International Conference on Intelligence Analysis, McLean, VA, May 2-6, 2005.

A Distributed Dynamic Intelligence Aggregation Method

Mark Boslough
Evolutionary Computing & Agent Based Modeling Dept.
Sandia National Laboratories
Albuquerque, NM, 87185-0318, USA
mbboslo@sandia.gov

Keywords: Multiple Competing Hypotheses, Information Sharing and Collaboration, HUMINT, OSINT, terrorism

Abstract

This paper describes an intelligence aggregation system for generating actionable knowledge by addressing the dynamical and distributed nature of the problem. Ever-changing and compartmentalized information can be used for evaluating hypotheses by trading among analysts. Trading rules are designed for self-assembly of metadata structures that attach actual data to the trades, allowing analysts to associate hypotheses with the raw data. *Adaptive aggregation* refers to built-in error correction by weighting the most current and relevant information, thereby addressing the dynamic aspect. The raw data and sources of information remain distributed as required for reasons of security, privacy, or turf. Knowledge is generated collectively using a system of hypothesis generation, investment, and probability discovery through trading.

1. Introduction

The problem of intelligence analysis has been compared to the construction of a jigsaw puzzle from a box with missing pieces and with pieces that belong to other puzzles. This analogy is useful, but does not go far enough because it neglects the fact that the problem is both *distributed* and *dynamic*. The distributed nature means that multiple puzzles must be solved simultaneously by many people using multiple boxes of pieces. Moreover, the capacity for exchanging puzzle pieces and other information is limited. The dynamic nature means that nothing is static. Not only do the puzzles change with time, but the pieces and the assemblers do not stay the same either.

Distribution of data and expertise is an intrinsic property of intelligence gathering. There are many intelligence-gathering agencies in the United States, often with differing and even conflicting purposes, scopes, classification, methods, and cultures. Information gathered by law-enforcement and public health agencies compound this compartmentalization. Effective collaboration is a difficult and elusive goal for domestic organizations. When international agencies are included, collaboration is hindered by even more barriers.

Much progress is underway to improve data collection, distribution, and mining methods. The ultimate problem with this approach is that it is not scalable. Even if all the relevant information were available to an expert panel of analysts, it would be too much for them to digest. An effective means of "swarm intelligence" is required to overcome non-scalability, as well as the distributed and dynamic properties of the data. James Surowiecki, the author of a 2004 book The Wisdom of Crowds (Surowiecki, 2004), describes the circumstances under which a large group is smarter than a few experts:

> "There are four key qualities that make a crowd smart. It needs to be diverse, so that people are bringing different pieces of information to the table. It needs to be decentralized, so that no one at the top is dictating the crowd's answer. It needs a way of summarizing people's opinions into one collective verdict. And the people in the crowd need to be independent, so that they pay attention mostly to their own information, and not worrying about what everyone around them thinks."

According to the 9/11 report, "The biggest impediment to all-source analysis… is the human resistance to sharing information." This is true even when institutional impediments are removed, because it is based on human nature. Information is valuable and people are reluctant to "give it away for free," even when it is their job to do so. Market trading is one method that has a long historical record of discovering useful information by aggregating distributed information amongst competitors. Financial markets are successful at price discovery because they harness human nature by rewarding information exchange, and the reward is directly related to the relevance and value of the information. By modifying market institutions for the purpose of adaptive intelligence aggregation, we argue that actionable knowledge can be generated in real time by competing analysts and agencies by exploiting currently-existing motivations and behaviors.

2. Hypothesis evaluation

Financial markets are designed to allow the exchange of assets such as stocks, bonds, currencies, and commodities at a price that is deemed to be fair to both buyer and seller. Successful markets are very efficient at discovering the fair price, despite the fact that information is decentralized. In recent years, markets have been extended to derivatives and event outcomes, such as presidential elections, sporting events, and weather. In the latter, the assets that are bought and sold are contracts that pay out to the owner, depending on the state of the world on the expiration date. A correctly-priced contract would be an accurate evaluation of the probability of some hypothesis about the future state of the world based on distributed information.

