

When Do Markets with Simple Agents Fail?

Abraham Othman
Computer Science Department
Carnegie Mellon University
aothman@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
sandholm@cs.cmu.edu

ABSTRACT

We consider (prediction) markets where myopic agents sequentially interact with an automated market maker. We show a broad negative result: by varying the order of participation, the market's aggregate prediction can converge to an arbitrary value. In other words, markets may fail to do any meaningful belief aggregation. On the positive side, we show that under a random participation model, steady state prices equal those of the traditional static prediction market model. We discuss applications of our results to the failure of the 1996 Iowa Electronic Market.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Economics, Experimentation, Theory

Keywords

Prediction Markets, Market Equilibrium, Price Theory, Market Failure, Zero Intelligence

1. INTRODUCTION

Prediction markets are markets that trade on future events; by aligning traders' incentives with the elicitation of information, they can produce very accurate forecasts. The best example of this has been the success of the most venerable prediction market, the Iowa Electronic Markets. These markets, administered by academics at the University of Iowa, have consistently produced more accurate predictions for elections than polls or aggregates of polls—often months in advance of the actual vote [Berg et al., 2001].

One of the most notable aspects of markets is their sequentiality. Prices are constantly in flux as traders emerge, interact, and exit. Examining traders' behavioral incentives in markets (e.g., to bluff, withhold information, learn, or reveal truthfully) and the impact of those actions on observed prices is a topic of considerable recent research:

- Feigenbaum et al. [2003] examine an iterated market in which perfectly rational traders, each endowed with

Cite as: When Do Markets with Simple Agents Fail?, Abraham Othman and Tuomas Sandholm, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 865–872

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

a bit of information, submit posterior trades of their expectation of a boolean function based on those bits. They exactly characterize the space of functions for which prices will converge to the correct value, such that all agents will adopt the posterior. Those functions are *weighted threshold functions*. In order to subvert the famous no-trade theorem of Milgrom and Stokey [1982], which would apply to the perfectly rational traders in this setting, traders are compelled (forced) to submit exactly their posterior in every round.

- Chen et al. [2007] show that, if agents receive conditionally independent signals, it is an equilibrium to sequentially adjust market prices to their posterior (given their signal and previous market prices) at an agent's first opportunity when interacting with an automated market maker. It is well known that an automated market maker encourages direct, truthful reporting by myopic players, but their result is interesting because it applies also to more sophisticated agents. Convergence to a common posterior is implicit in this result. They also show that if the signals agents receive about the true state of the world are conditionally dependent, then the solution concept falls apart and a range of strategic behavior can emerge.
- Dimitrov and Sami [2008] examines strategic play by agents in a setting where a range of strategic behavior is allowed. They show converging equilibrium prices, but only in a relatively simple two-state two-agent model and only with the imposition of a discount factor.
- Ostrovsky [2009] exactly characterizes the set of securities for which every Bayesian perfect equilibrium aggregates information for all possible prior distributions of trader information. Their results do not rely on the independence of prior signals or on the imposition of a discount factor, and hold for any number of agents. However, the work is fundamentally nonconstructive; the existence of such BPE is not proven.
- Shi et al. [2009] offer a compelling description of a prediction market populated by simple traders. Within a traditional market interpretation of their *one-round prediction mechanism*, agents act to change market prices to reflect their private beliefs, participate exactly once according to some (pre-determined) order, and are unaffected by both the actions of past agents as well as the possibility of action from future agents.

Since agents are not budget constrained and the market interpretation of the one-round prediction mechanism is so simple, the observed prices from such a scheme are quite trivial; prices are simply the private belief of the last agent to participate in the market.

In some models of interaction prices converge toward a shared posterior consensus. However, recent research by Graefe and Armstrong [2008] has suggested that models of markets converging towards a shared posterior consensus are a poor fit for real interactions among human traders. In the laboratory market experiments they studied, most participants thought (incorrectly) that they could achieve a higher payoff by changing, rather than accepting, the final price reached in their experimental market. Traders did not embrace the final market price as a universal consensus, they were simply unable to change the final price because of budget constraints.

