 5. Experiments 3a-3e have the same game setup as the corresponding Experiments 1a-1e. The difference in result shows the effects of changing the broker mechanism.

humans.

Why do the experiments use stalemate as the primary criterion for market performance? The goal of an information sharing market is to encourage players to share information to the point where an important conclusion is reached. For the game Find the Ring, this conclusion is the identification of the ring. Hence, the Find-the-Ring market succeeds in its goal if the ring is identified and fails if game ends with the ring still unidentified, i.e. in stalemate. (Similarly, for the terrorism intelligence marketplace described above, the success of the market is determined by whether the pieces are fit together to uncover the plot.) Hence, frequency of stalemate is the primary measure of the success of the market.

Note that this is different from information aggregation (futures) markets, which have the goal of predicting the outcome of a particular event. Hence, the predictive power of such markets is the primary measure of their success. Because the two types of markets have different goals, they need to be judged on different criteria.

Is the lack of uncertainty in the information/clues a fundamental flaw of the game and thus the experiments? It is true that in the real world, information is uncertain, and so not including uncertain information in our Find the Ring game was a simplification. However, we claim that this is a secondary issue, with the primary issue being the piecing together of clues, with or without uncertainty. In the example above of the terrorism plot, it was known with a high degree of certainty that particular individuals had bought tickets for a flight, had enrolled in flight school, etc, but these facts just had not been connected.

Information uncertainty is definitely an issue worth exploring in the future. It does change strategies about what information to buy, e.g. sometimes making it worthwhile to buy “redundant” information to increase confidence in a conclusion. It also makes the penalty for wrong identifications more important. Perhaps most interestingly, it introduces the possibility of intentional spreading of false information and how to pre-

vent it, potentially through the use of a reputation system.

Why do the experiments not address the issue of not sharing the information too widely despite our claiming this as an important feature of the approach? While it is true that our approach does discourage too free a flow of information, there is nothing explicit in the market mechanisms preventing market participants from buying information for the wrong purposes. For example, the market does not prevent a spy from taking the money provided by a foreign government and using his access to the market to buy information desired by this government. Preventing wrongdoing requires a layer of monitoring and auditing of the marketplace not addressed in this paper.

What an information sharing market does provide is a much more restricted and purposeful flow of information than in an open database. This enables the auditing of all the transactions, both manually and by automated means (such as those used by credit card companies to check for fraudulent purchases).

What can we learn from Find the Ring, and what can we not learn? Find the Ring provides a simplified abstract version of an information sharing market. This makes it a good testbed for investigating economic and game theoretic issues because it eliminates many of the other factors that would exist in a real-world information sharing market and would complicate the analysis. The shortcoming of Find the Ring is that it does not allow investigation of some of these other issues. One particularly important such issue is how a potential buyer can assess the value of a piece of information before knowing what the information is. In Find the Ring, a mathematical analysis provides a way for the buyer to fully specify a mathematical set of criteria identifying the information he needs. However, in the real world, this problem does not yield to such a tidy solution.

When does it make sense to use an information sharing market rather than another technique? One key condition is that the information should be more than the sum of

its parts. (Otherwise, an information aggregation market is likely a better alternative.) The second key condition is that it is hard for one person to find another person with information needs complementary to his. (If the information is concentrated in a small group of people, it is usually better to just gather the group in a room.) Other conditions that are good for an information sharing market are (i) existing disincentives to sharing, (ii) the need to protect sensitive information, and (iii) the need for output of more than just true-false or multiple-choice answers.

8. CONCLUSIONS AND FUTURE WORK

In this paper, we have defined the concept of an information sharing market. Such a market allows individuals with complementary information to merge this information so as to discover conclusions not available in the separate pieces of information. This is particularly important when these individuals are separated by barriers such as organizational boundaries. An information sharing market provides a mechanism for such individuals to locate each other and incentivizes them to share their information, while still providing safeguards to protect security and privacy and discourage excessive distribution of information.