The rules of trading in financial markets are defined by a *market institution* that specifies which types of bids and other messages are legal and the details of the mechanism by which trades are executed. These institutions have evolved over the years to efficiently determine the price according to the intersection of supply and demand curves of an asset in the absence of a "central planner" with complete information. The existence of efficient markets that can successfully discover prices (and therefore evaluate hypotheses) by aggregating decentralized information suggests that market institutions can be engineered for the purpose of aggregating intelligence information. The efficiency of financial markets is based on market institutions that rely only on the self-interest of the individual traders, as opposed to "the common good". Nevertheless, there is no fundamental reason that market institutions cannot be modified to optimize a common goal of generating knowledge by collectively assigning relevance to information, in addition to price discovery (or hypothesis evaluation), while retaining the attributes that reward self-interest in order to make them effective.

For markets with purposes other than asset trading, the space of market institutions increases dramatically to include features that would not be possible to implement in financial markets. For example, in typical financial markets there are two commodities, one is the *good*, and the other is the *money*. A self-interested trader uses the market institution to trade money for goods (and vice-versa) in an attempt to maximize her net utility. If the goal on the other hand is intelligence aggregation, the traders are analysts who are employed to evaluate hypotheses by trading. In this case the market institution could be modified such that the money commodity is replaced by a point system, in which points translate to reputation, status, and future influence (for both the trader and his institution), for a motivational reward.

If the purpose of a market is asset trading, the operational costs of the market are borne by the traders in the form of transaction fees. For an effective intelligence aggregation system, the transaction fee might be replaced by a rebate that can be applied to subsequent trades as a means of encouraging and rewarding active exchange of information through trading by analysts.

2.1 Market Institutions

Development of new market institutions for intelligence aggregation requires an understanding of market microstructure theory, which is the study of the process and results of trading assets under the rules of specific market institutions (O'Hara, 1995). This field grew rapidly in the 1990s as researchers attempted to learn how prices emerge while buyers and sellers trade assets, and how the emergence of new price knowledge depends on the trading rules.

Historically, the price-setting rules of market institutions were not prescribed by market engineers. They developed over the years, and have continued to evolve as technology has changed the ability and speed by which traders are able to obtain price information, interact, and place orders. Much of the current research in the field is focused on how traders make use of price information to update their assessment of value and adjust their bidding behavior. Information flows into the market by means of the set of admissible messages specified by the market institution. Admissible messages are the actions that traders are allowed to execute. Most market institutions limit the messages to prices and quantities that buyers want to bid, and sellers want to ask.

Market institutions further define how a commodity is allocated amongst buyers and sellers and how the price is settled for a given combination of messages. Examples of widely used market institutions are the clearinghouse (CH) and the continuous double auction (DA).

The CH institution is a two-sided auction that is discrete in time, so all traders move synchronously in steps from a given allocation to the next allocation. Messages (bids from buyers and asks from sellers) are collected and bundled during the trading interval. Trades are executed periodically by clearing the market using an algorithm that matches buy and sell orders at a common price that is determined by the overlap in supply and demand curves that are defined by the orders.

The DA institution is a two-sided auction that is continuous in time. Admissible messages are bid and ask offers for a single unit of an asset, and acceptances of existing offers. Acceptance messages lead directly to the execution of a trade between a buyer and a seller, which can occur at any time. The New York and Chicago exchanges all use some form of the continuous double auction institution.

Notably absent from the set of admissible messages in both of these conventional market institutions is any information exchange other than price and quantity. Nevertheless, these market institutions are remarkably efficient at price discovery, which is directly related to hypothesis evaluation when the assets being traded are contracts for some future state of the world. For example, futures contracts on orange juice are based on information and expertise that is distributed among traders. One trader may have special knowledge about weather forecasts for Florida, another may know about political instability of other citrus-producing nations, and a third trader might have performed a detailed analysis of projected fuel prices. However, none of this detailed information is exchanged. Knowledge is generated by projecting this information onto the hypothesis statement (in this case, some future price of orange juice) through admissible messages, which are limited to trade offers.

Despite the extremely narrow communications bandwidth imposed by the market institution, laboratory experiments have shown that even with small numbers of traders with limited information, DA markets are consistently good at price discovery (or knowledge generation) and efficient allocation of goods (Friedman, 1993). The fact that markets are so successful has been called a "scientific mystery" by the top economists (Smith, 1982) and there is still no complete theoretical understanding of why they work so well.

