Zero-Intelligence Agents in Prediction Markets

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

ABSTRACT

We construct a novel agent-based model of prediction markets in which putative human qualities like learning, reasoning, and profit-seeking are absent. We show that the prices which emerge from a market populated by a class of distinctly inhuman agents, Zero-Intelligence agents with diffuse beliefs, replicate the findings of empirical market studies. We use this result to argue against the prevailing descriptive theories of price formation in prediction markets, which have stressed the role of expert, rational participants.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics, Theory

Keywords

Zero-Intelligence, Prediction Markets, Marginal Trader Hypothesis, Simulation

1. INTRODUCTION

Prediction markets represent a hybrid of asset markets and sportsbooks, wedding the mechanism of the former with the predictive power (and often dubious legality) of the latter. A prediction market serves to march buyers and sellers of a share that trades on a future event – for instance, the winner of an election, the film that will win the Oscar for Best Picture, or the next World Series champion. If the event in question occurs, the share expires at one dollar, if not, the share expires worthless. Prediction markets are small-scale, democratic, generally virtual exchanges. At nearly two decades old the most venerable prediction market, the Iowa Electronic Markets (IEM), operates under a no-action letter from the SEC which limits overall investments to 500 dollars per trader [7].

Because they involve real people who often make significant mistakes in trading, (c.f. [6, 7, 3]), prediction markets are the foremost testing ground for the so-called *Hayek Hypothesis*: that markets can work efficiently despite general ignorance on the part of participants in the trading environment [14]. More than just an online sportsbook, prediction markets represent a laboratory for theories of the efficiency of free markets as a whole.

Current research into how prediction markets reach their price predictions have centered around the *Marginal Trader Hypothesis* [6, 12], which attributes efficiency to a small pool of knowledgeable traders who are capable of setting prices and acting without bias. This hypothesis posits an "expert" approach to Hayek's theory: that select, savvy individuals drive markets toward efficiency. Remove these "perfect" individuals from the pool of traders and prediction markets would irretrievably lose their accuracy. In this paper we explore an alternative explanation of the Hayek hypothesis, a "general" approach, in which the action of the individual, imperfect trader drives price formation. We do this by combining two modeling frameworks, an intuitive private-value model of prediction markets first proposed by Manski [11], and a model of Zero-Intelligence agents originally used by Gode and Sunder [10] to study double auctions. By synthesizing these structures, we produce a model of prediction markets that does not rely on any "higher-order" human characteristics such as utility maximization, following trends, or learning. Previous studies of Zero-Intelligence agents have relied on computational simulation; we solve for a new way to calculate expected transaction prices mathematically. We show that the prices which emerge from our model closely resemble the results of empirical studies. We conclude that the pricing behavior seen in prediction markets can be produced without the presence of "expert-level" intelligence or reasoning, and state a case for abandoning the Marginal Trader Hypothesis.

2. ZERO-INTELLIGENCE AGENTS

The agents that we refer to as Zero-Intelligence (ZI) are derived from a concept first put forward by Gode and Sunder. In their description of a ZI agent, Gode and Sunder write that "it has no intelligence, does not seek or maximize profits, and does not observe, remember, or learn. It seems appropriate to label it as a zero-intelligence trader". Gode and Sunder used these computer agents to simulate market transactions in a double auction. They showed that a market consisting of particular type of agent, a Zero-Intelligence agent with a so-called "budget constraint", produced results that closely mirrored the allocative efficiency of a simultaneous experimental human exchange. As the authors concluded, "[the] primary cause of the high allocative efficiency of double auctions is the market discipline imposed

Cite as: Zero-Intelligence Agents in Prediction Markets, Abraham Othman, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.
Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

on traders; learning, intelligence, or profit motivation is not necessary".

Gode and Sunder's results, that efficiency can emerge from a mass of irrational actions, were provocative and have drawn scholarly attention from both the economic and computational sides of the agent-based community. In the fifteen years since their work was published, this attention has given the ZI concept notable baggage — criticisms that must be addressed in order to justify their use. We address two of the most potent critiques of ZI agents here. The first, from Cliff and Bruten [4] argues that, though ZI agents are irrational, Gode and Sunder's setting essentially guarantees convergence to prices that are close to efficiency. That is, the "books are cooked" to ensure a successful result for ZI agents. The second critique comes from Gjerstad and Shachat [9], who argue that ZI agents actually have non-zero intelligence, and that it is this intelligence which drives results towards efficiency.

