

Outcome Manipulation in Corporate Prediction Markets*

Marco Ottaviani[†] Peter Norman Sørensen[‡]

February 2007

Abstract

This paper presents a framework for applying prediction markets to corporate decision making. The analysis is motivated by the recent surge of interest in markets as information aggregation devices and their potential use within firms. We characterize the amount of outcome manipulation that results in equilibrium and the impact of this manipulation on market prices. (JEL: D71, D82, D83, D84)

Keywords: Prediction markets, information aggregation, forecasting, incentives.

*This paper has been prepared for the invited session on “Communication and Complexity in Organizations” at the 2006 Conference of the European Economic Association in Vienna. We thank Itay Goldstein, Jesper Krogstrup, and Stefano Sacchetto for helpful comments.

[†]Economics Subject Area, London Business School, Sussex Place, Regent’s Park, London NW1 4SA, England. Phone: +44-20-7000-8411. Fax: +44-20-7000-8401. E-mail: mottaviani@london.edu. Web: <http://faculty.london.edu/mottaviani>.

[‡]Department of Economics, University of Copenhagen, Studiestræde 6, DK-1455 Copenhagen K, Denmark. Phone: +45-3532-3056. Fax: +45-3532-3000. E-mail: peter.sorensen@econ.ku.dk. Web: <http://www.econ.ku.dk/sorensen>.

1 Introduction

Markets have long been praised for their ability to pool the many bits of information that are otherwise dispersed across agents in the economy.¹ More recently, the wisdom of markets as prediction tools has been popularized by best-selling author Surowiecki (2004) and is currently attracting attention in prominent business and policy circles.²

Prediction markets are incentive-based mechanisms designed to pool information about future events. Arguably, prediction market forecasts are often more accurate and less expensive than those obtained with other more traditional forecasting methods, such as expert opinions or pools (see e.g., Berg and Rietz 2006). It is then natural to take these markets inside the firm and use them to produce information relevant for corporate decision making. Championed by companies such as Hewlett-Packard, Yahoo, and Google, prediction markets are now being used by a small but growing number of organizations.³ These developments have the potential to change drastically the way information is processed in organizations.

In prediction markets, assets are created whose final cash value is tied to a particular event (e.g., will a new factory be opened by the end of the quarter?) or a parameter (e.g., how many units of a product will be sold during the next quarter?). Prediction markets are based on the idea that the equilibrium price should reflect the information possessed by the market participants. Essentially, prediction markets are particularly simple financial markets that are created with the purpose of collecting information, but serve no liquidity purposes.

Despite the hype in the press, there is a limited amount of theoretical analysis in this area. In this paper, we present a modeling framework that can guide practitioners and researchers to understand the role and improve the design of corporate prediction markets. As Wolfers and Zitzewitz

¹See e.g. Hayek (1945) and Grossman (1976).

²For example, see the collection of essays in Hahn and Tetlock (2006a) and King's (2006) article in *Business Week*.

³See Plott and Chen's (2002) report on Hewlett-Packard's pioneering experiment.

(2006) stress, prediction markets must overcome a number of challenges in order to be used as effective prediction tools. In a corporate setting, market designers are often concerned about the possibility of outcome manipulation. Being also member of the organization, market participants are often in a position to take actions that directly influence the chance of project success.

We consider a simple prediction market set up to gather information about the chances of a binary event, such as the timely completion of a corporate project. In our model, more optimistic traders bet on success while pessimists bet against success. As a side effect of the prediction market, optimists have an incentive to improve the chances of project success, while pessimists manipulate the chances in the opposite direction.

We show that this game among traders leads them to wasteful activities partly off-setting each others' manipulations. Moreover, in the model we find that the manipulations of optimists and pessimists typically do not perfectly off-set each other. Hence, introducing the prediction market helps the corporate executive to become more informed about the chance of project success, but it also directly influences this chance.