2.2 Asymmetric Information

In most of the literature on asset pricing, it is assumed that traders have access to the same information. This assumption simplifies the theory, but does not reflect the fact that different people have different information and hold different opinions about the information they do have in common. The theoretical underpinnings of actual markets must take into consideration the fact that information is distributed. According to Brunnermeier (2001), the fact that information is dispersed among many traders means that prices have a duel function. Prices provide a measure of scarcity or bargaining power, and also serve as a conveyor of information.

The information encapsulated by prices feeds back into buyers' and sellers' beliefs about the value of the goods that are being traded. Prices carry information about what everyone else thinks the value should be. Because traders have a financial stake in their beliefs, prices are weighted in favor of those who think they have the best or most current private information, or who have the most confidence in their opinions about public information. Market participants base their decisions not only on their own private information and beliefs, but on the information communicated by publicly-available prices and price histories.

In the real world, prices are affected by news and information, and are constantly changing. Because of the feedbacks conveyed by price, markets are nonlinear dynamic systems. The flow of information causes traders to constantly re-evaluate their beliefs and make decisions to buy or sell, thereby conveying information about their changing beliefs to other traders. Even participants who have no new information, or conflicting information, are aware by observing price histories that others are responding to changes in their own information set. Without having any private information whatsoever, an observer can infer from the price history when relevant new information has become available, without the information itself being revealed. This attribute of markets suggests that they can be used to evaluate hypotheses that are defined at a high level, while protecting information that must remain compartmentalized for security or other reasons.

3. Information to knowledge

The job of the intelligence analyst is to take data and extract information that can be used to generate actionable knowledge. Intelligence information is asymmetrically distributed for a variety of reasons. Because of the hierarchical structure of intelligence agencies, the conversion of information to knowledge tends to take place in parallel, with little if any communication. This "stovepiping" is analogous to the way in which data flows in massively parallel computing environments for problems that are "embarrassingly parallel." This class of problems requires no temporal synchronization and little if any communication among computation nodes until data assimilation at the end of the run. A problem is embarrassingly parallel if it has a structure that allows it to be decomposed into parts that have computational independence (i.e. the intermediate results of one part of the problem do not affect other parts). Intelligence problems do not tend to be embarrassingly parallel, so full compartmentalization of analysis is inappropriate. Low bandwidth message passing using a protocol defined by a specially engineered market institution may provide a method for aggregation of intelligence information into knowledge in a way that retains necessary compartmentalization of data.

3.1 State of the world

Brunnermeier (2001) describes a model for how asymmetrically distributed information can be assembled into knowledge.

Individual analysts do not have complete information, and cannot independently determine the state of the world, ω, which contains a complete description of a system. For purposes of illustration, the world can be represented by a standard six-sided gaming die. A die placed on a flat surface and oriented to the compass directions has 24 possible states (for each of the six possible upfacing numbers, there are four possible rotations).

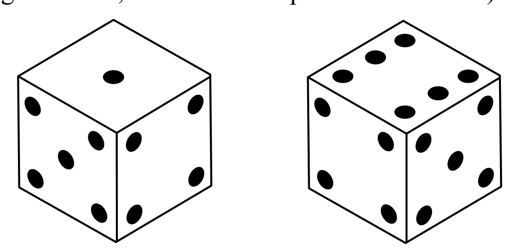

Figure 1. Two possible states (ω_{154} and ω_{645}) of the "world".

Figure 1 shows two possible states of the "world" defined by a single die. To establish a convention, assume that north is to the upper left, so the vantage point for this illustration from slightly above the southwest corner of the die. These two states can be represented by the symbols ω_{154} and ω_{645} respectively, where the three indices refer to the upper, western, and southern face values, respectively. The set of all 24 possible orientations is the state space, Ω.

An analyst with a particular perspective has limited information, from which she can rule out particular states of this world. Suppose the world is actually in state ω_{154} (the position shown on the left-hand side of Figure 1), and Alice can only view the southern face of the cube. She is immediately able to eliminate 20 states, and Alice's possibility set is $P_A(\omega_{154}) = \{\omega_{154}, \omega_{564}, \omega_{624}, \omega_{214}\}$. Alice can work out a possibility set for every possible observation she can make from her viewing angle, and impose an *axiom of truth*, which is a statement about her knowledge about the true state of the world. The axiom of truth simply says that the true state ω must be a member of the possibility set $P_A(\omega)$.

