

Trade-offs in Internal Information Markets

Eric Schultz

April 20, 2006

There have been several prominent information markets some of which are currently in operation. The most prominent markets include the Iowa Electronic Markets, Foresight Exchange, Hollywood Stock Exchange, and TradeSports, and the propositions traded range from political, scientific, economic, sports, and cultural events (Passmore and Cebeci 2005). Many corporations are taking note of the value of information markets and experimenting with them as a tool for decision-making. For instance, HP experimented with markets to forecast sales and Eli Lilly experimented with predicting successful drugs (Hahn and Tetlock 2006).

Information markets are much like commodity futures markets with the notable exception that the goods traded are not physical but information. Typically information markets trade on the prediction of some future state of the world (Passmore and Cebeci 2005). To understand how information markets may work in organizations it's important to understand the role information plays in markets. Hayek first identified information's role and has persuaded economists to believe that information markets can indeed successfully aggregate knowledge distributed amongst market participants (Passmore and Cebeci 2005).

Complete knowledge about particular circumstances, local conditions, temporary opportunities and relative importance never exist in a integrated form but are dispersed amongst all individuals in incomplete sometimes contradictory bits. Every individual possesses some of these bits that are unique and advantageous over everyone else. Markets solve the problem of utilization of dispersed knowledge that does not belong to one individual or is terribly inefficient to be handled by a central authority, and the market price is the aggregated result of every individual

using his or her unique knowledge in bargaining or trading (Hayek 1945).

Organizations, much like the markets Hayek describes, also have relevant bits and pieces of intuition, local data, insights and other subjective information distributed throughout (Plott and Chen 2002). This type of knowledge is called tacit knowledge, and Choo elaborates that tacit knowledge is hard to transfer or verbalize (Choo 2006). Markets provide the mechanism for aggregating intuitions, hunches, and insights as well as cultural and explicit knowledge. This isn't to say that markets provide the solution to all knowledge management issues but can be very useful in certain contexts.

Hirshleifer further explains the role of information as that which changes an individual's belief distribution over possible states of the world (Hirshleifer 1973). Information markets primarily trade belief distributions (or predictions), but a common theme in Information Economics is that market prices also reveal private information held by the participants (Manski 2005). This dynamic of market price and revealed information creates a feedback loop where market price can simultaneously be used by traders to update their beliefs, and updated beliefs change market prices making them more accurate (Gjerstad 2005). In addition, the potential profits from trading also provide incentives for information discovery that sustains the feedback loop. It may even be the case that traders believe that the market price never fully reflects all of the market's knowledge, and in equilibrium, inefficiency in pricing is sufficient to induce traders to discover more information. This feedback loop also causes market prices to respond rapidly to new information (Wolfers and Zitzewitz 2004). For organizations this is a promising way to transfer knowledge to all of the

participants as well as being adaptable to a changing environment.

Hayek and Hirshleifer provide a good foundation for understanding information in markets, but we can extend on this using Weick's sensemaking model to understand the role of markets in organizations. Weick describes several properties in which individuals make sense of their environment, and it's reasonable that markets and prices are mechanisms that exhibit similar properties. Sensemaking is grounded in identity construction. The environment is what individuals react to and project themselves into while simultaneously constructing their identity from. The market is one such environment where individuals can confront what they know and who they are (or could be) as well as taking action that can shape the market. Sensemaking is retrospective in that individuals select from alternative meanings in order make sense of past events. In markets individuals construct plausible reasons to explain what occurred in past transactions with the market. It is also not until the market has stopped trading an asset that individuals *truly* know what to think of their actions. Why did I lose on that prediction? Could I have looked for better information? I won! I must have a good strategy. Sensemaking is enactive in that individuals produce part of their environment. This is almost self-evident in markets since they are almost entirely formed based individual actions and values. Sensemaking is a social activity where individuals react with or anticipate the actions of a group for understanding. Individuals participate in markets by trading on the anticipation of behaviors of others, submitting assets that they think others will value, or acquiring knowledge from those that have it. Sensemaking is an ongoing continuous flow that never starts or stops. Participants in markets are in a constant state of adjusting their beliefs in order reconcile what they see as the market price. Sensemaking is focused

on and by extracted cues, points of reference for linking ideas into meaning. Similarly, markets provide cues (e.g, price and bidding behaviors) as well as inducing individuals to seek out other bits of information in order to make better predictions. Sensemaking is driven by plausibility not accuracy. No market participant can know all of knowledge that the entire market does. Participants have to confront the market with their bits of incomplete information and make the most plausible prediction.

