

A Manipulator Can Aid Prediction Market Accuracy

By ROBIN HANSON† and RYAN OPREA‡

†George Mason University and

‡University of California, Santa Cruz

Final version received 15 May 2008.

Prediction markets are low volume speculative markets whose prices offer informative forecasts on particular policy topics. Observers worry that traders may attempt to mislead decision makers by manipulating prices. We adapt a Kyle-style market microstructure model to this case, adding a manipulator with an additional quadratic preference regarding the price. In this model, when other traders are uncertain about the manipulator's target price, the mean target price has no effect on prices, and raising the variance of the target price can *increase* average price accuracy, by boosting the returns to informed trading and thereby incentives for traders to become informed.

INTRODUCTION

Observers have long been impressed by the ability of speculative markets to aggregate available information; it is hard to find information that is not already embodied in the prices of thick markets (Lo 1997). Recently new markets have emerged, known as 'prediction markets', 'information markets', or 'idea futures', which attempt to harness this power to make forecasts on particular topics of interest (Berg and Rietz 2003; Chen and Plott 2002; Spann and Skiera 2003; Wolfers and Zitzewitz 2004). Such markets promise to help firms, governments and researchers to make better informed decisions (Hanson 1999).

In 2003, prediction markets generated some controversy and caused some concern. During the furor over the DARPA Policy Analysis Markets (PAM), otherwise known as the 'terrorism futures' market, critics worried that PAM would encourage betting on the details of individual terrorist attacks.¹ Critics particularly feared that terrorists might commit more acts of terror in order to win bets about those acts, or that terrorists might intentionally lose bets in order to deceive decision makers.²

Since existing thick financial markets already respond to major terrorist attacks, it seems unlikely that new prediction markets, which are typically very thin, would substantially affect incentives in extra-market behaviour. A terrorist who wanted to profit from his activities would have done far better in a natural stock market. The typical thinness of prediction markets, however, does seem to make it cheap for participants to mislead markets with their trades. And at least one apparently successful attempt to manipulate political prediction markets has been reported (Hansen *et al.* 2004).

Many others, however, have reported failed attempts to manipulate prices with trades, historically (Strumpf and Rhode 2004), in the field (Camerer 1998) and in the laboratory (Hanson *et al.* 2006; Oprea *et al.* 2007). A recent review article concludes that, 'none of these attempts at manipulation had much of a discernible effect on prices, except during a short transition phase' (Wolfers and Zitzewitz 2004). Why is not this sort of manipulation³ as prevalent as many fear?

One explanation is that a manipulative trader is in essence a 'noise' trader, in the sense that his trades are based on considerations other than his best estimate of asset value. Standard models of market microstructure find that when potentially informed

traders have deep pockets relative to the volume of noise trading, increases in trading noise do not directly affect price accuracy. In addition, by inducing more traders to become better informed, an increase in noise trading indirectly *improves* the accuracy of market prices (Kyle 1989; Spiegel and Subrahmanyam 1992). If the presence of manipulative traders similarly induced more effort by informed traders, this could help explain the typical failure of manipulation attempts.

This paper presents a formal model which, in certain standard settings, illustrates this view of manipulative traders as noise traders. We start with one of the simplest standard models of market microstructure, the single-period simplification of Kyle (1985, 1989) as it appears in popular textbooks on market microstructure (Brunnermeier 2001; O'Hara 1997). This model contains noise traders, potentially informed traders and a competitive market maker. We adapt this standard model to the case of a thin prediction market, and then add a manipulator, i.e. a trader who has an additional preference over the market price, or equivalently over the beliefs of neutral observers influenced by the price.

We are interested in the case of a thin prediction market, i.e. where the quantity traded is small, where relevant information can be obtained with some effort, and where there may be no liquidity traders to attract rational speculators. We thus consider the case of risk-neutral traders who can buy information, and who may be, but need not be, mildly irrational. That is, each trader's information depends on his or her choice of information-gathering effort, and we allow trading choices to be made using a standard general model of irrational behaviour, the quantal response noisy game theory equilibrium (Goeree and Holt 2001; McKelvey and Palfrey 1995).