We have introduced the game Find the Ring as a simple example of an information sharing market. It has served as a useful tool for investigating such markets. Because the game is well constrained, it is amenable to both theoretical analysis and experimental investigation. We have used both theory and experiment to investigate certain properties of the game. While some conclusions are specific to this particular market, we also discovered some important properties that should apply generally to information sharing markets.

For short-term future work, there are some incremental changes that can be made to the game Find the Ring for further investigating properties of information sharing markets. One such change is the introduction of shares of the reward as a form of currency, which would provide a possible solution to the low-capital stalemate issue. Another possible enhancement discussed above is the introduction of uncertainty in the information.

A bigger, long-term step for the future is to introduce a different game. As stated above, use of the game Find the Ring limits our ability to investigate certain aspects of real-world information sharing markets. We therefore would like to create a more realistic game involving an information sharing market. An example would be a game with (simulated) terrorist intelligence data bought and sold in a market with the players acting as agents of intelligence and crime-fighting organizations.

REFERENCES

- [1] Adler, M. and E. Ziglio: 1996, *Gazing into the Oracle: the Delphi Method and its Application to Social Policy and Public Health*. London: Jessica Kingsley.
- [2] Anderson, L. and C. Holt: 1997, 'Information Cascades in the Laboratory'. *American Economic Review* **87**(5), 847–862.
- [3] Bert, J., R. Forsythe, F. Nelson, and T. Reitz: 2001, 'Results from a Dozen Years of Election Futures Market Research'. In: C. Plott and V. Smith (eds.): *Handbook of Experimental Economic Results*. Berlin: Elsevier Press.
- [4] Chen, K.-Y., L. Fine, and B. Huberman: 2003, 'Predicting the Future'. *Information Systems Frontiers* **5**(1), 47–61.
- [5] Hanson, R.: 1999, 'Decision Markets'. *IEEE Intelligent Systems* **14**(3), 16–19.
- [6] Hanson, R.: 2003, 'Combinatorial Information Market Design'. *Information Systems Frontiers* **5**(1), 105–119.
- [7] Harman, J.: 2002, 'Terrorism and Homeland Security'. <http://www.house.gov/harman/terrorism/infoshar.html>.
- [8] Kumagai, J.: 2003, 'Mission Impossible? (Special Report: Intelligence and Technology)'. *IEEE Spectrum* **40**(4), 27–31.
- [9] LeBaron, B.: 2001, 'A Builder's Guide to Agent-based Financial Markets'. *Quantitative Finance* **1**(2), 254–261.
- [10] LeBaron, B., W. B. Arthur, and R. Palmer: 1999, 'Time Series Properties of an Artificial Stock Market Model'. *Journal of Economic Dynamics and Control* **23**, 1487–1516.
- [11] Lo, A., N. Chan, B. LeBaron, and T. Poggio: 1999, 'Information Dissemination and Aggregation in Asset Markets with Simple Intelligent Traders'. *Computing in Economics and Finance*, Vol. 653.
- [12] Noeth, M., C. Camerer, C. Plott, and M. Webber: 1999, 'Information Aggregation in Experimental Asset Markets: Traps and Misaligned Beliefs'. Social Science Working Paper 1060, CalTech.
- [13] Pennock, D., S. Lawrence, L. Giles, and F. Nielsen: 2001, 'The Real Power of Artificial Markets'. *Science* **291**(5506), 987–988.
- [14] Pennock, D. and M. Wellman: 1997, 'Representing Aggregate Belief through the Competitive Equilibrium of a Securities Market'. *Thirteenth Conference on Uncertainty in Artificial Intelligence*. pp. 392–400.
- [15] Resnick, P., R. Zeckhauser, E. Friedman, and K. Kuwabara: 2000, 'Reputation Systems'. *Communications of the ACM* **43**(12), 45–48.
- [16] Roll, R.: 1984, 'Orange Juice and Weather'. *American Economic Review* **74**, 861–880.
- [17] Smith, E., J. D. Farmer, L. Gillemot, and S. Krishnamurthy: 2003, 'Statistical Theory of the Continuous Double Auction'. *Quantitative Finance* **3**, 481–514.