Preconditions for Information Aggregation in Prediction Markets

Saša Pekeč

Decision Sciences
Fuqua School of Business
Duke University
pekec@duke.edu

Abstract. We study necessary and sufficient conditions for information to aggregate in a stylized one-shot unit-demand uniform-price prediction market. We show that, in the increasing symmetric equilibrium, the market price properly aggregates information if and only if (i) the number of realized trades is non-negligible compared to number of market participants, and (ii) there is no uncertainty about the proportion of market participants that will end up trading. The latter condition heavily depends on distributional assumptions on the structure of privately held information, as well as on participation decisions of all potential market participants. Therefore, information aggregation is not automatic and requires these factors to be properly aligned.

1 Introduction

Use of reliable tools for estimating and forecasting uncertainty is critical in implementation of many decision-making models. It is often the case that information from variety of sources has to be aggregated to compile an estimate or a forecast for an uncertainty parameter of the model. Standard statistical techniques rely on aggregating large sets of data that and typically disregard potential differences in quality among data points. (For example, the idea of independent draws from a distribution is the key for most of limiting results in statistics that sampling procedures rely on.) Some expert opinion elicitation procedures developed by decision analysts attempt to take into account the quality of information to be aggregated: a typical example are different scoring rules. The latter approach works well for eliciting information from an individual but it does not generalize to aggregating information from a large number of diverse sources. Prediction markets seem to be merging desirable properties of all these techniques: they are designed with the goal of aggregating diversely held information while implicitly controlling for quality by requiring bets to be placed on information provided.

For the simplest example of a prediction market, consider a binary event E (e.g., $E =$*Brazil will win 2014 World Cup.* or $E =$*Barack Obama will be reelected in the 2012 US presidential election.* or $E =$*Drug XYZ will get FDA approval.* or ...) and denote two possible outcomes as $E = 1$ and $E = 0$. The goal is to estimate the probability v that $E = 1$. The prediction market for E would allow creating and trading the **contract** that pays \$1 if $E = 1$, and that pays \$0 if $E = 0$. Note that for an individual whose estimate of v is $\hat{p}$, the expected value of the contract is

$$E[contract] = \hat{p}.$$

Therefore, a risk-neutral individual i holding a belief $\hat{p}_i$ should be willing to buy the contract at any price less than $\hat{p}_i$ and sell the contract at any price higher than $\hat{p}_i$. The prediction market matches potential buyers and sellers and creates multiple copies of the contract and prices it to clear the market. The market price of the contract represents the aggregate opinion of all market participants about v. In other words, the market price p IS the aggregate estimate of the probability v that $E = 1$. If any market participant has a different estimate of v, they could profit by either buying or selling the contract which allows participants to "put their money where their mouth is" and thereby provides opportunity to communicate the perceived quality of $\hat{p}$.

Even this simple model of a prediction market as the model for aggregating information on v relies on some assumptions:

- an agent i holding belief $\hat{p}_i$ should always be willing to trade the contract,
- all agents should be risk-neutral,
- there is no discounting,
- there has to be a trading counterparty for an agent i in order for belief $\hat{p}_i$ to be aggregated into market price v.

In addition, for proper understanding of the information aggregation process in such a market, rules for price formation and market-clearing have to be specified. This is typically modeled via a double-auction:

- agents interested in buying the contract submit prices they are willing to pay (bids): $b_1 \geq b_2 \geq \ldots \geq b_n$ (ordered without loss of generality)
- agents interested in selling the contract submit prices they are willing to sell (asks): $a_1 \leq a_2 \leq \ldots \leq a_m$ (ordered without loss of generality)

The market can then clear at the largest k such that $a_k \leq b_k$, while the contract prices can be anywhere within intervals $[a_1, b_1] \supseteq [a_2, b_2] \supseteq \ldots \supseteq [a_k, b_k]$. Since the purpose of the market is to aggregate information, this is best achieved by uniform price setting, i.e., setting $p \in [a_k, b_k]$.