Into this basic prediction market model we add a manipulator who, in addition to a usual profit-based payoff, has a quadratic preference for the market price to be close to a target price. (This generalizes a previous model which had a linear manipulator preference (Kumar and Seppi 1992).) Other traders know the strength of this quadratic preference, but have only a noisy clue about the manipulator's target price. To study a best case for the possibility of manipulation, we consider a single fully-rational manipulator.

We find that the manipulator's mean target price has no effect on the market price, and that variance in the manipulator's target price has no direct effect on average price accuracy. However, by increasing the expected rewards to informed trading, a larger manipulator variance motivates other traders to gather information and so indirectly *increases* the accuracy of the market price as an estimate of fundamental asset value. Thus in this standard market microstructure model of thin prediction markets, with rational or irrational traders who can obtain information with effort, manipulator bias that is within the range of biases that traders suspect might exist will on average improve price accuracy.⁴

Our model illustrates potentially positive effects of manipulation attempts on price formation, effects which have been ignored previously in the prediction market literature. However, since this is not a fully general model, it cannot by itself support strong general claims about the price effects of manipulation. For example, the model assumes risk-neutrality, normally distributed values and signal errors, interior choices of information quantity, no transaction costs of trading, no budget constraints, and a single rational manipulator with quadratic manipulation preferences and a commonly known strength of desire to manipulate. While we believe that these assumptions are natural ones for a first modeling effort, one should remember that some of the model's findings may not be robust to changes in some of these assumptions.

I. MODEL

Consider a one-period version of Kyle (1985, 1989) that appears in popular textbooks on market microstructure (Brunnermeier 2001; O'Hara 1997). In this model, there is a single asset of uncertain value, and a single risk-neutral informed agent who has a noisy clue about this asset value. This informed agent chooses a quantity of the asset to buy or sell. A 'noise' or 'liquidity' trader asks to trade a random quantity at the same time, and these two orders are combined into a net market order, which is observed by a competitive risk-neutral market maker. This market maker then sets the price for these trades to be the expected value of the asset conditional on seeing this net market order.

Let us modify this standard model to describe a thin prediction market with a manipulative trader. That is, consider a single asset whose true value is drawn as $v \sim N(\bar{v}, S_v)$, i.e. from a normal distribution with mean \bar{v} and (finite) variance $S_v > 0$. We will allow for the possibility of a liquidity trade quantity l drawn as $l \sim N(\bar{l}, S_l)$, but all our results will hold when $S_l = 0$. (Unless specified otherwise, all parameters are drawn independently.)

There are T risk-neutral traders, labeled $i \in \{1, 2, \dots, T\}$, each gaining a trading profit

$$\pi_i(x_i) = x_i(v - P)$$

if he buys a quantity x_i when the asset price is P .

We also add a special trader, $i = 0$, with an extra quadratic preference regarding the market price P .⁵ His profit can be written in two equivalent forms, either in terms of a target price t or in terms of a bias w , as

$$(1) \quad \begin{aligned} \pi_0(x_0) &= x_0(v - P) - k(t - P)^2 \\ &= x_0(v - P) - k(\bar{v} - P)^2 + wP - \hat{k}, \end{aligned}$$

assuming $k \geq 0$. One can translate between these forms using $w = 2k(t - \bar{v})$ and $\hat{k} = k(t^2 - \bar{v}^2)$. Since an additional preference term gives this trader a special reason to manipulate the price, we call this trader⁶ the 'manipulator.'

We will use the second form, shown in equation (1), because this explicitly includes the case of a linear manipulation preference, where $k = 0$ and $w \neq 0$. Such a linear preference has been considered previously, in a model where a trader seeks to manipulate a spot market in order to raise the value of assets previously acquired in a futures market (Kumar and Seppi 1992). We assume that k is common knowledge, but that 'bias' w (or equivalently the target price t) is private information to the manipulator, and is only commonly known to have been drawn as $w \sim N(\bar{w}, S_w)$.

Of the T traders, let us assume that N of them can acquire information about the true asset value v , and about the manipulator's bias w . Specifically, each informed trader i observes the clues $a_i = v + \varepsilon_i$ and $b_i = w + \delta_i$, where $\varepsilon_i \sim N(\bar{\varepsilon}, S_\varepsilon)$ and $\delta_i \sim N(\bar{\delta}, S_\delta)$. If we define relative clue accuracy as $\eta_i \equiv S_v / (S_v + S_\varepsilon)$ and $\rho \equiv S_w / (S_w + S_\delta)$, then these clues give traders improved estimates $E[\Delta v | a_i] = \eta_i \Delta a_i$ and $E[\Delta w | b_i] = \rho \Delta b_i$, where for any variable f the deviation from its unconditional mean is $\Delta f \equiv f - E[f]$.