Another stream of economic research suggests, however, that markets need not reach a shared consensus among rational traders to produce desirable results. In particular, the famous *Hayek hypothesis* contends that markets are good aggregators of information even if participants are unsophisticated [Smith, 1982]. An interesting test of this hypothesis was performed by Gode and Sunder [1993], who examined how simple computational agents can produce market prices very close to those observed from profit-motivated human traders. Their work introduced the concept of *Zero Intelligence (ZI) agents*, who do not remember past actions, learn, or attempt to maximize their utility.

In the vein of examining the emergent properties from the interaction of simple traders, work by Manski [2006] has contended that markets populated by simple traders should reach roughly informative price predictions. Manski demonstrated that a market composed of simple, budget-limited agents should produce roughly efficient prices, though we should expect prices to exhibit the *longshot bias*, whereby less likely events (longshots) are overpriced and more likely events are underpriced. This prediction has been validated by several quantitative studies of prediction market data [Wolfers and Zitzewitz, 2006, Corwin and Othman, 2008]. More recent work by Othman [2008] that simulates the interactions of bids and asks in a standard double auction shows that expected transaction price (i.e., the mean price at which a trade occurs) with ZI traders should be close to the predictions of Manski’s model, although actual expected prices have no simple closed form and are the product of a complex iterated integral. Our work moves these simple static models into a more realistic dynamic setting.

What can observed prices look like over time when simple agents interact? We construct two different interaction models that parallel the classic worst-case and average-case dichotomy from the analysis of algorithms. We show that under an adversarial ordering model, market prices can be arbitrarily uninformative. In contrast, we also show that under a random ordering model, we should expect the market prices we observe to be informative—to be close to equilibrium prices in Manski’s standard static model of price formation. We discuss how our results are consistent with the performance of the 1996 Iowa Electronic Market, perhaps the most notable failure of prediction markets in practice [Berg et al., 2001].

2. MODELING MARKETS AND AGENTS

We begin by discussing how our market model works, and then how the agents in our model interact with the market.

2.1 Pricing rules

The markets we consider can be regarded as stylized abstractions of a broad class of actual, real-world markets. This paper only considers markets over *binary* events, which involve a partition over the future into two exhaustive sets. An example is a contract that pays a dollar if Manchester United defeats Arsenal in their next match. Prices in markets for binary events can be specified by a single scalar value, because no-arbitrage conditions imply the price for the complementary event.

We consider interaction with the market through the mediation of a market maker operating according to a *pricing rule*. The pricing rule defines a structured way of adjusting prices in response to instantaneous aggregate demand; it is a function, $p(q)$, that maps from aggregate demand to a price. Note that the price offered by the market maker is valid only for the purchase or sale of an infinitesimal quantity; as the agent interacts with the market maker, the price changes. This is why we say that pricing rules work based on instantaneous aggregate demand—the prices change instantly in response to demands.

Another way of looking at our structure is that it is analogous to a bookmaker setting a line for a match. The bookmaker seeks to equalize the number of dollars on each side of his proposition, so that if he sees too many bets on one team he adjusts his line, making further investment in that team necessarily riskier. A bookmaker makes money through the vigorish (bid/ask spread) he fixes when taking bets. A pricing rule can be considered the bookmaker’s behavior in the limit as the bid/ask spread approaches 0.

Some existing prediction markets fall within our model exactly. Traditional two-sided markets clearly do not operate this way, but our model is a good abstraction away for analyzing one fundamental concept, that of *slippage*. In a traditional market, an agent cannot purchase an unlimited quantity of goods at the lowest asking price. Rather, as he purchases more, ask orders are exhausted and the asking price will rise, causing the agent’s expected gain to “slip”. Pricing rules capture this intuition while providing a tractable formalism for how prices change. Slippage is such a fundamental notion to our understanding markets that automated market makers can be considered qualitatively (though not quantitatively) identical to traditional double auction models [Ostrovsky, 2009].

Not all functions make good or logical pricing rules. In this work, we only consider rules that satisfy technical conditions that fit a reasonable interpretation of how markets should work. We call such rules *normal*.