In addition to aggregating knowledge and facilitating sensemaking, organizations can utilize information markets to improve forecasting, decision-making, and risk management (Hahn and Tetlock 2006). Passmore and Cebeci propose other alternative uses. Information markets can be used to forecast sales that predict whether a new product will generate an appropriate amount of revenue. Information markets can be used to determine the expected impact of a new product or service. This may be very useful in circumstances where expert opinion, surveys, polls are used but provide little value except that there is no alternative. This is to say market predictions of a new cancer drug may not be suitable or at the very least palatable, but this doesn't necessarily exclude that market predictions may be suitable for other products. Information markets can also be used to manage projects in order determine accurate deadlines, determine the likelihood of political, social, or economic events that may impact the success of products or services, and identify expertise (Passmore and Cebeci 2005). Corporations are trending towards greater experimentation with markets, which may indicate many other unforeseen uses may emerge (Hahn and Tetlock 2006).

Information markets also may help ameliorate issues plaguing common organizational

information aggregation techniques (Plott and Chen 2002). Organizations have created many ways to answer questions about future events such as surveys, focus groups, expert opinions, and business meetings. These techniques can be too costly to conduct continuously and are usually implemented without much thought regarding quality or appropriateness (Passmore and Cebeci 2005). Capturing subjective information is difficult; business meetings are inefficient due to location, space, and timing, and quotas and budgets give incentives to individuals to hide information (Plott and Chen 2002). Cass Sunstein also indicates that there are other issues when aggregating knowledge through group decisions. These issues include amplification of individual errors, excessive focus on common knowledge where unshared key information induces groups to select inferior options, over-reliance on influential individuals or the majority in spite of private knowledge, and group polarization (Hahn and Tetlock 2006). The variety of processes developed for improving information flow suggests, though, that this is important for organizational success (Plott and Chen 2002).

In 1996 a joint research project between Charles Plott, Caltech, and HP demonstrated the potential of internal information markets. The research project was an experiment incorporating an information market into HP for the purpose of predicting sales forecasts. The experiment consisted of twelve predictions conducted over a three-year period and the participants were expected to predict monthly sales for a month three months in the future. Participants were anonymous, geographically diverse, and came from Marketing, Finance, and the HP Labs. The market predictions were compared to the official HP forecast, a management tool which quotas are measured and compensation paid. 6 out of 8 comparisons showed that market predictions were

more accurate than the official sales forecast. This experiment provides encouraging results for implementing internal information markets.(Plott and Chen 2002).

Designing internal information market entails more than what can be done technically. Technology alone will not induce individuals to gather, share, or transfer knowledge (Davenport and Prusak 1998), and market participants need knowledge. Market participants need to have differentiated knowledge while not being completely uniformed about the assets being traded. If all the participants have the same common knowledge about a given asset then there is little that can be achieved through the market. Common knowledge means that there's no information for the market to aggregate (Plott and Chen 2002).

On the other hand, participants need to share a common language in order to share knowledge (Davenport and Prusak 1998). Plott emphasizes some uninformed traders may improve trading in the market, but when the market consists of all uniformed traders no trading will occur (Plott and Chen 2002). Participants confronted with an unknown asset will be unable to credibly evaluate or understand it without supporting information resources or a common language. The intuition is straight forward. For instance, if Marketing begins trading an asset it may be very difficult for other business units such Accounting or R&D to credibly understand what exactly the asset means. More importantly, if these other business units have no available resources to utilize in order to resolve the uncertainty of Marketing's asset then they will cease to trade on a what could be valuable asset. Organization that predominately find a source of efficiency through organizational boundaries and autonomous business units will need to confront creating a common

language that supports the market.

Davenport and Prusak emphasizes the necessity for organizations to dedicate resources in developing new knowledge and transferring that knowledge to others that may benefit from it. For instance, R&D may be vary adept to creating new knowledge, but organizational boundaries and language barriers may inhibit other business units from utilizing that new knowledge. Establishing a market and individuals who take on the role of knowledge brokers may ameliorate inefficient knowledge transfer. Davenport and Prusak further explains the role of knowledge brokers. These individuals are usually people who enjoy making connections between people allowing them to transfer knowledge between each other. Librarians are usually conducive to this type of role. In a study conducted of NYNEX's corporate library activities, the activity users valued most from the library was knowledge brokering. Knowledge brokering may be an odd concept in traditional organizations but is a prevalent activity supporting markets. Davenport and Prusak also explain that typically the individuals that facilitate knowledge brokering (e.g. librarians) are viewed as "gossipy" and "nosy" and are typically undervalued (Davenport and Prusak 1998).