To the best of our knowledge, ours is the first formal analysis of outcome manipulation in prediction markets. See Hanson (2006) and Wolfers and Zitzewitz (2006) for informal discussions of outcome manipulation. The different problem of price manipulation—according to which traders have an interest in affecting the market price because they want to affect decisions based on that price—has been analyzed theoretically by Hanson and Oprea (2004) and Hahn and Tetlock (2006b) and empirically by Rhode and Strumpf (2006).⁴

⁴See also Goldstein and Guembel's (2006) interesting analysis of manipulation when stock prices affect the allocation of resources. By selling short, a manipulator can reduce the intrinsic value of the underlying asset and then profit by buying later at a lower price. In their model, manipulation is "trade based", but triggers indirectly a change in the underlying outcome which is determined by the firm's investment decision (which, in turn, is based on the information contained in the stock price). In our model instead, traders affect *directly* the outcome.

While couched in the context of prediction markets, our model contributes also to the wider literature on manipulation in financial markets. Outcome manipulation is a version of “action-based manipulation” (i.e., manipulation based on actions that change the actual value of the assets), in the nomenclature proposed by Allen and Gale (1992). See Vila (1989) for an alternative model of action-based manipulation, in which market participants cannot observe whether a manipulator is present. In our model instead, all traders are allowed to manipulate the outcome and their presence is commonly known.

Our model is based on Ottaviani and Sørensen’s (2005) analysis of rational expectations equilibrium when risk-averse traders have heterogeneous priors and private information. While Ottaviani and Sørensen (2005) focus on the bias that is generated by the presence of wealth effects, in this paper we abstract from wealth effects by considering the special case with constant absolute risk aversion preferences.

The paper proceeds as follows. In Section 2 we present the model. In Section 3 we analyze the model, focusing on the issue of outcome manipulation. We conclude in Section 4.

2 Modeling Framework

We consider a prediction market written on the realization of a binary event. To fix ideas, suppose that E is the event that a corporate project is completed by a given deadline. Following the rules of the Iowa Electronic Markets (see Berg and Rietz 2006), there are two Arrow-Debreu assets. Asset 1 pays out 1 currency unit if event E is realized, and 0 otherwise. Asset 2 pays out 1 currency unit if the complementary event E^c results, and 0 otherwise. Traders can enter the asset market by first buying an equal number of both assets, the bundle priced at 1. Subsequently, they can trade the assets with other traders. Hence, markets clear when the aggregate demand for asset 1 precisely equals the aggregate demand for asset 2. By arbitrage, the sum of the prices of the two assets is equal to

one because any trader can obtain back the dollar invested by purchasing both assets. Hence, we focus on the price of the asset paying in event E and denote this price by p .

A key feature of prediction markets is that the traders do not have a liquidity motive to trade. Hence, we assume that the initial wealth endowment w_{i0} of trader i is constant with respect to events E and E^c . Trader i has subjective prior belief q_i . In addition, before trading, trader i privately observes signal s_i . Although traders have heterogeneous priors, we assume that they interpret information in the same way—in the language of Milgrom and Stokey (1982), traders have concordant beliefs.

Trader i revises the subjective prior q_i to form the posterior belief π_i by incorporating the information contained in private signal s_i as well as in the market price. Trader i 's preferences are represented by the subjective expected utility

$$\pi_i u_i(w_i(E)) + (1 - \pi_i) u_i(w_i(E^c)).$$

We confine attention to the analytically convenient case of Constant Absolute Risk Aversion (CARA) utility, $u_i(w) = -\exp(-A_i w)$, where A_i is the Arrow-Pratt coefficient of absolute risk aversion.