Definition 1. A *normal* pricing rule is a function $p(q)$ that is continuously differentiable, strictly increasing, and onto $(0, 1)$.

Agents interact in the marketplace by spending their money to obtain shares that will have value at expiry. Say an agent seeks to move the market from q aggregate demand to $q' > q$ aggregate demand (he is betting for the event in question to

occur). The agent pays

$$\int_q^{q'} p(q) dq$$

and receives a payout of $q' - q$ dollars if the event in question occurs. Now consider an agent wishing to bet against the event, moving the aggregate demand from q to $q'' < q$. That agent pays

$$\int_{q''}^q 1 - p(q) dq$$

and receives a payout of $q - q''$ if the event in question does not occur. We denote by $c(p, p')$ the net cost of moving the market from price p to price p' . An agent who moves the market from price p to price p' and then back to price p has no expected benefit and no net cost—his costs are precisely equal to the payout he is guaranteed to receive regardless of whether the event occurs or not.

One property of pricing rules is that they are history-independent. The market maker will adjust prices based only on the current price, regardless of whether that price was generated after the participation of five agents or five thousand agents.

Because a normal pricing rule is strictly increasing, it is invertible. We denote the inverse of the pricing rule as $q(p)$. The difference in outstanding coupons (or shares) to be redeemed if the event in question occurs between a price of p and a price of p' is given by $q(p) - q(p')$.

PROPOSITION 1. *It takes a finite amount of money to move a normal pricing rule from any initial price in $(0, 1)$ to any other price in $(0, 1)$.*

PROOF. The net cost of changing the market between two prices is bounded above by the difference in quantities associated with each price. Because the rule is onto $(0, 1)$, this difference in quantities is necessarily finite. ■

We consider all markets as starting from an initial point of .5. Given a normal pricing rule, such a consideration is without loss of generality for our results.

2.2 Agent actions and beliefs

Agents are myopic and act according to three guidelines:

1. Agents have fixed beliefs about the value of a contract and are risk neutral.
2. Agents have a fixed, finite budget.
3. Agents participate exactly once, acting in the market and then exiting.

These assumptions are found widely throughout the literature. For instance, Shi et al. [2009] explores and justifies a model where agents follow the first and third properties (but the agents in their work implicitly have arbitrarily large budgets). Lambert et al. [2008] explore and justify a model where agents follow the first and second properties, but their agents cannot follow the third because the self-financed wagering mechanism they investigate is not an explicitly sequential market. Similarly, Manski [2006] and Othman [2008] both use the first two assumptions, but their static models rule out the third.

3. MARKET FAILURE AND UNINFORMATIVE PRICES

This paper investigates the results of different orderings of agents participating in markets. We denote by $p(o)$ the price after the participation of all agents in the specified order o , and we denote by $p_k(o)$ the observed price after the first k agents in the ordering o participate.

We are now ready to give our main results regarding how participation order affects prices. One might hope that regardless of agent ordering, the final price we observe after the participation of all the agents will be roughly the same. We show that this is not the case. With as few as two agents we can see final prices that diverge arbitrarily.

PROPOSITION 2. *For any $\epsilon > 0$, there exists a set of two agents with finite budgets such that $|p(\{1, 2\}) - p(\{2, 1\})| \geq 1 - \epsilon$.*

PROOF. Let both agents have a budget of $\max\{c(1 - \epsilon/2, \epsilon/2), c(\epsilon/2, 1 - \epsilon/2)\}$ with beliefs $b_1 = \epsilon/2, b_2 = 1 - \epsilon/2$. Then the participation order $\{1, 2\}$ yields a market price of $1 - \epsilon/2$, while the participation order $\{2, 1\}$ yields a market price of $\epsilon/2$. ■

In this example, not only do agent valuations change with ϵ , but the market in Proposition 2 is the smallest possible—only two agents. A skeptic could take the position that a large market with many agents will not suffer from these problems. Indeed, there is a fair amount of literature that suggests large, diverse prediction markets will be successful (e.g. [Surowiecki, 2004]), and on a more abstract level the concept of a “thick market” is associated with efficiency [Plott and Sunder, 1988, Yeh and Chen, 2001, Gan and Li, 2004].