An event E is a set of states. For example, the event E_6 can be defined as a roll of the die, which comes up six. This event is true regardless of the orientation of the die with respect to the compass, so it consists of a set of states. In other words, $E = \{\omega_{645}, \omega_{624}, \omega_{632}, \omega_{653}\}$. A knowledge operator K(E) can be defined to provide a different representation of an analyst's information. Suppose Alice wants to list all the possible states she can observe, from her southern vantage point, from which she can determine that the world is in event E_6. That would be Alice's knowledge operator, $K_A(E_6) = \{\ \}$. For Alice there is no state of the world that can give her complete certainty about event E_6.

Consider Bob, whose instruments can simultaneously view the western and southern faces, but cannot distinguish faces with values above four. Bob's knowledge operator is $K_B(E_6) = \{\omega_{624}, \omega_{632}\}$. Intelligence problems of interest tend to have the property that all analysts have null knowledge operators. The events of interest do not have observable signatures from which they can be unambiguously deduced.

3.2 Hierarchical knowledge discovery

Extending the example, suppose Carol is yet another analyst, and that the corners of the cube are painted eight different colors. Carol has a table of which colors correspond to which corners (defined by adjacent faces), and can see the color of the upper southwest corner, but nothing else. She sees a green corner, and her chart tells her that it is adjacent to the four, five, and six faces. Carol's possibility set is $P_C(\omega) = \{\omega_{645}, \omega_{564}, \omega_{456}\}$. Her vantage point, color table and observational data are classified, but event estimates are classified at a lower level so her agency circulates a report indicating that the probability of E_6 is 33%.

Bob makes a simultaneous measurement with his flawed instrument from his southwestern perspective, and sees that the western face is a four, but the southern face is ambiguous (either a five or a six). Bob's possibility set is $P_B(\omega) = \{\omega_{645}, \omega_{246}\}$. Bob's position and instrumental limitations are classified, but he is able to get his 50% probability estimate for event E_6 to colleagues on a need-to-know basis. Alice sees a five on the south face, calculates her possibility set to be $P_A(\omega) = \{\omega_{645}, \omega_{415}, \omega_{135}, \omega_{365}\}$, and her assessment of 25% for event E_6.

At the top of the hierarchy, decision makers look at the various numbers as independent estimates and dismiss the likelihood of E_6 as no more than half. From their perspective, there is insufficient information about the state of the world to act.

3.3 Distributed knowledge discovery

Suppose on the other hand that Alice, Bob, and Carol are permitted to communicate by a system that allows them to pass messages to one another in the form of bids and offers on 24 contracts for each of the possible states of the world. Perhaps there are many other analysts with different types of information and different beliefs and levels of confidence. Uninformed traders might estimate that each of the 24 states has an equal probability of occurring, and place their trade orders accordingly. Misinformed traders might bid more points on contracts for states that the better-informed analysts have eliminated. Bob's best strategy would be to place orders to buy shares of contracts on the two states in his possibility set, and sell short the other 22 state contracts.

As trading continues, it will reach a competitive equilibrium in which the price of the true state (ω_{645}) approaches the face value of the contract, and the other 23 approach zero. When the true state of the world is finally revealed, the analysts who hold the correct contract are rewarded with the payoff. They also get to keep the points they gained by short-selling contracts for states they had eliminated from their possibility sets. The analyst with the best information in this example was Bob, whose trading strategy would have allowed him to accumulate the most points and thereby have more buying power and influence in subsequent auctions. His gains were at the expense of misinformed, uninformed, less-informed, or incorrect analysts, whose losses reduce their ability to influence the generation of knowledge. Analysts (and groups of analysts) who are consistently able to convert their information to relevant knowledge will continue to become stronger voices, whereas the voices of consistently incorrect or misinformed analysts will become increasingly muted.

3.4 Scalability of distributed analysis

In the simple single-die example, the true state of the world would be obvious to anyone with access to all the data. This problem was contrived to be simple enough for a single individual to understand at an intuitive level, decompose into possibility sets, and assimilate the full set of compartmentalized observational data. Many real-world problems are not amenable to analysis by a single individual, due to the massive quantity of data and specialized knowledge that is necessary to extract useful information and interpret it. The market institution provides the message-passing protocol that allows knowledge to emerge from distributed data, even when the data set is so large that it would overload any single analyst if it were available in its entirety. The scalability of parallel analysis also has an analog in computer science.