Davenport and Prusak views the organizational market as informal, transparent, and existing without any type of explicit pricing mechanism (Davenport and Prusak 1998). This obviously does not have to be the case. Establishing a formal market allows undervalued knowledge brokers to visibly display their value to the organization through several activities. First, motivated by the expectation of making profits, knowledge brokers facilitate the knowledge transfer others need in order to become good market traders. For instance, brokers may buy into

assets earlier on and then produce and disseminate knowledge in order to drive the price up (Hahn and Tetlock 2006). Second, knowledge brokers are most likely going to be the individuals that earn the most in the market. Lastly, the organization can credibly identify the individuals that have real expertise in being knowledge entrepreneurs. The last activity is important when organizations need to identify who should be assigned as a full-time knowledge broker or any other role that requires boundary-spanning knowledge-transferring activities. How do we identify these individuals? They are most likely going to be the people that accumulate the most wealth through the market. This a reasonable assumption in that market wealth correlates well with individuals that are well versed in gathering information and connecting with others for knowledge transfer. Organizations can determine who has superior knowledge or who is better at gathering information by keeping track of profits and final holdings (Chen, Fine and Huberman 2003). Hahn and Tetlock support this claim as well. Traders that have the most valuable information make bigger bets and influence prices the most (Hahn and Tetlock 2006).

Christoffel, Franke, and Kotkamp conducted a simulation study that further establishes the value and implications of knowledge brokers, or intermediaries, in information markets. The most encouraging result from the study was that knowledge brokers created a clear and permanent increase in utility for all traders (Christoffel, Franke and Kotkamp 2001).

There are other implications for knowledge brokers that are worth noting for organizations looking to establish individuals with these roles. Knowledge brokers that become specialized earn less profit. The intuition here is that as traders become more informed on a specific assets they rely

less on brokers for trading (Christoffel, Franke and Kotkamp 2001). Organizations may benefit more from establishing brokers that are specifically *not* divided by area of specialization. For instance, having specific marketing brokers, accounting brokers, and R&D brokers will not be as useful in the market and may only reemphasize the organization boundaries that inhibit knowledge transfer in the first place.

Christoffel, Franke, and Kotkamp also indicate that brokers gain more profit from continual renewal of traders. Again, traders become less reliant over time on brokers as they learn about assets and the market (Christoffel, Franke and Kotkamp 2001). This results indicates that organizations may look for mechanisms that help brokers meet new traders. This certainly could be accomplished in many ways. For instance, the organization may establish an "ask a broker?" feature, and queries are sent to the most available broker ensuring a quasi-random connection between traders and brokers. Davenport and Prusak suggest creating an informal marketplace such as meetings, conferences, or presentations where traders and brokers can meet and transfer knowledge (Davenport and Prusak 1998). In addition, training brokers in the importance of making new connections between traders and brokers and allowing them to establish their own ways of accomplishing this may suffice.

We could also extend Christoffel, Franke, and Kotkamp results in order to derive some interesting insights. Brokers facilitate a common understanding between disparate business units allowing trading to occur in the market. Over time business units become aware of pertinent knowledge transferred from other business units. Ideally, a business unit may begin to filter ideas

not only based on what they know but also what they know other business units have said in past experiences. We can imagine this may create more selective robust ideas over time. In addition, in an idealized scenario we could envision that knowledge transfer facilitated by brokers is begins to be allocated efficiently throughout the organization. For instance, Marketing doesn't necessarily need to know everything that R&D knows, but it's not obvious what R&D knowledge Marketing does need to know and vice-versa. Brokers can help establish an efficient allocation of knowledge transfer allowing business units to coordinate and understand each other's concerns.

Another important issue that arises in designing internal information markets is the role of money. In particular, does "play money" work as well as real money? Organizations would obviously prefer a design that uses play money if the same results could be generated. On the other hand, it is often believed by economists that markets where traders must risk real money will be more accurate.

There are some other implications that make play money more appealing to organizations. Executive management will be more likely to follow market predictions if the cost of paying participants is lower than the value of the prediction. It might be plausible that if executive management had to pay a hefty amount of money for a market prediction they may opt to enact the alternative in order to avoid the pay off. For instance, the market is tasked to predict whether the organization should implement X or Y technology, and the successful market participants are paid on the technology that is implemented. The majority of the market says X technology should be implemented, but executive management realizes the payoff would be too large in regards to the

value of implementing such technology, and, thus chooses the alternative. This would create confusion or disappointment amongst the participants.

Another implication is that the predictions traded on the market are handpicked based on some discernible value that is greater than the payoff. This means that any predictions where the value of knowing the prediction can not be credibly determined are not allowed to be traded. In addition, the value of a prediction would need to be known to executive management before it would be allowed to enter the market. This almost completely counters the value of an information market to aggregate ideas and predictions that aren't obvious to executive management. Play money enables a wider variety of assets to be traded since pay-off is not such a dominant factor.