In contrast to this intuition, our next result shows how prices in large markets can be uninformative. In particular, by changing the participation order of a set of agents, prices can converge to any arbitrary value. Every agent still acts in the market, all that changes is the order in which they do so. This is a surprising result that suggests conceptions of efficiency in large, dynamic markets do not carry a theoretical basis.

PROPOSITION 3. *There exists a countable set of agents with equal budgets such that, for all $\epsilon > 0$ and $x \in (0, 1)$, there exists a budget amount b , participation order o and finite T such that for all $t > T$, $|p_t(o) - x| \leq \epsilon$.*

PROOF. Without loss of generality, we force ϵ small enough so that $0 < x - \epsilon < x + \epsilon < 1$. If the given ϵ is too large, we simply satisfy a tighter bound.

Agents have valuations corresponding to unsimplified rational numbers; that is, for every ordered pair of positive integers $p < q$ there is an agent with value p/q . Now divide agents into three regions, as in Figure 1.

One region consists of all agents with values p/q such that $x - \epsilon/2 < x < x + \epsilon/2$. The other two regions are those agents with values above and below this region. We will refer to the three regions as the “upper”, “middle”, and “lower” regions. Note that each of the number of agents with valuations in each of these regions is countable, and as such there exists a mapping between agents with valuations in each region and the positive integers.

Endow every agent a budget of $\min\{c(x - \epsilon/2, x - \epsilon), c(x + \epsilon/2, x + \epsilon)\}$. Our ordering process is as follows: First, we will




Figure 1: The three regions used in the proof of Proposition 3.

take agents from the middle region until the current price is between $x - \epsilon/2$ and $x + \epsilon/2$, which will take a finite number of steps. Then, we will take agents in a specified order, ensuring both that every agent will visit the market and that prices remain within our bound.

By corollary 1, it takes only a finite amount of money to move the observed market price to any arbitrary value of $(0, 1)$. As a result, after a finite number of agents in the middle region participate, the observed market price will be between $x - \epsilon/2$ and $x + \epsilon/2$, the middle range.

Once prices enter the middle region, agents repeat the participation order described below *ad infinitum*:

1. The next agent from the upper region participates.
2. At least one agent in the middle region participates. Participation of further agents ceases when the current price is between $x - \epsilon/2$ and $x + \epsilon/2$.
3. The next agent from the lower region participates.
4. At least one agent in the middle region participates. Participation of further agents ceases when the current price is between $x - \epsilon/2$ and $x + \epsilon/2$.

Now provided that the initial starting price is between $x - \epsilon/2$ and $x + \epsilon/2$, the observed price after the participation of any agent will not be more than ϵ away from x . By construction of the budgets, a single agent will be able to take a price from $x - \epsilon/2$ to $x - \epsilon$, or $x + \epsilon/2$ to $x + \epsilon$. But after the participation of agents from the middle region, the observed price will move to within the middle region again.

Finally, every agent will participate in the market. In each cycle at least one agent from each region participates. Since the regions are countable, it follows that by repeating the process *ad infinitum*, every agent will eventually participate. ■

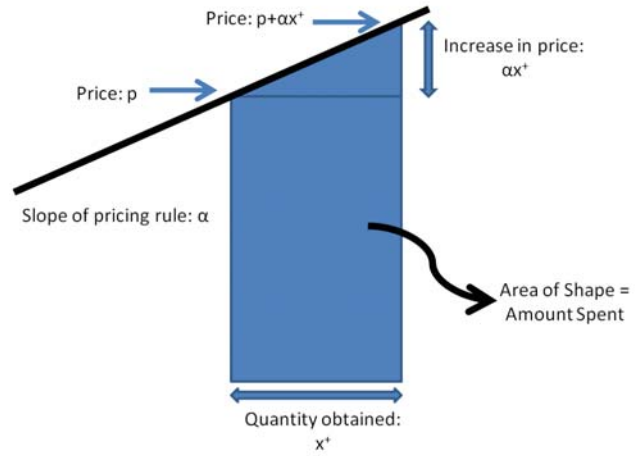


Figure 2: An agent with belief greater than p raises the price by αx^+ .