Gustafson (1988) first described what is now known as "Gustafson's Law" which has become the paradigm for scalability of massively parallel computing and overcame pessimistic predictions of its potential power.

The earlier pessimistic viewpoint was expressed by Amdahl's Law, which stated that the parallel efficiency of a program is subject to severe limitations when there is a significant proportion of serial code (code that cannot be parallelized and must be executed by every processor). As the number of processors tends toward infinity, the serial fraction dominates each processor's time and nearly all of the work is redundant. The analog in the intelligence problem with massive data would be the situation where every analyst has access to everyone else's data, and works independently toward an estimate of every aspect of "state of the world," rather than keeping both data and analysis distributed. In this extreme case, the analog to the parallel component of code is the collection of the data, and the serial component is the requirement that every analyst must arrive at the same possibility set with all the data. Not only is the efficiency of the system severely limited, but there is a maximum problem size that can be analyzed beyond which a single analyst cannot comprehend (analogous to memory-bounded problem in parallel computing).

By keeping data and analysis distributed, an analog to Gustafson's Law becomes operative. For sufficiently large problems, the efficiency can be improved by increasing the number of processors. In the intelligence analog, this means that individual analysts can efficiently work on their part of the problem by combining their limited data set with the common (serial) data and determine their own possibility set. The class of problems that is most amenable for distributed analysis is that for which the sum of all data and processing requirements would overwhelm a single analyst. The answer emerges from a continuous aggregation of the distributed possibility sets by passing messages defined by a market institution.

3.5 Re-engineered market institutions

The clearinghouse and continuous double auction market institutions are widely used because they efficiently allocate goods and are successful at price discovery. These institutions rely on traders' self-interested incentive to maximize their own utility. As mentioned earlier, these are not the only possible institutions.

There is another body of literature devoted to "mechanism design" which treats the specific rules of trading as variables that can be optimized. Mechanism design is based on the foundations of game theory, and attempts to maximize variables such as market efficiency or expected revenue for traders. Mechanism design is limited by the need to make assumptions about the trading environment and distribution of information.

Efficient intelligence aggregation requires mechanism design that optimizes the ability to evaluate hypotheses that lead to actionable knowledge, while maintaining the incentive compatibility that leads to strong participation by analysts. Employed analysts have different incentives to succeed than traders in financial markets, and mechanism design needs to account for this difference in motivation and behavior.

Traders in financial markets who have private information have an incentive to keep that information to themselves. They have a motivation to buy undervalued or sell overvalued assets. An effective market institution for intelligence aggregation must remove this incentive by employing a revelation requirement. This requirement must reward analysts for making information available, while promoting the ultimate goal of generating actionable knowledge.

One approach to revelation would allow traders to keep information private until a trade, based on that information, is executed. This would encourage analysts to reveal data only when they think it is significant enough to change the value of a contract (i.e. affect the evaluation of a hypothesis). Both classified and unclassified information can be referenced by means of unclassified pointers. The set of time-stamped pointers yield a self-assembled metadata structure associated with each hypothesis, and allows high-level decision makers (with a need to know) to focus only on the data that actually influenced the valuation at any given time.

As dynamic and distributed information and analysis becomes available, the price history of a contract will unfold. Metadata that are associated with sudden price moves are those that other analysts perceive to be important, and decision-makers have access to them in real time.

Metadata revelation is one potential attribute of a market institution that is engineered to optimize intelligence aggregation. Using the process of mechanism design to develop a market system for this purpose will require experimental markets and agent-based modeling to more fully explore the space of possible market institutions.

3.6 Real-world example: WMD in Iraq

The real world is, of course, much messier than the "world" defined above for the single-die example. Analysts aren't simply getting peeks at numbers on the sides of a cube. In general, there is no way for them to develop explicit possibility sets and knowledge operators. The mathematical formalism that has been developed to understand real markets, nevertheless, has explanatory power. It provides theoretical justification for using market methods to aggregate information.