Traders are members of the corporation, and we assume that they have an opportunity to directly affect the chance that the corporate project succeeds. Specifically, trader i chooses a degree of manipulation $m_i \in \mathbb{R}$. Given signals $s = (s_1, \dots, s_I)$ and manipulations $m = (m_1, \dots, m_I)$, trader i 's subjective posterior chance $\pi_i(s, m)$ that event E is true satisfies

$$\log \frac{\pi_i(s, m)}{1 - \pi_i(s, m)} = \log \frac{q_i}{1 - q_i} + \sum_{i=1}^I \log \frac{f(s_i|E)}{f(s_i|E^c)} + \sum_{i=1}^I m_i. \quad (1)$$

Thus, the manipulation m_i describes how far trader i moves the log-likelihood ratios of each trader's posterior beliefs.⁵ We assume that manipulation is costly, so that $c(m_i)$ is subtracted from the individual's wealth,

⁵This is one of many possible formulations of outcome manipulation. We wish to give traders the chance to add and subtract from the probability of outcome success. Given that individuals have heterogeneous prior beliefs, our simple and manageable formulation allows the trader to simultaneously move all individuals' posterior beliefs equally far.

where c is a differentiable and strictly convex function with $c(-m) = c(m)$ for all m .

We study the market in a fully revealing rational expectations equilibrium. The price p then reveals the log-likelihood ratio for updating,

$$\ell = \sum_{i=1}^I \log \frac{f(s_i|E)}{f(s_i|E^c)} + \sum_{i=1}^I m_i. \quad (2)$$

Taking as given the price p and the residual information $\ell_{-i} := \ell - m_i$, trader i then chooses the asset position as well as the manipulation m_i to maximize his posterior expected utility. The chosen portfolios clear the market.

Before turning to the characterization of the equilibrium, we stress our assumption that traders behave perfectly competitively when taking as given the price, but yet have direct influence on the beliefs of others through the manipulation. By making this perhaps unrealistic assumption, we stress our focus on outcome manipulation rather than price manipulation. In addition, this assumption makes the model easy to analyze.

3 Outcome Manipulation

3.1 Individual Manipulation

The following result states the intuitive fact that a trader who adopts a non-neutral position in the asset market will choose to manipulate the outcome in the direction that favors the asset position's return. When x_{ik} denotes trader i 's final holding asset k , the net position is defined as $\Delta x_i = x_{i1} - x_{i2}$.

Proposition 1 *Trader i chooses a long net position, $\Delta x_i > 0$, if and only if the trader's manipulation satisfies $m_i > 0$.*

Proof. The portfolio (x_{i1}, x_{i2}) and manipulation m_i give final wealth in the two events

$$w_i(E) = w_{i0} - c(m_i) + x_{i1} - px_{i1} - (1-p)x_{i2} = w_{i0} - c(m_i) + (1-p)\Delta x_i,$$

$$w_i(E^c) = w_{i0} - c(m_i) + x_{i2} - px_{i1} - (1-p)x_{i2} = w_{i0} - c(m_i) - p\Delta x_i.$$

Given p and ℓ_{-i} , the trader chooses Δx_i and m_i to maximize posterior expected utility $\pi_i u_i(w_i(E)) + (1 - \pi_i) u_i(w_i(E^c))$ where, from (1),

$$\log \frac{\pi_i}{1 - \pi_i} = \log \frac{q_i}{1 - q_i} + \ell_{-i} + m_i. \quad (3)$$

The necessary first-order condition for Δx_i is

$$\pi_i (1 - p) u'_i(w_i(E)) = (1 - \pi_i) p u'_i(w_i(E^c)). \quad (4)$$

The necessary first-order condition for m_i is

$$\begin{aligned} \frac{d\pi_i}{dm_i} u_i(w_i(E)) - \pi_i u'_i(w_i(E)) c'(m_i) \\ = \frac{d\pi_i}{dm_i} u_i(w_i(E^c)) + (1 - \pi_i) u'_i(w_i(E^c)) c'(m_i). \end{aligned} \quad (5)$$

Using $dm_i = d\pi_i / ((1 - \pi_i) \pi_i)$ and (4), this can be rewritten as

$$c'(m_i) = \frac{(1 - p) \pi_i (u_i(w_i(E)) - u_i(w_i(E^c)))}{u'_i(w_i(E^c))}. \quad (6)$$