Note that the above description is overly simplistic because it treats the market as a one-shot (no later trades are possible) and it assumes single contract bid or ask for each agent (unit demand). Most of prediction markets implemented in practice and discussed in theory allow for multiple contract bids (which can simply be modeled as separate single bids) and for sequential nature of the market in which agents can submit bids at any time, so the market price p and the aggregate estimate of the probability v that $E = 1$ evolves over time. However, limitations for information to aggregate that will be discussed in this simplistic market, readily extend in enriched frameworks that allow for sequential trading and multiple demand (since single shot and unit demand remain special instances).

While interest in study and design of prediction markets started with Iowa Electronic Markets, that focuses on predicting US election outcomes, more than twenty years ago [7], the interest in prediction markets and its use in a variety of contexts has exploded in the last few years. Prediction markets are used beyond obvious betting applications (such as tradesports.com and intrade.com) and have been used in corporate context in companies such as BestBuy, Google, HP and others [6, 3, 17] and for predicting a variety of things ranging from electricity demand[4] to underpricing in IPOs[1]. (Some proposed uses turned out to be controversial such as using prediction markets to predict terrorist attacks [8].) Along with experimental and lab studies, there is now a body of methodological work related to prediction markets; check [18, 19, 21] for review). Despite prediction markets reportedly working well in practice and experiments, it is important to understand when can these markets be manipulated and when does information aggregation fail (e.g., see [19, 20, 9, 16]). The issues discussed include misreporting and manipulating the market-clearing procedure thereby distorting information aggregation, issues of initializing the market, agent's information structure and effects of already observed market information such as market prices prior to submitting bid, and the problem of participation of all information holders that was already mentioned.

In terms of microeconomic theory, a prediction market can be viewed as a double auction. There is a considerable literature on equilibria in double auctions. For example, [11] demonstrates the existence of a no-trade equilibrium (everyone asks 1 and bids 0) and at least one non-trivial equilibrium, and [5] establishes efficiency of the latter type of equilibria. However, a critical assumption in this literature is that market participants have private valuations for the object being traded. In contrast, in prediction markets, the contract has a common value and participants hold private beliefs (signals) about that value. Thus, insights from common value auctions are more appropriate in the prediction markets context.

This paper contributes to the stream of research that attempts to identify potential problems with information aggregation in prediction markets by examining plausible market clearing and pricing procedures and constraints these procedures impose for information to aggregate. We first develop a simple informational structure for market participants who are both willing to buy a contract below its expected value, and sell the same contract above its expected value, conditional on the value of their private signal. Next we model inherent uncertainty about the extent of market clearing (i.e., how many buyers will be matched by sellers), and this uncertainty has to be taken into account when analyzing equilibrium behavior of market participants. Finally, we analyze market clearing by decomposing it into analysis of two uniform price auctions, both with uncertainty about the number of items (contracts) that are being auctioned. Information aggregation results from auction theory are then applied into this setting, yielding that information aggregation could fail even in such simplistic prediction markets. This indicates that information could also fail to aggregate in richer models of prediction markets: incorporating additional realistic features such as dependence of private signals or risk-awareness or off-equilibrium behavior or participation decisions (on the buy and/or sell side), etc., would not make information aggregation any easier to achieve.

Section 2 describes the basic model. Section 3 describes relevant information aggregation results for auction procedures. Section 4 relates auction theory results to our prediction market model. Section 5 concludes.

2 A model of the one-shot market

We focus on necessary conditions for information aggregation in a one-shot unit-demand prediction market. There is only one type of the **contract** traded in the market. The seller of the contract pays \$1 to the buyer of the contract if event $E = 1$ happens (and this will be resolved after market clears), otherwise the contract is worthless.

The expected value of the contract is equal to the the probability of $E = 1$, as already noted in the introduction. This value is unknown and the principal feature of prediction markets is that the market price of the contract will reveal the value by properly aggregating diversely held private information about it through market activity.

The value is modeled as a random variable V with twice differentiable probability density function g. Let the support of g be convex, i.e., $\underline{v} < \overline{v} \leq \infty$, and $g(v) = 0$ otherwise.