We assume that each informed trader can reduce the variance S_i of his asset value clue via costly effort $c(S_i)$, where $c'(S_i) < 0$, $c''(S_i) > 0$, and $c(\infty) = 0$. Full trader profit is thus

$$\pi_i(x_i, S_i) = x_i(v - P) - c(S_i).$$

The N informed traders can choose $S_i \in [0, \infty)$, and get $S_\delta < \infty$, while the other uninformed traders are stuck with $S_i = S_\delta = \infty$. For simplicity, we also assume $S_0 = \infty$,

so that the manipulator has no private information on the asset value, though he does know w exactly.

The order of actions is as follows. First, every informed trader chooses his clue noise S_i . Second, every trader observes his clues a_i, b_i . Third, every trader chooses his market order x_i . No trader sees any other player choices before choosing S_i or x_i , and so informed trader effort $c(S_i)$ is hidden. Fourth, orders x_i are summed and added to the liquidity trade l to produce a total order

$$y = l + \sum_{i=0}^T x_i.$$

Finally, the market maker observes this total order y and sets the market price to be

$$P = E[v|y] + \theta,$$

where $\theta \sim N(0, S_\theta)$ describes error in the (unmodeled) price-setting process. All our results will continue to hold when $S_\theta = 0$.

The optimal actions, expressed in terms of expected profit, are thus

$$S_i^* = \operatorname{argmax}_{S_i \in R^+} [\bar{\pi}_i(S_i) \equiv E[\pi_i(x_i^*(S_i), S_i)|S_i]],$$

$$x_i^*(S_i) = \operatorname{argmax}_{x_i \in R} [\bar{\pi}_i(x_i) \equiv E[\pi_i(x_i, S_i)|a_i, b_i, x_i]],$$

$$x_0^* = \operatorname{argmax}_{x_0 \in R} [\bar{\pi}_0(x_0) \equiv E[\pi_0(x_0)|w, x_0]].$$

While all our results will hold when all traders choose optimal quantities x_i^* , we also allow irrational choice given by the widely used noisy game theory, the quantal response equilibrium (Goeree and Holt 2001; McKelvey and Palfrey 1995). This gives a probability density of

$$(2) \quad \Pr[x_i] = \frac{\exp(\bar{\pi}_i(x_i)/r_i)}{\int_{-\infty}^{\infty} \exp(\bar{\pi}_i(x'_i)/r_i) dx'_i},$$

where $r_i \geq 0$ is the irrationality of trader i . When $r_i = 0$, then $x_i = x_i^*$ for sure. To study a best case for manipulation, we assume $r_0 = 0$.

Ultimately, we want to know whether a stronger manipulator preference reduces or increases price accuracy. That is, how does the distribution of the error of the price as an estimate of value, $P - v$, including its mean $E[P - v]$ and mean square

$$\Phi = E[(P - v)^2],$$

respond to changes in the parameters \bar{w} and S_w describing the distribution of manipulator bias? (When we change the variance S_w we will assume that the clue error S_δ varies proportionally, so that the ratio S_w/S_δ remains constant.)

II. ANALYSIS

As is standard in such a model, we seek equilibria where the price and quantities are linear in the various clues and parameters. Such linear equilibria are possible because we have assumed jointly normal distributions, which gives linear conditional expectations, such as

$$(3) \quad E[v|y] = \bar{v} + \Delta y \frac{E[\Delta v \Delta y]}{E[\Delta y \Delta y]}.$$

For clarity, we focus on deviations from means. We thus seek equilibria where

$$(4) \quad P = \mu + \lambda \Delta y + \theta,$$

$$(5) \quad x_i^* = \alpha_i + \beta_i \Delta a_i + \gamma_i \Delta b_i,$$

$$(6) \quad x_0^* = \tau + \gamma_0 \Delta w$$

for certain values $\mu, \lambda, \alpha_i, \beta_i, \gamma_i, \tau, \gamma_0$. We furthermore seek equilibria which are symmetric within each class of traders. That is, for all informed traders $S_i = S_e, \eta_i = \eta, \alpha_i = \alpha, \beta_i = \beta$, and $\gamma_i = \gamma$, and for all other traders (besides the manipulator) $\alpha_i = \beta_i = \gamma_i = 0$.