Thus, trader i chooses a long net position, $\Delta x_i > 0$, if and only if $w_i(E) > w_i(E^c)$, if and only if $c'(m_i) > 0$, if and only if $m_i > 0$. ■

3.2 Equilibrium Characterization

The CARA assumption is convenient, because it allows for a simple characterization of the equilibrium price and trades. The relative risk tolerance of trader i is denoted by $\tau_i = (1/A_i) / (\sum_{j=1}^I (1/A_j))$.

Proposition 2 *Define the market prior q by*

$$\log \frac{q}{1 - q} = \sum_{i=1}^I \tau_i \log \frac{q_i}{1 - q_i}. \quad (7)$$

The market price p is the posterior probability of event E given this prior and incorporating the log-likelihood ratio ℓ that summarizes all the private signals and outcome manipulations. Trader i 's equilibrium position is

$$\Delta x_i = \frac{1}{A_i} \left(\log \frac{q_i}{1 - q_i} - \log \frac{q}{1 - q} \right). \quad (8)$$

Proof. With the CARA utility function, (4) is solved by

$$\Delta x_i = \frac{1}{A_i} \log \left(\frac{1-p}{p} \frac{\pi_i}{1-\pi_i} \right). \quad (9)$$

Since the two assets are supplied in equal amount, the price p clears the market when

$$\sum_{i=1}^I \Delta x_i = 0. \quad (10)$$

Combining (9) and (10), we obtain that the equilibrium price obeys

$$\log \frac{p}{1-p} = \sum_{i=1}^I \tau_i \log \frac{\pi_i}{1-\pi_i} = \sum_{i=1}^I \tau_i \log \frac{q_i}{1-q_i} + \ell, \quad (11)$$

where the log-likelihood (2) incorporates all the relevant private information as well as the aggregate outcome manipulation. Note that the price is an increasing function of ℓ and hence fully revealing. Finally, inserting (11) into (9), we obtain (8). ■

By (8), a trader's equilibrium net position is a function of the prior disagreement with the average belief. Notice that the outcome manipulations m do not enter here—this is the essential advantage of the CARA assumption. The result that the equilibrium price is equal to the market posterior belief (given the market prior (7) and incorporating all private signals and manipulations) is special to the CARA setting, as shown by Ottaviani and Sørensen (2005).

3.3 Aggregate Manipulation

As we have already observed, optimistic traders with $q_i > q$ choose $m_i > 0$. On the other hand, pessimists choose a negative amount of manipulation. In principle, these manipulations could cancel out. We provide a simple two-trader illustration to show that, as a rule, there is non-negligible aggregate manipulation in equilibrium.

Proposition 3 *Suppose that there are two traders with equal constant absolute risk aversion, $A_1 = A_2$, and different subjective prior beliefs. Let $\ell^* = \log(1-q) - \log q$ where q is the market prior given by (8). Then the aggregate manipulation $m_1 + m_2$ is positive when $\ell < \ell^*$ and negative when $\ell > \ell^*$.*

Proof. Using the CARA utility function and expression (8) for Δx_i , (6) can be simplified to

$$c'(m_i) = \frac{(1-p)\pi_i(q_i(1-q) - (1-q_i)q)}{A_i(1-q)q_i}. \quad (12)$$

Without loss of generality, suppose that $q_1 > q_2$. As we have observed, then $m_1 > 0 > m_2$. Using convexity and symmetry of the cost function, notice that $m_1 + m_2 > 0$ if and only if $c'(m_1) + c'(m_2) > 0$. Since $[(1-\pi_i)p]/[\pi_i(1-p)] = [(1-q_i)q]/[q_i(1-q)]$, we can rewrite (12) as