A market-participant j observes a private *signal* X_j, a random scalar distributed with cumulative distribution function $F(x|v)$ and density $f(x|v)$ conditional on $V = v$, where f is assumed to have a third derivative with respect to v

We also assume that f satisfies the Strict Monotone Likelihood Ratio Property (SMLRP):

$$\frac{f(x|v)}{f(x|v')} > \frac{f(x'|v)}{f(x'|v')} \quad \forall x > x', v > v', \tag{1}$$

which can be interpreted as a condition that guarantees at least some informativeness of the signals with respect to v, and is a standard technical condition needed to ensure equilibrium existence in uniform auctions.

After observing their private signal x_j, all market participants report their willingness to buy a contract at price b (bid) and willingness to sell the contract at price a (ask). Obviously, a rational participant would want to buy low and sell high, i.e., $a > b$ (and this will be the case in the equilibrium presented later). Note that we could assume that all participants report a and b, since $b \leq \underline{v}$ and $a \geq \overline{v}$ cannot result in a loss (and will not clear in the market either).

Without loss of generality, all reports can be represented as

$$\text{bids} : b_1 \geq b_2 \geq \ldots \geq b_N$$

and

$$\text{asks} : a_1 \leq a_2 \leq \ldots \leq a_N.$$

(Subscript does not relate to the identity of the participant, i.e., a_i and b_i could have been submitted by different participants.)

Let k be the largest integer such that $a_k \leq b_k$. Then, k contracts get traded by matching the seller a_i with buyer b_i, $i = 1, \ldots, k$. The contract price is uniform and set to be within $[a_k, b_k]$. (The exact pricing rule becomes irrelevant for asymptotic analysis as long as $\lim_{N \to \infty} b_k - a_k = 0$.)

Note that all market participants face a problem of transforming their signal into their bid b and ask a. Also note that the buyer's problem is symmetric to the seller's problem. In what follows, we will mainly focus on the buyer's problem: buyers face an auction in which there is a principal uncertainty about the number of contracts k that will be sold to buyers, since these depend on asks submitted by the sellers.

We next review information aggregation results for uniform auctions with the uncertain number of contracts sold.

3 Information aggregation in uniform auctions

We consider auctions in which k homogeneous objects (contracts) are sold to the k highest bidders at a uniform price. In order to make analysis tractable, it is important to set the auction price to be independent of any of the winning bids. This is customary done by setting the price at the level of $(k+1)^{st}$-highest bid (ties broken randomly). However, the whole analysis goes through for any price that is independent of and below the k-th highest bid.

Bidders are drawn from a pool of $N \leq \infty$ potential bidders in accordance with an exogenous symmetric stochastic process, $\Omega = \{(n_1, \pi_1), ..., (n_M, \pi_M)\}$, specifying that n_i bidders are present with probability π_i. The n_i are labeled ascendingly: $n_i < n_{i'}$ for $i < i'$. Let $\overline{n} = \sum_{i=1}^{M} \pi_i n_i$, the expected number of bidders. If the number of contracts to be sold is $\mathbf{k} = \{k_1, \ldots, k_M\}$ to sell k_i items if n_i bids are submitted, $k_i < n_i, \forall i$, the resulting stochastic process is

$$\Omega_{\mathbf{k}} = \{(n_1, k_1, \pi_1), ..., (n_M, k_M, \pi_M)\}.$$

$\Omega_{\mathbf{k}}$ is assumed to be common knowledge.(This model of exogenous bidders' participation is standard: the single-object case $k_1 = \ldots = k_M = 1$ has been first studied in [13].) Keeping in mind clearing in prediction markets, note that this model (that still focuses on only one side of clearing) allows not only for uncertainty on the number of contracts being traded, but also about participation, i.e., some market participants might choose not submit their bid b (willingness to pay) and/or ask a (willingness to sell).

We follow the presentation in [10]: for a given k_i and n_i, the fraction of winning bidders is called the *success ratio*, $s_i = k_i/n_i$. The policy $\Omega_{\mathbf{k}}$ represents *proportional selling* if there is (almost) no uncertainty about the success ratio, i.e., if all success ratios $s_i = k_i/n_i$, $i = 1, \ldots, M$, are approximately equal.(Since k_i and n_i are integers, only approximate equality is attainable. However, for large n_i this integer constraint becomes negligible.)