The question of whether or not weapons of mass destruction existed in Iraq in the year 2003 would have been amenable to an interagency intelligence aggregation market. Unlike the 24 possible "states of the world" in the idealized single-die problem, the WMD problem could have been decomposed to two possible states with a single hypothesis statement: "The existence of Iraqi WMD will be confirmed by the end of the year." With appropriate definitions and criteria for confirmation, this statement could have become an interagency contract from which the real-time quantitative evaluations (and links to the information on which they were based) could have been made continuously available to decision makers. Analysts with independent information or understanding that was discounted by formal reporting channels (e.g. aluminum tubes as evidence for WMD development) would have reason to believe that contracts were overpriced. They would have placed sell orders, and any associated drop in the hypothesis evaluation (as measured by contract price) would have been linked to the aluminum tube analysis. At the end of 2003, contracts sold by skeptical analysts would have evaluated to zero, and those analysts would have kept the points they earned by selling to competitors who had ignored their data.

4. Implications for intelligence

Intelligence aggregation has the potential to solve many of the major challenge problems defined by the Intelligence Technology Innovation Center (ITIC) Knowledge Discovery and Dissemination (KDD) research program.

Cross-media intelligence value estimation

The hierarchical level of knowledge assembly is higher than that of the body of data. Information can be extracted and analyzed by domain experts who then transform it to a hypothesis statement that can be used to define a contract that can be traded with other market participants. Hypotheses provide the common interface for information exchange through price discovery. Participating analysts provide cross-media intelligence value estimates by offering and bidding on hypotheses that make use of the data within their domain, regardless of the form.

Confidence measures of uncertainty

The primary purpose of the intelligence market is price discovery, which is a direct measurement of aggregate confidence. An engineered market institution can define a protocol for revelation, attaching metadata in a way the leads to the self-assembly of a data structure. Participating analysts can invest and divest in hypotheses that are created by the market. Hypotheses have values that are quantitative and are emergent properties of the collective wisdom of all the analysts. The most important hypotheses are those that have high value and volume and lead to actionable knowledge. Decision-makers can monitor a steady "ticker" stream of quantitative price and volume data associated with the hypotheses that are of interest to them.

Evidence-Hypothesis-Action spaces

The linkages between evidence and hypotheses are provided through trading histories and metadata. Values of hypotheses are constantly evolving, but the links to data are preserved. When analysts announce a change of opinions by divesting in a hypothesis, the market responds immediately. This trading information can be used by other analysts to revisit the original data on a need-to-know basis to keep raw data compartmentalized. Likewise, the decision-makers can use trading metadata to connect actions to original intelligence.

Human elements of efficient analytic teams

Hypotheses are high-level statements, so analysts with different backgrounds and with access to orthogonal data sets can collaborate on price discovery in an open auction system. Implementation of a merit-point system provides the incentive for analysts from different intelligence organizations to bid the values of various hypotheses values up and down. High-level collaboration emerges as a consequence of individual and organizational competition in a way that is directly analogous to price discovery in commodities markets. Motivations and behaviors such as competition and information compartmentalization are exploited in a way that is beneficial. Investors in useful hypotheses that turn out to be true are rewarded for their "good" behavior by collecting points from those who divested. Investors in false hypotheses are penalized by forfeiting their points to the divested analysts. Those who have a good track record accumulate the most points and therefore have more influence in subsequent hypothesis formation and evaluation.

5. Conclusions

An interagency intelligence aggregation market, using an engineered market institution that is optimized for the purpose of hypothesis evaluation, can provide a means for continuously updated quantitative estimates based on collective analysis of distributed information. The method is most effective if diversity, decentralization, independence of thought, and distribution of analysis are maintained.

Acknowledgments

This research was performed at Sandia National Laboratories. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000. Report number: SAND2005-0995C.

References

Brunnermeier, M. 2001. *Asset Pricing under Asymmetric Information*. Oxford University Press.

Friedman, D. 1993. The Double Auction Market Institution: A Survey. Friedman, D. and Rust J. (eds.) *The Double Auction Market*. Addison-Wesley. pp. 3-25.

Gustafson, J. 1988. Reevaluating Amdahl's Law. *Communications of the ACM* 31(5): 532-533.

O'Hara, M. 1995. *Market Microstructure Theory*. Blackwell Publishing, Ltd.

Smith, V. 1982. Microeconomic Systems as an Experimental Science. *Amer. Econ. Rev.* 72: 1119-1152.

Surowiecki, J. 2004. *The Wisdom of Crowds*. Doubleday.