4. INFORMATIVE PRICES AND THE STANDARD STATIC MODEL

In this section, we show there is a sense in which large markets achieve a level of efficiency, so that their outcomes are not arbitrary with respect to the underlying distribution of agent beliefs.

Definition 2. In the ϵ -tweak process, agents interact with the market maker with uniform probability and spend $\epsilon \rightarrow 0$ dollars at each interaction.

PROPOSITION 4. Let agent valuations be distributed according to the continuous, gapless distribution $F(\cdot)$. With the use of any normal pricing rule, the unique fixed point of the ϵ -tweak process is the p^* that satisfies $p^* = 1 - F(p^*)$.

PROOF. Let us investigate the forces at play given the current price $p \in (0, 1)$. Since the pricing rule is differentiable, we can approximate the price function locally by a line with slope α , and since the pricing rule is strictly increasing, $\alpha > 0$.

As shown by figure 2, the result of an agent with belief greater than p interacting in the marketplace is to raise the price by αx^+ , where x^+ satisfies

$$x^+(p + \alpha x^+/2) = \epsilon$$

This solves to

$$\alpha x^+ = \sqrt{2\alpha\epsilon + p^2} - p$$

Now consider

$$\frac{\partial \alpha x^+(\epsilon, p)}{\partial \epsilon} = \frac{\alpha}{\sqrt{2\alpha\epsilon + p^2}}$$

And evaluating at $\epsilon = 0$

$$\frac{\partial \alpha x^+(0, p)}{\partial \epsilon} = \frac{\alpha}{p}$$

Now by similar, symmetric reasoning, the result of an agent with belief less than p interacting with the market lowers the price by αx^- , where

$$\alpha x^-((1 - p) + \alpha x^-/2) = \epsilon$$

And by repeating the above steps with x^- we have

$$\frac{\partial \alpha x^-(0, p)}{\partial \epsilon} = \frac{\alpha}{1 - p}$$

The expectation of the price movement in the ϵ -tweak process is the expectation of the result of an individual with a belief moving the market times the probability that individual is selected. Thus, the expectation of the price movement ϵ -tweak process is

$$\alpha \left(\frac{1 - F(p)}{p} - \frac{F(p)}{1 - p} \right)$$

The only p^* that makes this expression equal to 0 is $p^* = 1 - F(p^*)$. ■

The assumption that this is a large market is implied by the existence of a continuous valuation distribution, as well as the implication that there exists enough money invested in the market to move prices to the attractive fixed point. The significance of this result is that the equilibrium price of our dynamic model is identical to the unique equilibrium of the standard static model of prediction markets.

Definition 3. The *standard static model of prediction markets* involves a pool of risk-neutral traders with uniform wealth and a Walrasian market maker who sets a single price to balance supply and demand. The equilibrium price in this model is the p^* such that $p^* = 1 - F(p^*)$ [Manski, 2006].

To briefly explain the model of Manski [2006], a risk-neutral trader will always maximize their holdings, spending their entire endowment either for or against a contract depending on their private belief and the price offered. Letting traders all have unit wealth and setting price p , a trader with belief $b_i > p$ will buy $1/p$ shares, while a trader with belief $b_i < p$ will sell $1/(1 - p)$ shares. To balance supply and demand, then, the auctioneer sets the price p^* so that $p^* = 1 - F(p^*)$.

COROLLARY 1. *The fixed point of the ϵ -tweak process is identical to the price set by the Walrasian market maker in the standard static model of prediction markets.*

Because p can only be an equilibrium price if it represents the $1 - p$ -th percentile of beliefs, this fixed point is not necessarily as informative as we might hope. An equilibrium price of 50 cents could indicate a mean of beliefs anywhere between 25 and 75 cents. Moreover, the $1 - p$ -th percentile of beliefs is above the median (and, generally, the mean) for $p < .5$, and below the median (and mean) for $p > .5$. The result is a longshot bias, where unlikely events are overpriced and likely events are underpriced. Empirical studies of prediction market data have shown that prices are generally efficient but that there exists a small but statistically significant longshot bias. Wolfers and Zitzewitz [2006] demonstrate the bias on Iowa Electronic Market data on political elections, while Corwin and Othman [2008] demonstrate the bias on data from betting on professional basketball games.