Symmetry allows focusing on realized bidder 1. (Note the departure from convention $b_1 \geq b_2 \geq \ldots$ from the previous section. Bidder 1 here is not necessarily the bidder submitting the highest bid. It is just a bidder we focus our analysis on.) Bidder 1 is said to be *pivotal* if his bid is on the boundary between winning and losing. Let the *pivotal rival* to bidder 1 be the rival bidder with the k_i^{th}-highest signal, of the $n_i - 1$ rivals. Among all n_i bidders, the *price-setter* is the bidder submitting the $(k_i+1)^{st}$-highest bid. Given $\Omega_{\mathbf{k}}$, the identities of these order statistics may be stochastic, but are well-defined.

An auction is defined by $\mathcal{A} = \{\Omega_{\mathbf{k}}, g(\cdot), f(\cdot)\}$, specifying the stochastic structure of bidders and objects sold, the prior distribution of the common value, and the conditional distribution of any one bidder's signal.

To derive equilibrium function, consider the case of a known number of bidders (n) and objects (k), which corresponds to $\Omega_{\mathbf{k}} = \{(n,k,1)\}$. Let v_{nk} denote the usual conditional value function

$$v_{nk}(x,y) = E\left[V \mid X_1 = x, Y_{n-1}^k = y\right],$$

where Y_{n-1}^k is the k^{th}-highest of $n-1$ signals. Then

$$b_{n,k}(x) = v_{nk}(x,x), \tag{2}$$

the expected value given that bidder 1 is pivotal in an n-bidder auction for k objects, is the unique symmetric equilibrium bid function ([14, 15])

The bidding function for an unknown number of bidders that is a weighted average of the bidding functions for each (n_i,k_i) pair. The weights are Bayesian updatings of the probabilities of n_i bidders and k_i objects (the i^{th} component of $\Omega_{\mathbf{k}}$) under the assumption that bidder 1 is pivotal. Denote

$$w_i(x) = \frac{C(i)\int_{\underline{v}}^{\overline{v}} f^2(x|t)F^{n_i-k_i-1}(x|t)\big(1-F(x|t)\big)^{k_i-1}g(t)dt}{\sum_{j=1}^{M} C(j)\int_{\underline{v}}^{\overline{v}} f^2(x|t)F^{n_j-k_j-1}(x|t)\big(1-F(x|t)\big)^{k_j-1}g(t)dt}. \tag{3}$$

where

$$C(j) = \frac{\pi_j n_j (n_j - 1)!}{(k_j - 1)!(n_j - k_j - 1)!}$$

Theorem 1. *(Harstad et al. [10]) If there exists a symmetric equilibrium $b^*: \Re \to \Re$ in increasing strategies for $\mathcal{A} = \{\Omega_{\mathbf{k}}, g(\cdot), f(\cdot)\}$, then it is*

$$b^*(x) = \sum_{i=1}^{M} w_i(x) b_{n_i,k_i}(x), \tag{4}$$

where b_{n_i,k_i} and w_i are defined by (2) and (3). Conversely, if b^* is increasing, then it is the unique symmetric equilibrium in increasing strategies.

Information aggregation will be investigated on sequences of auctions $\mathcal{A}_1, \mathcal{A}_2, \ldots$, denoted by $\{\mathcal{A}_\zeta\}_{\zeta=1,2,\ldots}$, along which the numbers of bidders and objects grow large while, for expositional clarity, keeping the probability vector $(\pi_1, \ldots, \pi_M)$, common value distribution, and bidders' signal distribution fixed. Thus, all sequences of auctions $\{\mathcal{A}_\zeta\}$ considered below satisfy:

$$\begin{aligned} &[\text{i}]\ n_{1_\zeta} \to \infty \text{ as } \zeta \to \infty \\ &[\text{ii}]\ \pi_{i_\zeta} = \pi_i > 0 \text{ for all } i = 1, \ldots, M, \text{ all } \zeta = 1, 2, \ldots \\ &[\text{iii}]\ g_\zeta(\cdot) = g(\cdot) \text{ and } f_\zeta(\cdot) = f(\cdot) \text{ for all } \zeta = 1, 2, \ldots \end{aligned} \tag{5}$$