$$c'(m_i) = \frac{\pi_i(1-p) - (1-\pi_i)p}{A_i}. \quad (13)$$

Thus, $c'(m_1) + c'(m_2) > 0$ if and only if $(\pi_1 + \pi_2)/(1 - \pi_1 + 1 - \pi_2) > p/(1-p)$. From (11) we have $p^2/(1-p)^2 = [\pi_1/(1-\pi_1)][\pi_2/(1-\pi_2)]$. Using this fact, we obtain that $(\pi_1 + \pi_2)/(1 - \pi_1 + 1 - \pi_2) > p/(1-p)$ if and only if

$$[\pi_1(1-\pi_1) - \pi_2(1-\pi_2)][\pi_1 - \pi_2] > 0.$$

Since $q_1 > q_2$, we always have $\pi_1 > \pi_2$, so the second factor is positive. The sign of the first factor is positive when π_1 is closer than π_2 to $1/2$. Precisely when $\ell < \ell^*$, we have $p < 1/2$, and hence $\pi_1 > p$ closer than $\pi_2 < p$ to $1/2$. In sum, we have proved that $m_1 + m_2 > 0$ if and only if $\ell < \ell^*$. The same line of arguments shows that $m_1 + m_2 < 0$ if and only if $\ell > \ell^*$. ■

Corollary 4 *Consider two traders with equal risk aversion and different priors. When the signal-based information $\log(f(s_1|E)/f(s_1|E^c)) + \log(f(s_2|E)/f(s_2|E^c))$ is below ℓ^* , then $m_1 + m_2 > 0$ and $\ell < \ell^*$.*

Proof. First, if $m_1 + m_2 \leq 0$ then $\ell < \ell^*$, but that would contradict the proposition. Second, if $m_1 + m_2 > 0$ is so large that $\ell \geq \ell^*$, there is again a contradiction to the proposition. ■

When the realized signal-based information $\log(f(s_1|E)/f(s_1|E^c)) + \log(f(s_2|E)/f(s_2|E^c))$ is precisely ℓ^* , the situation is perfectly symmetric. The price is $1/2$, and the traders' posterior beliefs are equally far from $1/2$. Both traders attempt to manipulate the outcome in their desired

direction, but in the symmetric situation they choose equal magnitudes of manipulations, and hence $m_1 + m_2 = 0$.

When instead the signal-based information is below ℓ^* , traders' subjective posterior beliefs are lower, and the optimist's belief must be closer than the pessimist's to $1/2$. Manipulations move the log-likelihood ratio of the subjective probability by a given distance, but the trader with beliefs closer to $1/2$ sees a faster change in the outcome's probability. This gives a greater incentive to manipulate, so the optimist chooses a greater manipulation than the pessimist, and hence $m_1 + m_2 > 0$. Still, the corollary establishes that the positive amount of manipulation is not so strong as to completely overwhelm the signal-based information, for in equilibrium $\ell < \ell^*$.

3.4 Discussion

We have found that the agents choose to manipulate the chance of the project's success, when they are given the opportunity to do so. While each trader finds it in his private interest to manipulate the market, the manipulations of traders on the other side of the market counteract some of this effort. Thus, part of the traders' costly manipulations are socially wasted.

Outcome manipulation may also interfere with the intentions of the corporate decision maker, who cannot merely regard the prediction market as a tool for collecting information. On the other hand, the market does convey information to the decision maker, and it does permit traders to gain from trades on belief differences. We leave to future research a complete analysis of the welfare implications of opening the market.

In our setting with CARA preferences, the price resulting in the prediction market can be usefully interpreted as a posterior chance of the event given all information and manipulations. The results of Ottaviani and Sørensen (2005) suggest that this neutrality result is a peculiar property of the CARA model. More generally, we expect the extent of manipulation to interact with information and create a bias in the interpretation of the

price as a probability.⁶

4 Conclusion

While economists have long praised the ability of markets to aggregate dispersed information, they have also identified important barriers to the aggregation of information—such as the no-trade theorem and the occurrence of herd behavior. This paper identifies outcome manipulation as another potential obstacle to the use of markets within corporations.