Linear strategies make profit functions $\bar{\pi}_i(x_i)$ quadratic, and so quantal response behaviour of equation (2) makes trading errors $e_i = x_i - x_i^*$ normally distributed. Thus $e_i \sim N(0, S_{e_i})$, with a variance of $S_{e_i} = -r_i/\bar{\pi}_i''$, expressed using the constant second derivative of $\bar{\pi}_i(x_i)$. We can combine these trading errors to obtain a total noise trade of $e \equiv \sum_{i=0}^T e_i$ with variance $S_E \equiv \sum_{i=0}^T S_{e_i}$. This allows us to redescribe the total order as

$$y = l + e + \sum_{i=0}^T x_i^*,$$

a form into which we can substitute our linear equations (4), (5) and (6).

Making such substitutions in our price error form gives

$$(7) \quad \Phi = S_\theta + S_v(1 - N\lambda\beta),$$

and in the market maker equations (3) and (4) gives $\mu = \bar{v}$, implying $E[P - v] = 0$, and

$$(8) \quad \frac{1}{\lambda} = \frac{S_E + S_X}{\beta N S_v} + \frac{\beta(N S_v + S_e)}{S_v},$$

where non-error trading noise is

$$S_X \equiv S_l + N\gamma^2 S_\delta + (\gamma_0 + \gamma N)^2 S_w.$$

Substituting into the expected trader profit forms gives

$$\begin{aligned} \bar{\pi}_0(x_0) &= w\bar{v} - \lambda(x_0 - w)(\Delta x_0 + N\gamma\Delta w) - \hat{k} - kS_\theta - k\lambda^2(S_l + S_E) \\ &\quad - k\lambda^2(\Delta x_0^2 + 2\gamma N\Delta w\Delta x_0 + \beta^2 N(NS_v + S_e) + \gamma^2 N(N\Delta w^2 + S_\delta)) \\ \bar{\pi}_i(x_i) &= x_i[(1 - \lambda\beta(N - 1))\eta_i\Delta a_i - \lambda(\Delta x_i + (\gamma_0 + \gamma(N - 1))\rho\Delta b_i)] - c(S_i) \\ \bar{\pi}_i(S_i) &= \beta_i(1 - \lambda(\beta + (N - 1)))S_v - \lambda(\beta_i^2(S_v + S_i) + \gamma^2 S_\delta + \gamma(\gamma_0 + N\gamma)S_w) - c(S_i). \end{aligned}$$

Note that when choosing S_i , agent i must distinguish the equilibrium value β from the value β_i that will describe his behaviour if he makes an out of equilibrium choice of $S_i \neq S_e$.

For the uninformed traders $\bar{\pi}_i(x_i) = -\lambda x_i \Delta x_i$, making the first-order condition (FOC) on x_i be $x_i^* = 0$, which gives $\alpha_i = \beta_i = \gamma_i = 0$ as we had assumed. The FOC on

x_i for informed traders and manipulators give $\alpha_i = 0$, $\tau = \bar{w}$ and

$$\begin{aligned} \beta_i &= (1/\lambda + \beta(N-1))\eta_i/2, \\ (9) \quad \gamma_i &= \frac{-\rho}{N\rho + 2(2-\rho)(1+k\lambda)}, \\ \gamma_0 &= \frac{N\rho + 2 - \rho}{N\rho + 2(2-\rho)(1+k\lambda)}, \end{aligned}$$

$$(10) \quad 1/\lambda\beta = N + 1 + 2S_e/S_v.$$

The second-order condition (SOC) for manipulators is $0 \geq \bar{\pi}''_0 = -2\lambda(1+k\lambda)$, and for all other traders is $0 \geq \bar{\pi}''_i = -2\lambda$. If we assume a strict interior optimum of x_i^* , then $\lambda > 0$. And since $r_0 = 0$, we have $S_E = r/2\lambda$, for $r \equiv \sum_{i=1}^T r_i \geq 0$.