The sequence of auctions $\{\mathcal{A}_\zeta\}$ **aggregates information** if

$$\forall \varepsilon > 0 \;\; lim_{\zeta\to\infty} \Pr\left[|P^\zeta - V| > \varepsilon \right] = 0.$$

Theorem 2. *(Harstad et al. [10]) Consider a sequence of auctions $\{\mathcal{A}_\zeta\}_{\zeta=1,2,\ldots}$ such that b^*_ζ is a symmetric equilibrium for each $\mathcal{A}_\zeta$. $\{\mathcal{A}_\zeta\}$ aggregates information if and only if there exists $s \in (0,1)$ such that $\lim_{\zeta\to\infty} s_{i_\zeta} = s$ for all $i = 1, \ldots, M$.*

Therefore, for information to aggregate, there cannot be any uncertainty about the success ratio. (Note that if the asymptotic success ratio is $s = 0$ or $s = 1$, the question of information aggregation depends on further assumptions about model elements; see [15]).

4 Information aggregation in the one-shot prediction market

Note that the information aggregation condition from the theorem above applies to both the buyers auction and the sellers auction. Since information aggregation cannot be assumed even in a single uniform auction with uncertainty about the number of items auctioned (contracts), it is not surprising that it will be hard to achieve it in a prediction market that requires information to aggregate in a double auction.

We first describe equilibrium behavior in the prediction market, which is a direct application of Theorem 1 applied to both buy side and sell side of the market. The sell side uses the fact that, analogous to the buy side, in equilibrium asking price a is set to be the expected value of the contract conditional on the signal being tied with the k^{th} lowest signal among $n-1$ rivals.

Theorem 3. *Let $\Omega_{\mathbf{k}}$ describe the number of market participants submitting bids and asks and the number of contracts traded. If there exists a symmetric equilibrium $(b^*(x), a^*(x))$ that is increasing in signal value x, then it is*

$$b^*(x) = \sum_{i=1}^{M} w_i(x) b_{n_i,k_i}(x),$$

$$a^*(x) = \sum_{i=1}^{M} w_i(x) b_{n_i,n_i-1-k_i}(x), \tag{6}$$

where $b_{n_i k_i}$ *and* w_i *are defined by (2) and (3). Conversely, if* $b^*(x)$ *and* $a^*(x)$ *are increasing in* x*, then* $(b^*(x), a^*(x))$ *is the unique symmetric equilibrium in increasing strategies.*

Recall that, for any pair (n_i, k_i) and any signal x, $b_{n_i,k_i}(x)$ is the expected value of the contract assuming that x is tied with the k^{th} highest of the remaining $n-1$ signals, and $b_{n_i,n_i-1-k_i}(x)$ is the expected value of the contract assuming that x is tied with the k^{th} lowest of the remaining $n-1$ signals. Since k, the number of contracts traded, is at most $n/2$, assuming that x is among k highest signals implies $b_{n_i,k_i}(x) < x$ and, similarly, assuming that x is among k lowest signals implies $x < b_{n_i,n_i-1-k_i}(x)$. Therefore,

Corollary 1 $b^*(x) < x < a^*(x)$.

Thus, in the equilibrium, market participants will attempt to buy a contract at a low price below their private signal and sell the contract at a high price above their signal. A trade between two participants s and t happens only if $a^*(x_s) < b^*(x_t)$, i.e., only if intervals $[b^*(x_s), a^*(x_s)]$ and $[b^*(x_t), a^*(x_t)]$ are disjoint (which implies that signals x_s and x_t are not close, i.e., trade happens only if privately held signals are sufficiently different). Thus, the number of contracts traded depends on the number of such disjoint pairs of intervals, all separated by the market price p.

Proposition 2 *The number of contracts cleared in a uniform-price prediction market with* n *participants with private signals* x_j *is a random variable*

$$k = \max_p \min\{|\{a^*(x_j) \leq p\}|, |\{b^*(x_j) \geq p\}|\}$$

The uncertainty about the number of contracts traded is non-negligible even when there is no uncertainty about the number of market participants (i.e., when everyone submits both b^* and a^*), and the probabilities of market clearing $k = 0, 1, 2, \ldots$ contracts depends on distributional assumptions.