Alternatively, the information market could be designed so that participants are expected to use their own money for trading. This, though, is unlikely to be successful. Participants that are willing to take the risk of losing money in the market would more than likely opt to take their money to an outside more financially lucrative market (e.g. the stock market). This means that those employees that are good at making predictions will exit the market, and the internal market would be left with those that are only average at making predictions. In addition, participants may more willing to engage the market if they are not confronted with the risk of actually losing money. Some participants are risk-averse but could contribute good valuations only if the market does not require they use their own money.

Another issue regarding money is that the organization would have to incur substantial

costs in creating a system for moving employee's market money around. Unless the organization is already a financial institution, this would be a costly system for what is essentially a non-core capability. In addition, mistakes, system failure, and fraud become serious issues when operating an information market using real money. At the end, these costs can quickly exceed the value created through the market (Servan-Schreiber, Wolfers, Pennock and Galebach 2004).

Pennock, Lawrence, Giles, and Nielson conducted a study of the Foresight Exchange, an artificial market used to predict the likelihood of unresolved science questions being discovered. This analysis examined historical price information generated through the FX market for questions that could be answered definitely "yes" or "no" after the market closed. This study showed that price, despite being "play money", did in fact correlate well with the observed outcomes. For instance, high prices indicated higher likelihood of discovery and, of course, lower prices indicated lower likelihood of discovery (Pennock, Lawrence, Giles, and Nielsen 2001).

Servan-Shreiber, Wolfers, Pennock, and Galebach also show that play money performs just as well as real money. Real money is generally regarded as a mechanism which provides large incentives for information discovery by forcing individuals to bet on outcomes they truly are confident in. Servan-Shreiber, Wolfers, Pennock, and Galebach suggest the amount traders are willing to bet might also result from how wealthy the individual is. This means that individuals can place larger bets using wealth generated through other means such as other skills, luck, or inheritance. In markets without real money traders can only amass wealth based solely on the ability to predict outcomes. The incentives to individuals in a play money market are usually based

on reputation and such incentives can still drive intense trading yielding accurate predictions (Servan-Schreiber, Wolfers, Pennock, and Galebach 2004).

Servan-Shreiber, Wolfers, Pennock, and Galebach compared results from two popular sports trading information markets: TradeSports.com and NewsFutures.com. TradeSports.com is a real-money site based in Ireland, and NewsFutures.com is a play-money site based in the United States. Both sites are similar in most aspects such as contracts, media exposure, and number of participants. Most importantly, the sites differ mainly on the use of real money or play money. In both markets traders can bid anywhere from 0 to 100 reflecting the probability of a particular team winning. Out of 208 games studied, 65.9% of TradeSports' teams won and 66.8% of NewsFutures teams won, and respectively the average trading prices were 65.1 and 65.6. Servan-Shreiber, Wolfers, Pennock, and Galebach's analysis is more elaborate, but their research in general indicates that the prediction accuracy of both markets is not affected by the use of real money or play money (Servan-Schreiber, Wolfers, Pennock, and Galebach 2004).

Organizations may have many incentives to use play money, and recent research indicates that prediction accuracy is not inhibited by the use of play money. There may be some possible negative implications, though, that haven't addressed. Markets based on play money may be perceived by participants as entertainment or having low business value and, thus, do not engage in or take the market seriously. Some market designs such as NewsFutures compromise by offering limited prizes to winners. It may still be difficult to convince employees that an artificial market with fun prizes has real business value especially when faced against other activities that have been

established as worthwhile.

Let us consider an example described by Daveport and Prusak that may clarify why employees may not value the market. Davenport and Prusak describe a scenario involving Mobil Oil where the Mobil Oil engineers discovered a serious cost saving technique that reduced the amount of steam needed in drilling. They realized the potential and decided to pass on the innovation to other drilling operations via a memo. Not surprisingly no one adopted the new technique. The information manager reasoned afterward that a memo was not an effective medium for transferring this type of knowledge. Another interpretation is that the value of the technique could not be conveyed through a memo. If the technique was truly cost-saving why did the manager choose the least costly form of communication, a memo? Why not validate the technique with a form of communication such as a video presentation that has some cost associated with it? Any manager can make claims through a memo, but only managers that have true cost saving techniques will expend some cost (e.g. a video presentation) in order to convince the other drilling operations of the validity of the technique. Similarly, employees may enact the same logic when confronted with an artificial market with prizes. If the market was so truly revenue generating or cost saving why is the organization offering limited prizes and play money?