Equation (10) can be substituted into equation (7) to give our first result.

Theorem 1. The error $P - v$ of price as an estimate of asset value is distributed with mean zero, and variance

$$\Phi = S_\theta + S_v \left(1 - \frac{N}{N + 1 + 2S_e/S_v} \right),$$

which (for fixed number of traders N , asset variance S_v and price setting variance S_θ) depends only on (and is increasing in) the equilibrium asset value clue variance S_e .

Thus there is on average no net bias, and the manipulator bias distribution $w \sim N(\bar{w}, S_w)$ can only affect price accuracy by affecting the informed trader's choices of clue variance S_e .

The FOC and SOC for the choice of S_i are $\bar{\pi}'_i(S_i) = 0$ and $\bar{\pi}''_i(S_i) \leq 0$, where these expressions hold constant β , η , S_e , λ , γ , γ_0 , while allowing β_i , η_i to vary with S_i . We assume that $c''(S_i)$ is large enough to induce a strict interior optimum, so that $\bar{\pi}''_i(S_e) < 0$, $S_e < \infty$.

Solving equations (8), (9) and (10) gives a quadratic equation, whose solution is

$$(11) \quad \beta = \beta_0 \pm \sqrt{\beta_0^2 + S_X/N(S_v + S_e)}$$

where

$$\beta_0 = \frac{r(N + 1 + 2S_e/S_v)}{2N(S_v + S_e)}.$$

The negative sign solution violates the SOC on x_i^* ($\lambda > 0$), so only the positive sign solution is valid.

Equations (10) and (11) show that an equilibrium exists, with λ and β finite and non-zero, when *any* one of S_b , S_w or r is non-zero (and all are finite). Thus liquidity trading is not required to produce a prediction market equilibrium; a manipulator with unknown bias or any trader irrationality can substitute as a source of noise trading.

We are interested in how the price error Φ changes as we change the parameters $(k, \bar{w}, S_w, S_\delta)$ that describe the manipulator. We already know that \bar{w} has no effect on prices. Let us now consider a proportional variation of S_w and S_δ together. That is, let us hold ρ constant while we increase S_w , keeping a proportional $S_\delta = S_w(1 - \rho)/\rho$. This

variation plausibly describes a change to a situation with ‘more’ manipulation and similar clues about such manipulation.

We can show that as S_w and S_δ increase together in this way, the equilibrium trader clue error S_e will decrease, lowering average price error Φ .

Theorem 2. Changes in the mean manipulator bias \bar{w} have no effect whatsoever on prices. When the variance of manipulator bias S_w and bias clue error S_δ increase in the same proportion, trader asset clue variance S_e and price error Φ both decrease.

Proof. The first claim follows trivially because \bar{w} only affects τ , which does not affect price P . Regarding the second claim, if we define $\hat{\pi}(S_i) \equiv \bar{\pi}_i(S_i) + c(S_i)$, then the FOC for S_i can be written $c' = \hat{\pi}' = -\lambda\beta_i^2$ and the SOC can be written $c'' \geq \hat{\pi}''$, where only β_i, η_i vary with S_i . If we differentiate this FOC with respect to S_w and collect terms appropriately, we find

$$(12) \quad (\hat{\pi}'' - c'') \frac{dS_e}{dS_w} = - \frac{\partial}{\partial S_w} \Big|_{S_e} \hat{\pi}' = \frac{\partial}{\partial S_w} \Big|_{S_e} \left(\frac{\beta_0 + \sqrt{\beta_0^2 + (S_X/[N(S_w + S_e)])}}{N - 1 + 2/\eta} \right).$$

In the far right term, when S_e is held constant then only S_X depends on S_w . So if we assume a strict interior optimum of S_i , with the SOC holding strictly ($\hat{\pi}'' - c'' > 0$), we must have

$$-\text{sign} \left[\frac{dS_e}{dS_w} \right] = \text{sign} \left[\frac{\partial S_X}{\partial S_w} \Big|_{S_e} \right].$$