Even in the simplest case of two market participants (here either a single contract gets traded or no contracts get traded) and a uniform distribution of v and a uniform signal distribution $f(x|v) = U[v-c, v+c]$ (where $c < 1/2$ is some constant), it can be shown that there exists a $d < c$ such that $b^*(x) = x - d$ and $a^*(x) = x+d$ (see, e.g., [12]). Then, the probability of a trade is the probability of $[x_1 - d, x_1 + d]$ and $[x_2 - d, x_2 + d]$ being disjoint; it is straightforward to calculate this to be $(1-d/(2c))^2$. Therefore, probability of no trade is $d/c-(d/2c)^2$. Precise calculation of probabilities of $k = 0, 1, 2, \ldots, n/2$ trades is much harder for large n and for value and signal distributions that are hard to handle analytically. What is important for our analysis is to note that these probabilities exist and that the uncertainty about the number of trades depends on distributional assumptions.

Given $\Omega_{\mathbf{k}} = \{(N, 1, \pi_1), (N, 2, \pi_2), \ldots, (N, \lfloor N/2 \rfloor, \pi_{\lfloor N/2 \rfloor})\}$, b^* and a^* are given by Theorem 3 and define the distribution of k, the number of contracts traded, as described in Proposition 2. That distribution has to match the distribution implied by $\Omega_{\mathbf{k}}$, i.e., probability of k contracts traded should be equal to π_k. If this were not the case, then $\Omega_{\mathbf{k}}$ could not have been common knowledge in the first place. One could attempt to find exact values for $\Omega_{\mathbf{k}}$ by a sequence of iterative applications of Theorem 3 and replacing π_i in $\Omega_{\mathbf{k}}$ according to the probabilities from Proposition 2 along the way. However, there is no guarantee of convergence since the increasing symmetric equilibrium might not always exist. In such ill-behaved instances, the presented model cannot describe even basic parameters of the market, let alone information aggregation. These instances are characterized by ex ante symmetric market participants somehow having to select qualitatively different strategies and/or by strategies in which higher signal values might mean lower bid or ask values. On the positive side, our model provides necessary and sufficient conditions for information aggregation in the prediction markets in which the increasing symmetric equilibrium exists.

Theorem 4. *The one-shot prediction market in which the increasing symmetric equilibrium exists, aggregates information if and only if*

$$0 < \lim_{N\to\infty} \frac{\underline{k}_N}{N} = \lim_{N\to\infty} \frac{\overline{k}_N}{N} < 1$$

where $\underline{k}_N$ and $\overline{k}_N$ are the minimum and the maximum of the support of the number of contracts traded with N market participants.

This is a straightforward application of Theorem 2 to $\Omega_{\mathbf{k}}$ with uncertainty only about the number of contracts traded (as described by Proposition 2), and using the fact that $\Omega_{\mathbf{k}}$ is the same for buy side auction and sell side auction. Note that the condition for information aggregation is not automatic and depends on the properties of distributions describing the contract value and private signals.

5 Concluding remarks

Information aggregation is not automatic in a one-shot prediction market with ex ante symmetric unit-demand market participants. The critical obstacle to information aggregation in the increasing symmetric equilibrium (when it exists) is inherent uncertainty about the number of contracts being traded since this number depends on the private signals. The extent of this uncertainty depends on the structure of the model primitives, namely on the distribution of privately held signals. Only when this uncertainty is negligible relative to the number of market participants, information will aggregate (Theorem 4). It is not surprising that information aggregates only if a non-negligible number of market participants end up trading but the condition that there can be no (asymptotic) uncertainty on the ratio of market participants involved in a trade appears restrictive.

The information aggregation result can be extended to allow for participation uncertainty, i.e., allowing for common knowledge about the probabilities π_i that there will be $n_i \leq N$ bids, $m_i \leq N$ asks and k_i contracts traded. Then, for the buy-side auction $\Omega_{\mathbf{k}} = \{(n_1, k_1, \pi_1), (n_2, k_2, \pi_2), \ldots\}$ and for the sell-side auction $\Omega_{\mathbf{k}} = \{(m_1, k_1, \pi_1), (m_2, k_2, \pi_2), \ldots\}$ and information aggregates if and only if both buy-side and sell-side auction aggregate information, i.e., Theorem 2 has to hold for both auctions as N grows large. Thus, one can attempt to achieve the desired numbers of bids and/or asks in the system with the aim of achieving information aggregation. This could potentially be done by enticing participation on either buy or sell side or by committing to add liquidity when necessary.