The results in play money versus real money are at best conflicting, and there is little guidance as to which is the best design. Most likely, though, this really depends on the type of organization the market exists in. Real money may work well in organizations where employees are motivated by real money. For instance, organizations that are oriented around a sales force earning

commissions on sales generated might use real money. Play money may work well in organizations driven by creativity, entertainment, or abstract thinking. Organizations in between may just have a luke-warm engagement with the information market.

Even when markets are designed with consideration of supporting information resources and appropriate pay-offs participants may still find it in their best interest to exploit the market. Market manipulation occurs when individuals choose not to reveal his/her true valuation, and these differences in valuations lead to price distortions. If the price distortions are large enough then the ability of the market to predict outcomes is compromised in favor of an individual's or group's objective. Ideally, participants attribute their *true* value to an asset on the market. Information markets will only predict as well as the value individuals attribute to assets and while there are other reasons where information markets predict sub-optimally (e.g. market thinness and inadequate knowledge of participants) lack of revealed true valuations is the Achilles' heel of a would be working market.

How serious is misrepresentation of value? This issue is paramount when we consider decision-making and policy-making that is contingent on market predictions. If manipulation is successful then some individuals or groups can indirectly influence decisions or policies that are in their best interest but not in the best interest of the entire organization. In the famous information markets being studied, such as the Iowa Electronic Market, the market predictions are generally independent from the outcomes. This is to say that what president is predicted to be elected by the Iowa Electronic Market probably does not have much influence on which president is actually

elected. On the other hand, Defense Advanced Research Projects Agency planned a prediction market called Policy Analysis Market (PAM) that would aid policy makers in predicting geopolitical events. In this market decision-making would have become contingent on market prediction, and the program was shut down by the Senate precisely over concerns of indirect policy manipulation by terrorists (Hanson, Oprea, and Porter 2006).

Hanson, Oprea, and Porter claim that market price manipulators are unable to distort prices effectively. Generally, any price distortions in one direction is overcompensated by other traders in the opposite direction. Hanson, Oprea, and Porter constructed a laboratory experiment to test how well manipulators could distort the market price. In this experiment subjects were endowed with experimental currency, allowed to bid on the value of an asset, and could redeem real money that was converted from experimental currency after the experiment. The asset was assigned one of three different values, and the subjects were differentially informed on what value the asset would *not* take. Once aggregated the subjects had enough information to predict the true value of the asset. In addition, some subjects received an additional payoff based on the median contract price of the asset giving these subjects an incentive to manipulate the price upward by overstating his/her true valuation. Lastly, two experiments were conducted one in which all subjects were aware that some subjects may have an incentive to manipulate the market price, and one where the other subjects were not aware of this (Hanson, Oprea, and Porter 2006).

Hanon, Oprea, and Porter's experiment led to some interesting results. First, and not surprisingly, subjects that had an incentive to manipulate the market price did in fact try to

manipulate the price. Secondly, manipulation did not seem to have any effect on the market price. Lastly, non-manipulators tend to accept trades equal to or lower than expected. According to Hanson, Oprea, and Porter's experiment individuals will alter their trading behavior if they are aware that others may be trying to manipulate the market or at the very least have an incentive to do so. These augmented behaviors tend to cancel each other out and does not seem to affect the market price (Hanson, Oprea, and Porter 2006).

Other theoretical research by Hanson and Oprea indicates that manipulators may in fact increase the accuracy of the market. Generally traders that seek to manipulate the market can indirectly induce others to become more informed, and, thus, these informed traders receive greater payoffs. Hanson and Oprea's model illustrates a thin market where traders can obtain more information with some effort, a manipulator that has a preference for the market price to be at a target price, and the other traders have a noisy clue of what the manipulator's target price is. On average, market accuracy will increase if the manipulator's bias is within the range of biases that other trader's suspect (Hanson and Oprea 2005).

Hanson, Oprea, and Porter's recent research is encouraging in that it suggests that manipulation may have no discernable effect on price distortion. Although, most of this research tends to be based on a single individual or a group of independently acting individuals that have incentives to distort the market price. It's reasonable to believe in organizations were groups of individuals can successfully collude market distortions will occur. More importantly, non-manipulators may not balance the price distortion caused by the colluding groups.

Collusion tends to be difficult to maintain and typically cartels don't usually last long. There are notable exceptions like De Beers. The reason that collusion is difficult is best exemplified by the prisoner's dilemma, a classic game theory model. The basic problem in this model is that each participant would be better off if they could cooperate but without any binding agreement each will choose to not cooperate (McGuigan, Moyer, and Harris 2005). Groups in organization may be tempted to collude in order to manipulate the information market. As discussed before, it may be in a group's best interest to indirectly influence the outcome of an information market. Without credible binding agreements one group may agree to push up the price of another group's asset but after agreement the former group doesn't have much incentive to follow through. The latter group recognizes this and would likely do the same. Both groups would rather that the other push up the value of their respective asset and then spend money on other assets. In addition, depending on how the market is designed, neither group may be able to credibly discern whether the other followed through or if each group was really overstating their true valuation.