4.1 What causes informative prices?

There are two differences between the ϵ -tweak model and the previous model of naïve behavior. The first is that agents participate in random order, and the second is that agents spend infinitesimal amounts when they participate. We may

ask which of these effects is more responsible for producing a good result. Put another way, imagine an adversary intent on making the result of our market uninformative gives us a choice between two options:

- He can order participation to his choosing, but agents spend a vanishingly small amount each time they participate.
- Agents participate randomly, but spend their entire budgets.

Which of these setups will most confound our adversary?

The answer is the second. In the first plan, our adversary could simulate the negative results of the previous section by having agents move again and again in small increments. On the other hand, the fixed point of having agents move in a random order but exhausting their entire budgets is identical to the one in which they spend a vanishingly small amount, with two differences:

1. The pricing rule is linear only for small investments, as an approximation facilitated by its differentiability. For larger investments it may curve, which will result in agents getting more (or less) quantity than would be expected from a linear approximation to the pricing rule.
2. For non-infinitesimal investments, a non-vanishing portion of the agent pool may reach their valuations before they have spent all their money in the market. As a result, they will not move the market as much, in aggregate, as they may be expected to under the ϵ -tweak process.

These effects destabilize the fixed point of the ϵ -tweak process we explored in this section, which was both unique and attractive. But these effects are second order in the true mathematical sense of the term. As a result, we can expect that the fixed point(s) of a random, full budget expenditure process will be close to $p^* = 1 - F(p^*)$. Though a full budget expenditure process may have a fixed point, the actual observed price at any instant could be quite far from this value. The deviation depends on the ability of an individual agent to move the market a significant amount by their participation. The closer agents get to not being able to greatly impact market prices with their individual expenditure, the closer we should expect the actual observed price to be close to the informative fixed point.

5. ORDERING PROBLEMS IN PRACTICE

We have demonstrated that the order in which agents participate can lead to market prices that are arbitrarily unrepresentative of underlying beliefs. But this does not mean we should expect to see unrepresentative markets in practice, particularly because Proposition 4 suggests that the random participation of large numbers of agents does produce meaningful prices. What our results do suggest is that prediction markets *can* become unhinged. In this section, we argue that the failure of the 1996 IEM vote-share market is consistent with agent ordering skewing the final outcome.

The market offered contracts in the two candidates, Clinton and Dole, with shares that paid out at the percentage of the two-party vote share each candidate received. Unfortunately, we only have access to daily summary prices; the




Figure 3: *Estimated margin of victory for Clinton over the last hundred days before the election in the 1996 IEM vote-share market. For reference, Clinton's actual margin of victory is given by the straight line. Time proceeds from left to right.*

full trading data from the markets is proprietary to the IEM researchers. Regarding the 1996 market, those researchers have said:

[I]n 1996, the market diverged from the correct outcome in the final days to close at midnight on the eve of the election with prices further from the election outcome than they had been since the Super Tuesday primaries in March. Apparently a large cash influx by new traders late in the campaign drove the price movement. [Berg et al., 2001]

This quote, along with the daily trading data, constitutes what we know about what happened in 1996.

Figure 3 plots the last hundred days of the market in terms of Clinton winning percentage, which is given by the Clinton share prices minus the Dole share price. As is evident from the plot, prices are volatile but relatively accurate until roughly 2-3 weeks before the election, at which point they begin a sharp, rapid increase. At the final trade the day before the election, the market's projected edge for Clinton was *nearly double* his actual victory margin, off by almost nine percentage points. This is an enormous error; the largest margin of victory in a modern presidential contest was less than 30 percentage points.