One important issue that is beyond the scope of the presented approach and that has therefore not been discussed is the case in which available information consists not only from the private signal but also from the publicly available historical data. In most prediction markets, participants can observe trades, and often even bids and asks in the system that have not cleared yet. Such information would clearly influence the bidding: in case of the market price p close to the probability v of $E = 1$, one would conjecture that such information would be beneficial for information aggregation; similarly if market is currently mispriced, such information would hinder proper information aggregation. When historical market information is available, properly pricing early contracts in the market (or refraining from trade until sufficiently large number of boils and asks is in the system) seems to be of critical importance.

References

1. JE Berg, GR Neumann, and TA Rietz. 2009. Searching for Google's Value: Using Prediction Markets to Forecast Market Capitalization Prior to an Initial Public Offering, *Management Science*, 55, 348–361.

2. C Camerer. 1998. Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting. *Journal of Political Economy*, 106(3), 457–482.
3. KY Chen, C Plott. 2002. Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. *CalTech Social Science Working Paper* No. 1131.
4. P Cramton and LI de Castro. 2009. Prediction Markets to Forecast Electricity Demand. *University of Maryland Working Paper.*
5. MW Cripps and JM Swinkels. 2006. Efficiency of Large Double Auctions. *Econometrica*, 74, 47–92.
6. R Dye. 2008. The promise of prediction markets: A roundtable. *McKinsey Quarterly* Issue 2, 82–93.
7. R Forsythe, F Nelson, GR Neumann and J Wright. 1992. Anatomy of an Experimental Political Stock Market. *American Economic Review*, 82, 1142–1161.
8. R Hanson. 2006. Designing real terrorism futures. *Public Choice*, 128, 257–274.
9. R Hanson, R Oprea and D Porter, 2006. Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior and Organization*, 60, 449–459.
10. RM Harstad, A Pekeč, I Tsetlin. 2008. Information aggregation in auctions with an unknown number of bidders. *Games and Economic Behavior*, 62, 476–508
11. MO Jackson and JM Swinkels. 2005. Existence of Equilibrium in Single and Double Private Value Auctions. *Econometrica*, 73, 93-139.
12. P Klemperer. 1999. Auction theory: A guide to the literature. *J. Econ. Surveys*, 13, 227286.
13. S Matthews. 1987. Comparing auctions for risk averse buyers: A buyer's point of view. *Econometrica*, 55, 633–646.
14. PR Milgrom. 1981. Rational expectations, information acquisition, and competitive bidding. *Econometrica*, 49, 921–943.
15. W. Pesendorfer and JM Swinkels. 1997. The loser's curse and information aggregation in common value auctions. *Econometrica*, 65, 1247–1281.

16. PW Rhode and KS Strumpf. 2008. Manipulation in Political Stock Markets. *University of North Carolina Working Paper*
17. K Strumpf. 2009. Introduction to Special Issue on Corporate Applications of Prediction Markets. *Journal of Prediction Markets* Vol.3, No.1, i–viii
18. G Tziralis and I Tatsiopoulos. 2007. Prediction Markets: An Extended Literature Review. *Journal of Prediction Markets*, Vol.1, No.1, 75–91
19. J Wolfers and E Zitzewitz. 2004. Prediction Markets. *Journal of Economic Perspectives* 18(2), 107–126.
20. J Wolfers and E Zitzewitz. 2006. Five Open Questions About Prediction Markets. In *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*, eds: R. Hahn and P. Tetlock, The AEI=Brookings Press, Washington, D.C.
21. J Wolfers and E Zitzewitz. 2008. Prediction Markets in Theory and Practice. In *The New Palgrave Dictionary of Economics*, 2nd ed., ed. L. Blume and S. Durlauf.