It appears from the literature that attempts to manipulate the market are not very effective, but this may be an over-generalization. In organizations where cooperation, reciprocity and reputation are enacted in the best interest of the organization adds a very real social cost to individuals or groups wishing to collude. These manipulators risk reputation and future cooperation by revealing that they wish to collude. In addition, there is also a cost to individuals or groups to augment their behavior away from the cultural norm in order to collude. This means that if they are not accustomed to colluding, trying to figure out how one goes about doing such a thing is

going to be more difficult. On the other hand, in organizations where information hoarding and non-cooperative behaviors are prevalent collusion may fail because of non-binding agreements (i.e. the prisoner's dilemma). Collusion may fail, but in these types of organizations the entire market may fail as well. In an extreme example, buyers and sellers may completely mistrust each other because both believe that a transaction would only occur if the other had an informational advantage. In such an extreme example, no one wants to buy or sell anything because they believe they will always be paying too much or selling too low and, thus, the market disappears. Akerlof illustrates this in the market for "lemons". Buyers don't know the quality of a good but know that there are high quality and low quality goods. These buyers are only willing to buy at the average price, but the seller is not willing to sell the high quality goods at the average price. Buyers recognize that sellers are unwilling to sell the high quality goods and accordingly adjust what they will pay. Both the seller and buyers know this situation is at hand, and the market either disappears or diminishes into a market of low quality goods (Akerlof 2002). It's reasonable that this principle can apply to information markets or the disappearance of such a market where individuals gain power through informational advantages over others.

Organizations where individuals and groups have successfully established relationships that are beneficial through collusion will probably succeed in distorting the information market towards their interests. If collusion is a cultural norm information markets will emphasize that behavior by allowing individuals or groups greater control in influencing decision-making. In addition, an information market may also serve to mask such behavior under the guise of a credible and reliable form of information aggregation. As mentioned before, it's not easy to determine if

individuals are really stating their true valuation of an asset. Another issue worth considering is that the manipulation is not a static behavior. This means that over time individuals or groups may learn how to exploit the market, and what may have started out as a valid form of information aggregation may depreciate into an outlet for manipulation. This conclusion is especially problematic when an organization becomes overly reliant on the information market. Previous attributed success of the information market may lead an organization to become myopic to the continual exploits occurring. This results in an organization eventually becoming disconnected from the environment it was trying to make sense of in the first place.

Another issue to be considered in information markets where manipulation still produces accurate results is the social behavior that it induces. While the market may be accurate individuals and groups will adopt behaviors in attributing less than accurate value to assets. These behaviors will more than likely be enacted exogenous to the market where distortions may not be balanced out. For instance, individuals and groups may continually undervalue ideas or goods produced outside of their group because on average they suspect that these other groups are overvaluing them. To be succinct, the organization and individuals may lose skills in appropriately evaluating knowledge.

In addition to manipulation, there are a myriad of subtle issues that can negatively affect internal information markets. Do employees have an appropriate framework for understanding what the market is? Orlikowski calls this technological framing. Individuals understand new technology by framing it in terms of existing technology. This is problematic when the new technology is quite different and there is not an appropriate frame to augment. For instance,

Orlikowski studied a consulting company that installed Lotus Notes company-wide. Lotus Notes was so new and different that many employees ended up treating it like e-mail or other familiar technologies (Orlikowski 1992). Incorporating markets into organizations is not common, and it's likely that individuals will act in unpredictable ways. Orlikowski suggest that organizations can ameliorate these issues by providing sufficient communication, information, and training about the technology (Orlikowski 1992).

Are employees interested? Individuals in the Iowa Electronic Market self-selected themselves as interested in participating (Plott and Chen 2002). This may not be the case inside organizations. Employees may feel pressured to participate but are really not that interested. Another possibility is that not enough employees are interested in participating or are just inexperienced. Lack-luster participation or very thin markets can impact the accuracy of market predictions. Organizations may be wise to establish the value of the market to employees and assess how receptive employees will be. Even when thin markets or inexperienced participants plague the internal market Chen, Fine, and Huberman propose a non-linear aggregation technique that can improve on these market inadequacies (Chen, Fine, and Huberman 2003). This is interesting in that the characteristics of the market participants may not completely determine the potential of the market. Better methods for aggregating have been or could be developed to improve results.