What is remarkable about this market is *when* the prediction began to sour. The end of the 1996 campaign was quiescent, with no major news stories or scandals breaking in the last few weeks of the contest. If we believe, like much of the theoretical literature contends, that market prices represent a shared posterior consensus among rational agents, prices should have settled, not diverged, at the end of the campaign. Theories of market efficiency would suggest exactly the opposite from what happened: that prices should perhaps be inaccurate months before the election, but as agents learn more from each other and information shocks play out, prices should converge towards an accurate prediction. Of course, when prediction markets work as well as they normally do, these theories are quite consistent with the observed data.

In contrast, agent ordering predicts prices to have the

greatest potential of being skewed at the end of trading. Consider a group of similar but extreme-minded market participants. How can these participants have the greatest impact on the final result of a market? (We ask this question being agnostic to whether these agents actually desire to manipulate final or intermediate market prices.)

The answer is by participating at the end of the market, after all the agents with different views have acted; the participation of the other agents drives prices as low as possible for the pool of similarly-minded agents. Consequently, our pool of traders with skewed beliefs effects the maximum amount of deviation from the market.

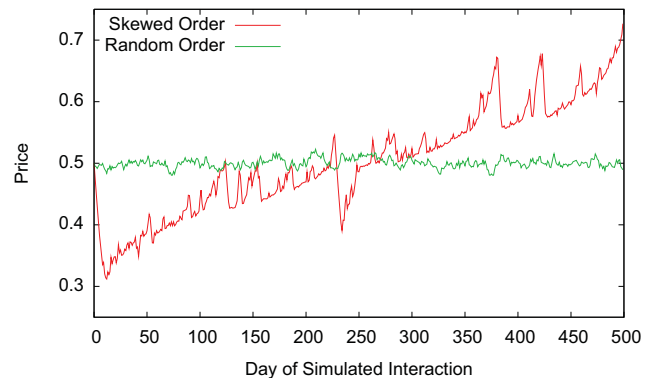


Figure 4: *A simulation of market prices shows the impacts of agent ordering. The green line is the simulated daily prices from random participation of agents with a true belief median and mean of 0.5. The red line is the simulated daily prices where agents with high beliefs are skewed to participate in the last hundred simulated days. Exactly the same agents participate in both cases, but the final price of the skewed ordering is much higher.*

Figure 4 is the results of simulating the daily prices of a market under two different ordering schemes. We simulate a daily price by taking a snapshot after the participation of 25 agents with small budgets in the market. The distribution of agent beliefs is given by a normal distribution with mean 0.5 and standard deviation 0.1. Both trading schemes consider an identical pool of agents, but in one case the order in which agents appear is skewed so that agents with higher beliefs are more likely to appear in the final hundred days of simulated trading. In particular, we probabilistically sorted agents so that for any two agents, the probability that the agent with higher belief participates after the agent with lower belief is 90%. The final prices in the skewed ordering are much higher than the final prices from random ordering.

The 1996 market is a phenomenon of natural science, and as such we cannot *prove*, in the theoretical sense of the other sections of this paper, why it failed. However, an agent ordering explanation is consistent with what happened in the 1996 market. Of course, our results rely on the notion of the market possessing an automated market maker, whereas the IEM is a traditional double auction. However, as Ostrovsky [2009] discusses, automated market makers are good approximations of traditional double auction market models, and many results from traditional models carry over qualitatively into automated market maker formats.

An agent ordering model suggests that most of the time,

we should expect markets to work well, but that, on occasion, we can expect sharp divergences from good outcomes at the close of contracts. The 1996 market provides a striking example of the latter case, but the wealth of data we have from other prediction markets—from sporting events and political elections—suggests that this kind of failure is very rare. Finally, the statement from IEM researchers, which constitutes our only information about the way agents actually behaved to develop the observed prices, is entirely consistent with agent ordering. On the other hand, a theories of market efficiency would suggest that savvy traders, particularly as prices diverged sharply from their previous (efficient) levels, would arbitrage away the actions of these new traders. Therefore, unlike more complex models of rational behavior, agent ordering provides a simple explanation consistent with the data as to why prices in the 1996 IEM diverged in the closing weeks away from their correct predictions.