Substituting proportional variation of S_w and S_δ into our expression for S_X , and differentiating, we find

$$\begin{aligned} \frac{\partial S_X}{\partial S_w} \Big|_{S_e} & \left(1 + 4kS_w(2 - \rho) \frac{N\rho(1 - \rho) + (2 - \rho)^2}{(N\rho + 2(2 - \rho)(1 + k\lambda))^3} \frac{\partial \lambda}{\partial S_X} \Big|_{S_e} \right) \\ & = \frac{N\rho(1 - \rho) + (2 - \rho)^2}{(N\rho + 2(2 - \rho)(1 + k\lambda))^2}. \end{aligned}$$

If we take the derivative of λ with respect to S_X , we can rearrange and find that S_X increases in S_w , and hence S_e and Φ decrease in S_w , when

$$1 > \frac{\beta - 2\beta_0}{\beta - \beta_0} \times \frac{2(2 - \rho)k\lambda}{N\rho + 2(2 - \rho)(1 + k\lambda)} \times \frac{S_w}{S_w + S_l \frac{(N\rho + 2(2 - \rho)(1 + k\lambda))^2}{N\rho(1 - \rho) + (2 - \rho)^2}}.$$

The right-hand side here is a product of three non-negative terms, each of which is no greater than one, and one of which (the middle) is strictly less than one. QED.

III. DISCUSSION

These results may seem surprising at first, but should become less so once one recognizes the commonality between manipulators and other noise traders, and once one distinguishes between *ex ante* and *ex post* manipulation.

In ordinary financial markets, many traders base their trades on private information about common asset value, and so market makers who set prices try to infer common asset value from the set of trades. Such markets, however, also include *noise trades*, i.e. trades best explained by something other than information about common asset value. Holding all else constant, each noise trade randomly disturbs the trade set, and would seem to reduce the ability of observers to infer information about the asset value from that trade set.

Since the largest financial markets have the most and largest noise trades, one might expect such markets to also have the least accurate prices. But in fact the largest financial markets tend on average to have the *most* accurate and informative prices. The key to resolving this conflict is to see that all else is not equal. Other traders who anticipate a larger volume of noise trades change their trading behaviour in response.

First, market makers who set prices adjust their interpretation of the trade set to take into account any expected imbalance of noise trades. For example, if on a certain day, tax incentives will make noise traders tend to sell more than they buy of a particular asset, then observing more total sales than buys will not by itself be taken as a negative sign about that asset. It will only be a bad sign if there are even more sales, relative to buys, than expected.

Second, non-noise traders with private information about common asset value will increase the volume of their trades to compensate for the larger volume of noise trades. Since this larger volume increases their profit from using any particular source of information, these traders also respond by increasing their investment in information. Together, these effects can easily explain why high volume financial markets have better price accuracy.

Now consider a manipulator. All else equal, a manipulator who has a different target price or a different strength of preference for that target will make a different manipulation trade. If this realized manipulation preference is uncorrelated with the common asset value, the resulting trade *is* a noise trade, in the relevant sense of being best explained by something other than information about the asset value. So as with other kinds of noise trades, we expect market makers who set prices to correct for any expected imbalance in manipulation trades. If manipulators are expected to on average buy to push prices up, then market makers will only let prices rise if there are even more extra buys than expected.

Holding constant all other trading behaviour, it is clear that *ex post* manipulation must be feasible, in the sense that the final price will in fact vary predictably with details of the manipulator's actual trade. If the manipulator buys more, for example, then the price will in fact be higher. This does not, however, imply the feasibility of *ex ante* manipulation, i.e. that prices become on average more distorted or less accurate with the introduction of a more motivated or varied manipulator. When we average over the 'noise' of different possible manipulator desires, prices can become *more* accurate due to other traders anticipating and responding to manipulator trades.

Of course the fact that we have a particular model illustrating these results hardly implies that these results always hold in every context. Our model assumes risk-neutrality, normally distributed values and signal errors, interior choices of information quantity, only quantal response-type irrationality, no meta-signals about the signals of other

agents, no transaction costs of trading, no budget constraints, and a single rational manipulator with quadratic manipulation preferences and a commonly known strength of desire to manipulate. However convenient these assumptions may have been for solving the model, one can reasonably question the empirical relevance of models based on them.

We have, for example, assumed away budget constraints, yet budget constraints are clearly an important feature of many real prediction markets, such as the Iowa Electronic Exchange. Budget constraints limit both the trades of a manipulator, and of those who intend to counter manipulator trades, and so manipulation may be more feasible when the first group is less constrained than the second. For example, when all traders face the same budget limit, then there could be problems if most traders were part of a conspiracy pursuing a common manipulation objective.