Is there an opportunity cost to participating in the market? Individuals might be better off doing some other task (Plott and Chen 2002). Even when individuals can earn money or prizes

does participating in the market incur costs of not attributing time to more valuable projects? If individuals are expected to accomplish projects and participate in the market they are more than likely going to pick the activity they are expected to accomplish. Executive management might consider directing managers to allot specific time for trading in order to diminish the opportunity costs individuals face.

Do individuals think that there are others participating? Having the market open all the time may be ideal, but the lack of activity may appear as if no one is participating, and individuals may lose interest after seeing little activity. In Plott's experiment the market was designed to be open during lunch and evenings in order to keep employee interest (Plott and Chen 2002). A design like this may also help allot appropriate time for employees to participate in the market. Also, it may sound counterintuitive, but attracting uniformed traders helps to establish a well-functioning market by providing profits to informed traders. Creating a "buzz" is important (Hahn and Tetlock 2006). In a similar example, Siemens tackled this issue when they launched ShareNet, an award winning knowledge management system. In order to get everyone to participate Siemens specifically assigned individuals to be knowledge managers tasked to promote, monitor, and contribute to the system (MacCormack 2002). Like Siemens, one solution for internal information markets could be to assign a small set of individuals to act as knowledge brokers. Alternatively, starting an internal information market as limited research project might be sufficient to create a "buzz" throughout the organization. The research project could be designed so that individuals self-select themselves as interested in the project. This means the participants that are involved are those that are truly interested and are the appropriate individuals to promote the system. In addition, we could imagine

that the research project offers a limited number of invitations that could be traded in the market as well. Trading invitations in the market could be useful because it may indicate how interested in the market individuals are. The organization can then credibly decide how to approach launching the internal information market to the entire organization.

Are individuals biased? Prices become distorted if individuals are not capable or understand how to trade unbiasedly. Forsythe, Nelson, Neumann and Wright noticed traders can display a bias and unrealistic hopeful belief about a preferred outcome that is "irrational" since it yields the trader less profit than unbiased traders (Berg, Forsythe, Nelson and Rietz 2000). Similarly, traders could have a tendency known as the "favorite-longshot bias" where they over-price low probability events (Wolfers and Zitzewitz 2004). Another version of bias is when traders become over-reliant on other traders' actions usually referred to as prediction bubbles. For instance, there is a monopolist who is very informed, and all other traders know this and rely on the monopolists behavior for guidance. In this case the monopolist could begin to manipulate the market (Hahn and Tetlock 2006).

Are risk attitudes interfering with the market? Risk attitudes can also alter bidding behavior that cause price distortions in the market. If an individual is extremely risk-averse then s/he will spread the bets across all possibilities. On the other hand, an extremely risk-loving individual will place all of his or her money on one possibility. The ideal situation is an individual that is risk-neutral and places bids on his or her true valuation of an asset. Again, Chen, Fine, and Huberman's non-linear aggregation method can help subtract trader risk attitudes (Chen, Fine and Huberman

2003). In addition, organizations could structure the market in order to minimize the employee's sense of risk. For instance, an organization could simply use play money.

Are the contracts in the market appropriate? There are various ways to design contracts and each way can induce different types of forecasts. Some typical contracts include linking payoffs to a specific event occurring, a continuous variable, a combination of both, and contingency between several events (Wolfers and Zitzewitz 2004). More complex contracts can be developed, and in internal markets complex contracts play a vital role since the knowledge being aggregated may not fit neatly into typical contracts. For instance, Hollywood Stock Exchange created a special contract for American Idol. This is most likely a result of the nature of the show, and the type of knowledge that people really want to trade on.

We can also imagine that while a variety of contracts have great potential this also opens the possibility for confusion and manipulation. There is an important question that arises. Who controls the market, decides what can be traded, and determines the contracts used? On one hand, a completely free market allows for the potential of valuable assets to be traded. Will this eventually cascade into a market trading assets that are not related to the business? In addition, can employees really be expected to design contracts and understand all of the details of trading in the market. On the other hand, management or some other business unit might control the market to alleviate the burden of contracting. Will the owners of the market manage it in a way that is beneficial to everyone or will they use it to shift power to themselves? For instance, the Sales business unit establishes the market. They decide that what is really relevant to the organization is sales

forecasting, and, thus, what becomes important for deciding the future of the organization is sales. Other information that could have been traded which was valuable wasn't because it was locked out by Sales. Markus indicates that systems that are designed with the intention of shifting power will be resisted by those who lose power, and changing individuals or fixing technical issues will do little to resolve the resistance (Markus 1983). This may end up leading to massive market failure. Following the previous example, we could also imagine that Sales designs the market and contracts in order to induce lower expectations of sales forecasting. Lower expectations means they wouldn't have to work as hard, and their work would be masked by the credibility of the market.