We believe that prediction markets provide economists with a fascinating laboratory for testing and finessing theories of individual and collective decision making. In turn, economic analysis can offer insights for the practical development of prediction markets. As we have argued in this paper, the operation of prediction markets can interfere with more traditional ways of doing business. We leave to future work the empirical quantification of outcome manipulation in these markets and the formulation of solutions.

⁶See also Manski (2006), Wolfers and Zitzewitz (2005), Gjerstad (2005) on the interpretation of prediction market prices as probabilities.

References

- Allen, Franklin, and Douglas Gale (1992). "Stock-Price Manipulation." *Review of Financial Studies*, 5(3), 503–529.
- Berg, Joyce E., and Thomas A. Rietz (2006). "The Iowa Electronic Markets: Stylized Facts and Open Issues." In *Information Markets: A New Way of Making Decisions*, edited by Robert W. Hahn and Paul C. Tetlock. AEI-Brookings Joint Center for Regulatory Studies.
- Gjerstad, Steven (2005). "Risk Aversion, Beliefs, and Prediction Market Equilibrium." Working Paper, University of Arizona.
- Goldstein, Itay, and Alexander Guembel (2006). "Manipulation and the Allocation Role of Prices." Working paper, Wharton School and Said Business School.
- Grossman, Sanford J. (1976) "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information." *Journal of Finance*, 31(2), 573–585.
- Hahn, Robert W., and Paul C. Tetlock (2006a). *Information Markets: A New Way of Making Decisions*. AEI-Brookings Joint Center for Regulatory Studies.
- Hahn, Robert W., and Paul C. Tetlock (2006b). "Optimal Liquidity Provision for Decision Makers." Working Paper, University of Texas at Austin.
- Hanson, Robin (2006). "Foul Play in Information Markets." In *Information Markets: A New Way of Making Decisions*, edited by Robert W. Hahn and Paul C. Tetlock. AEI-Brookings Joint Center for Regulatory Studies.
- Hanson, Robin and Ryan Oprea (2004). "Manipulators Increase Information Market Accuracy." Working Paper, George Mason University.
- Hayek, Friedrich A. (1945). "The Use of Knowledge in Society." *American Economic Review*, 35(4), 519–530.
- King, Rachael (2006). "Workers: Place Your Bets." *Business Week online*, August 3.

Manski, Charles F. (2006). “Interpreting the Predictions of Prediction Markets.” *Economics Letters*, 91(3), 425–429.

Milgrom, Paul, and Nancy Stokey (1982). “Information, Trade and Common Knowledge.” *Journal of Economic Theory*, 26(1), 17–27.

Ottaviani, Marco, and Peter N. Sørensen (2005). “Aggregation of Information and Beliefs in Prediction Markets.” Working paper, London Business School.

Plott, Charles, and Kay-Yut Chen (2002). “Information Aggregation Mechanisms: Concepts, Design and Implementation for a Sales Forecasting Problem.” Social Science Working Paper 1131, California Institute of Technology.

Rhode, Paul W., and Koleman S. Strumpf (2006). “Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data.” University of North Carolina at Chapel Hill Working Paper.

Surowiecki, James (2004). *The Wisdom of Crowds*. Doubleday.

Vila, Jean-Luc (1989). “Simple Games of Market Manipulation.” *Economics Letters*, 29(1), 21–26.

Wolfers, Justin, and Eric Zitzewitz (2005). “Interpreting Prediction Market Prices as Probabilities.” Working Paper, Wharton School and Stanford GSB.

Wolfers, Justin, and Eric Zitzewitz (2006). “Five Open Questions About Prediction Markets.” In *Information Markets: A New Way of Making Decisions*, edited by Robert W. Hahn and Paul C. Tetlock. AEI-Brookings Joint Center for Regulatory Studies.