Consider also risk aversion, which can function much like a budget limit in limiting trades. On the one hand, prediction markets almost invariably feature small trading volumes and relatively small stakes. As Rabin (2000) argues, imputing substantial risk aversion to agents in such low stakes environments requires quite implausible global levels of risk aversion. On the other hand, many kinds of small stakes laboratory and field behaviour are often said to be best explained via something akin to risk aversion. Once we have better models of such behaviour, we will want to apply such models to small stakes prediction markets.

The assumption of normal distributions for uncertain asset value and signal errors is clearly less general than one would like. But, as is well known in the market microstructure literature, these models become vastly more complex with most other distributions. We still know relatively little about non-normal cases.

Our model focuses on a prediction market with a single manipulator, but we believe (though we have not proven) that similar, if more complex, results follow when our model is modified to have many manipulators with differing target prices t and strengths of desire k . The linear trading behaviour of each manipulator should add up to a linear behaviour for the group, a behaviour equivalent to that of a single manipulator with some effective target price \tilde{t} and strength of desire \tilde{k} . Normally distributed signals about individual targets should also be equivalent to a normally distributed signal about the effective group target.

IV. CONCLUSION

The novelty, complexity and anonymity of prediction markets have led some observers to fear that prediction markets will be misused. One typical concern is that traders will commit acts of sabotage or terrorism in order to earn money in the market. Another is that traders will take losses in these markets in order to deceive decision makers. While the low volume of prediction markets limits their likely effect on extra-market behaviour, that low volume seems to also make it cheap for traders to manipulate perceptions with their trades. This has led to fears that manipulators may decrease the predictive accuracy of prediction market prices.

Historical, field and laboratory data, however, have usually failed to find substantial effects of such manipulation on average price accuracy. Since previous models have found that increases in noise trading can increase the accuracy of thin markets, by increasing the rewards to informed trading, we might hypothesize that manipulators are like noise traders, in that both make trades based on considerations other than their best estimate of the asset value.

This paper has presented a standard Kyle-style market microstructure model that formally illustrates this idea. Adapting this standard model to the case of a low stakes prediction market, and adding in a manipulator, we find that a manipulator can substitute for a liquidity trader or for trader irrationality to produce a prediction market equilibrium. A manipulator with a known target price preference has no effect on the market price, but one whose target price is unknown is much like a noise trader with an unknown trading quantity. The prospect of trading against someone who trades on non-asset value considerations can entice other traders to become better informed, increasing average price accuracy.

While the social desirability of prediction markets remains an open question, this model suggests that concerns about manipulators reducing average price accuracy may be overstated.

ACKNOWLEDGMENTS

For their comments, we thank Bryan Caplan, Dan Houser, John Ledyard, Kevin McCabe, Dave Porter, Emile Servan-Schreiber, Koleman Strumpf, Alex Tabarrok, Eric Zitzewitz, several anonymous referees, and participants at the 2007 Growth of Gambling and Prediction Markets Conference. For financial support we thank the Center for Study of Public Choice, the Mercatus Center, and the International Foundation for Research in Experimental Economics.

NOTES

1. Actually, PAM would have focused on aggregate geopolitical trends, such as how the chance of political unrest in Saudi Arabia depends on whether US troops leave there (Polk *et al.* 2003).
2. Senators, reporters, and economists complained:

Terrorists themselves could drive up the market for an event they are planning and profit from an attack, or even make false bets to mislead intelligence authorities.

(Wyden and Dorgan 2003)

Would-be assassins and terrorists could easily use disinformation and clever trading strategies to profit from their planned misdeeds while distracting attention from their real target.

(Pearlstein 2003)

Trading . . . could be subject to manipulation, particularly if the market has few participants – providing a false sense of security or . . . alarm. . . the lack of intellectual foundation or a firm grasp of economic principles – or the pursuit of other agendas – has led to a proposal that almost seems a mockery of itself.

(Stiglitz 2003).