Another issue regarding power and markets is when management makes decisions that influence the outcome of the markets. For instance, if traders are tasked to predict whether a product will be launched in June, July, August, or September. When left alone the market predicts September, but let's say management realizes and acts on some opportunity that makes the product launch in July. The would-be winners predicting September are now losers and vice-versa. Management will inevitably be tempted to alter the flow of the market especially if the decision has good intentions. The real issue is that the participants are now going to incorporate the predicted actions of management into their trades, and in an extreme case the market does nothing more than predict what actions management will take (Hahn and Tetlock 2006). It's apparent that the market will cease to perform as expected.

There are undoubtedly countless other scenarios where individuals could use power and the markets in undesirable ways. Basically, these implications are really unexplored but may present

the most problematic of issues to organizations.

In conclusion, markets may be a viable alternative for facilitating knowledge in organizations, but there are costs, and, thus, there are trade-offs. Internal information markets may work well in some organizations. We might suspect, though, that in these organizations the cultural norms, behaviors, and values were already established, and a market just reinforces them. Other organizations will incur many costs in fixing the issues discussed beforehand. These costs may be so dramatic that they outweigh the value of the market. Organizations should not expect that a market alone will revolutionize the way the organization works.

## References

- Akerlof, G. (2002). The Market for 'Lemons': Quality, Uncertainty and the Market Mechanism. *Economic theory and the welfare state*. 1:308-20.
- Berg, J., Forsythe, R., Nelson, F., Rietz, T. (2000). Results from a Dozen Years of Election Futures Markets Research. Unpublished, Retrieved March 24, 2006 from [http://www.biz.uiowa.edu/iem/archive/BFNR\\_2000.pdf](http://www.biz.uiowa.edu/iem/archive/BFNR_2000.pdf).
- Chen, K., Fine, L., Huberman, B. (2003). Predicting the Future. *Information Systems Frontiers*, 5 (1): 47-61.
- Christoffel, M., Franke, T., Kotkamp, S. (2001). Trader-Supported Information Markets - A Simulation Study. *Lecture Notes in Computer Science*, 2115:101.
- Choo, C. (2006). The Knowing Organization, 2nd Edition. *Oxford University Press*.
- Davenport, T., Prusak, L. (1998). Working Knowledge: How Organizations Mangle What They Know. *Harvard Business School Press*.
- Gjerstad, S. (2005). Risk Aversion, Beliefs, and Prediction Market Equilibrium. Economic Science Laboratory, University of Arizona.
- Hahn, R., Tetlock, P. (2006). Information Markets: A New Way of Making Decisions. *The AEI Press*.
- Hanson, R., Oprea, R., Porter, D. (2006). Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior and Organization* (forthcoming).
- Hanson, R., Oprea, R. (2005). Manipulators Increase Information Market Accuracy. *mimeo*, George Mason University.
- Hayek, F. (1945). The Use of Knowledge in Society. *The American Economic Review*, 35(4): 519-530.
- Hirshleifer, J. (1973). Where are we in the Theory of Information? *The American Economic Review*, 62(2): 31-39.
- MacCormack, A. (2002). Siemens ShareNet: Building a Knowledge Network. Retrieved from

<http://harvardbusinessonline.hbsp.harvard.edu/>

Manski, C. (2005). Interpreting the Predictions of Prediction Markets. *NBER Working Paper #10359*.

Markus, M. (1983). Power, Politics, and MIS Implementation. *Communications of the ACM*, 26 (6): 430-444.

McGuigan, J., Moyer, R., Harris, F (2005). Managerial Economics, 10th Edition. *Thomson Wouth-Western*.

Orlikowski, W. (1992). Learning from Notes: Organizational Issues in Groupware Implementation. *CSCW 92 Proceedings*.

Passmore, D., Cebeci, E. (2005). Market-based information for decision support in human resource development. *Human Resource Development Review*, 4(1): 33-48.

Pennock, D., Lawrence, S., Giles, C., Nielsen, F. (2001). The Real Power of Artificial Markets. *Science*, 291:987-988.

Plott, C., Chen, K. (2002). Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. *Social Science Working Paper 1131, Pasadena, California, California Institute of Technology*.

Servan-Schreiber, E., Wolfers, J., Pennock, D., Galebach, B. (2004). Prediction Markets: Does Money Matter? *Electronic Markets*, 14(3).

Wolfers, J., Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspectives*, 18(2): 107-126.