6. CONCLUSION

Our work took a paradigm from computer science—worst-case and average-case analysis, and applied it to the study of pricing dynamics within prediction markets. By viewing agent actions as inputs and market prices as outputs, we achieved two distinct results. We showed that in the worst case, prices can be completely uninformative, converging to an arbitrary value. We then showed that in the average case, we can expect prices to be close—in a precise, mathematically rigorous way—to the informative equilibria of the standard static model of prediction markets.

Our model makes predictions that are supported by real-world studies. In particular, agent ordering provides a compelling explanation, consistent with the observed data, for both the failure of the 1996 IEM as well as for why that failure is so unique.

We feel that this work opens up new doors in the study of pricing dynamics. Models of markets often postulate looking only at equilibria, and indeed when these states are unique, attractive, and prescriptive such models can be compelling. However, the modeling compromises one makes in achieving these goals, such as ascribing perfect rationality to agents or completely ignoring dynamism, can be unrealistic. We know from traditional artificial intelligence that asynchronous settings can produce different emergent properties than synchronous settings (e.g., Huberman and Glance [1993]), and we know from theoretical computer science that there are significant differences between settings in which an adversary orders input and settings in which that input arrives randomly. This intuition proved correct when applied to our study of markets. The similarity between these results from the computer science literature and our findings suggests techniques from computer science will be fruitful in the future study of markets.

Acknowledgements

This work is supported by NSF grant IIS-0905390.

References

- J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a Dozen Years of Election Futures Markets Research. *Handbook of Experimental Economics Results*, 2001.
- Y. Chen, D. Reeves, D. Pennock, R. Hanson, L. Fortnow, and R. Gonen. Bluffing and Strategic Reticence in Prediction Markets. In *WINE*, 2007.
- I. Corwin and A. Othman. Time Inconsistency and Uncertainty Aversion in Prediction Markets. *Third workshop on prediction markets, part of ACM conference on Electronic commerce (EC2008)*, 2008.
- S. Dimitrov and R. Sami. Non-myopic strategies in prediction markets. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 200–209, 2008.
- J. Feigenbaum, L. Fortnow, D. Pennock, and R. Sami. Computation in a distributed information market. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 156–165, 2003.
- L. Gan and Q. Li. Efficiency of Thin and Thick Markets. *NBER Working Paper*, 2004.
- D. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, pages 119–137, 1993.
- A. Graefe and J. S. Armstrong. Can you beat the market? Accuracy of Individual and Group Post-Prediction Market Judgments. In *Third workshop on prediction markets, part of the ACM conference on Electronic Commerce (EC)*, 2008.
- B. A. Huberman and N. S. Glance. Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences*, 90(16):7716–7718, 1993.
- N. Lambert, J. Langford, J. Wortman, Y. Chen, D. Reeves, Y. Shoham, and D. Pennock. Self-Financed Wagering Mechanisms for Forecasting. In *Proceedings of the ninth ACM conference on Electronic Commerce (EC)*, pages 170–179, 2008.
- C. Manski. Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425–429, 2006.
- P. Milgrom and N. Stokey. Information, Trade and Common Knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.
- M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. In *Proceedings of the tenth ACM conference on Electronic Commerce (EC)*, pages 253–254, 2009.
- A. Othman. Zero-intelligence agents in prediction markets. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems (AAMAS)*, pages 879–886, 2008.
- C. Plott and S. Sunder. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*, 56(5):1085–1118, 1988.
- P. Shi, V. Conitzer, and M. Guo. Prediction Mechanisms That Do Not Incentivize Undesirable Actions. In *WINE*, 2009.
- V. Smith. Markets as Economizers of Information: Experimental Examination of the “Hayek Hypothesis”. *Economic Inquiry*, 20(2):165–79, 1982.
- J. Surowiecki. *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. Doubleday, 2004.
- J. Wolfers and E. Zitzewitz. Five Open Questions About Prediction Markets. Technical Report 1975, Institute for the Study of Labor (IZA), Feb. 2006.
- C. Yeh and S. Chen. Market Diversity and Market Efficiency: The Approach Based on Genetic Programming. *Artificial Intelligence and Simulation of Behaviour Journal*, 1(1):147–167, 2001.