3. Other kinds of ‘manipulation’ not considered here include hidden actions that influence events, deceptive cheap talk intended to influence prices (Allen and Gale 1992), and strategic contrary trading made by an informed trader to control the rate at which his information is revealed (Chakraborty and Yilmaz 2004).
4. Of course it is possible for the average social harm from price errors to increase even as the average price errors decrease, if social harm varies in complex ways with price errors. As we lack a model of social harm from price errors, we do not further consider this possibility here.
5. This ends up being the same as having a preference over the beliefs of outside observers.
6. We introduce only one manipulator for simplicity and to study a best case for manipulation. This single trader can represent a group of manipulators with maximal internal coordination.

REFERENCES

- ALLEN, F. and GALE, D. (1992). Stock-price manipulation. *Review of Financial Studies*, 5(3), 503–29.
- BERG, J. E. and RIETZ, T. A. (2003). Prediction markets as decision support systems. *Information Systems Frontiers*, 5(1), 79–93.

- BRUNNERMEIER, M. K. (2001). *Asset Pricing under Asymmetric Information*. New York: Oxford University Press.
- CAMERER, C. (1998). Can asset markets be manipulated? A field experiment with racetrack betting. *Journal of Political Economy*, **106**, 457–82.
- CHAKRABORTY, A. and YILMAZ, B. (2004). Manipulation in market order models. *Journal of Financial Markets*, **7**(2), 187–206.
- CHEN, K.-Y. and PLOTT, C. R. (2002). Information aggregation mechanisms: concept, design and implementation for a sales forecasting problem. Discussion paper 1131, California Institute of Technology.
- GOEREE, J. K. and HOLT, C. A. (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, **91**(5), 1402–22.
- HANSEN, J., SCHMIDT, C. and STROBEL, M. (2004). Manipulation in political stock markets – preconditions and evidence. *Applied Economics Letters*, **11**(7), 459–63.
- HANSON, R. (1999). Decision markets. *IEEE Intelligent Systems*, **14**(3), 16–9.
- , OPREA, R. and PORTER, D. (2006). Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, **60**(4), 449–59.
- KUMAR, P. and SEPI, D. J. (1992). Futures manipulation with ‘Cash Settlement’. *Journal of Finance*, **47**(4), 1485–502.
- KYLE, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, **53**(6), 1315–36.
- (1989). Informed speculation with imperfect competition. *Review of Economic Studies*, **56**(3), 317–55.
- LO, A. W. (1997). *Market Efficiency: Stock Market Behaviour in Theory and Practice*. Lyme: Elgar.
- MCKELVEY, R. D. and PALFREY, T. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, **10**, 6–38.
- O’HARA, M. (1997). *Market Microstructure Theory*. Oxford: Blackwell.
- OPREA, R., PORTER, D., HIBBERT, C., HANSON, R. and TILA, D. (2007). Can manipulators mislead market observers? <http://hanson.gmu.edu/judges.pdf>.
- PEARLSTEIN, S. (2003). Mislacing trust in the markets. *Washington Post*, 30 July.
- POLK, C., HANSON, R., LEDYARD, J. O. and ISHIKIDA, T. (2003). The policy analysis market: an electronic commerce application of a combinatorial information market. *Proceedings of the 4th ACM Conference on Electronic Commerce* (pp. 272–73). New York: Association for Computing Machinery.
- RABIN, M. (2000). Risk aversion and expected-utility theory: a calibration theorem. *Econometrica*, **68**(5), 1281–92.
- SPANN, M. and SKIERA, B. (2003). Internet-based virtual stock markets for business forecasting. *Management Science*, **49**(10), 1310–26.
- SPIEGEL, M. and SUBRAHMANYAM, A. (1992). Informed speculation and hedging in a noncompetitive securities market. *Review of Financial Studies*, **5**(2), 307–29.
- STIGLITZ, J. E. (2003). Terrorism: there’s no futures in it. *Los Angeles Times*, 28 July, p. B13.
- STRUMPF, K. S. and RHODE, P. W. (2004). Historical presidential betting markets. *Journal of Economic Perspectives*, **18**(2), 127–41.
- WOLFERS, J. and ZITZEWITZ, E. (2004). Prediction markets. *Journal of Economic Perspectives*, **18**(2), 107–26.
- WYDEN, R. and DORGAN, B. (2003). Wyden, Dorgan call for immediate halt to tax-funded ‘Terror Market’ scheme. Press release, 28 July.