2.1 The Mathematical Properties of ZI Agents

Cliff and Bruten use a dismissal of ZI agents as motivation for the creation of what they term "ZIP" (Zero-Intelligence Plus) agents, which combine features of Zero-Intelligence with basic learning mechanisms. Because of their simplicity and flexibility, ZIP agents have been widely adopted in situations in which adaptability is crucial, like evolutionary mechanisms or competitions (c.f. [5, 16, 2]).

ZIP agents represent an important addition to the literature independent of their motivation. At the same time, Cliff and Bruten's analysis of the probabilities associated with the behavior of ZI agents is flawed. After establishing that supply and demand curves provide a way by which to calculate the probability density functions (PDFs) of ask and bid orders, the authors write that "the PDF for transaction prices will be determined by the *intersection* of the PDFs of the offer prices and bid prices". The authors provide no further justification for this statement, but use the concept of intersecting density functions to solve for further properties of a ZI exchange. In turn, the authors use this analysis to argue that Gode and Sunder's setting was biased towards producing results close to efficiency.

By means of a worked counterexample, we show that the distribution of transactions is not given by the intersection method. Given a distribution over bid and ask orders, we show how to solve for the correct distribution of transactions, and we use this result to argue that Cliff and Bruten's more general point about ZI agents does not hold.

2.1.1 *A Worked Counterexample*

Imagine a large pool of an equal number of ZI buyers and sellers engaged in a continuous double exchange. Each agent, with value v_i, is called to trade with uniform probability. If a buyer, the agent places a bid order uniformly on $[0, v_i]$, and if a seller, the agent places an ask order uniformly on $[v_i, 1]$. A trade occurs when the highest-priced bid exceeds the lowest-priced ask at the price of the order which was placed first. As in Gode and Sunder, we restart the market after each transaction and buyers and sellers never switch roles.

Buyers and sellers each have three possible valuations in the market. Those valuations and the fraction of agents which possess them are given by the following table:

Buyers' Valuation	Fraction	Sellers' Valuation
1	4/7	0
.5	2/7	.5
.25	1/7	.75

Based on ZI agent behavior, these valuations induce the densities for bid and ask orders shown in Figure 1.

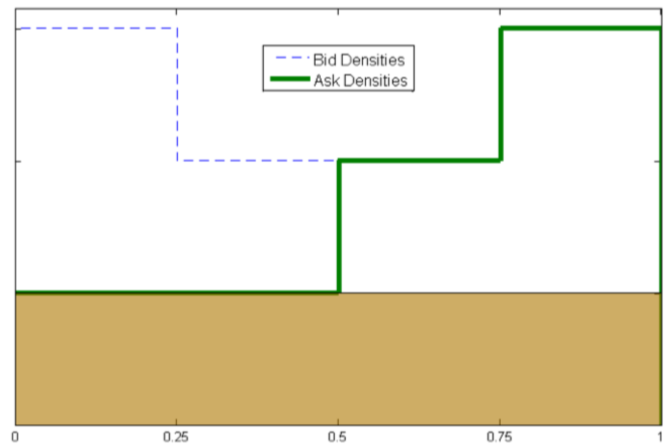

Figure 1: Probability density functions of Bid and Ask orders for the worked example. The shaded region represents the intersection of the two functions.

Were the argument in Cliff and Bruten correct, the density function for transactions would be proportional to the intersection of the two pictured densities — the shaded region, i.e. constant on $[0, 1]$. In fact, computational simulation of more than 10 million trades shows that transactions take place according to the following cumulative distribution function:

Input	Cliff and Bruten	Actual cdf
.25	.25	.2
.5	.5	.5
.75	.75	.8

As is evident, the distribution of transactions does not follow Cliff and Bruten's method.

2.1.2 *The Distribution of Transactions*

If transactions are not distributed according to the intersection of bid and ask densities, then how are they actually distributed? Let F_T, F_A and F_B represent the cumulative distribution functions over transactions, asks, and bids, respectively. By examining the results of millions of trades run on dozens of separate inputs, we were able to determine that:

$$F_T = \frac{F_A}{F_A + 1 - F_B} \tag{1}$$

Figure 2 provides a two-dimensional map of the transaction price cdf. As the plot evinces, the simple functional form of the distribution belies the complex relation between bids, asks, and transactions. The observant reader will note the plot contains a "hole" in the upper-lefthand corner, where $F_A = 0$ and $F_B = 1$. This reflects the intuition that if the highest bid order is placed lower than the lowest ask order, no transactions will occur at all.

We can gain further intuition into Equation 1 by taking its derivative, producing the probability density function for

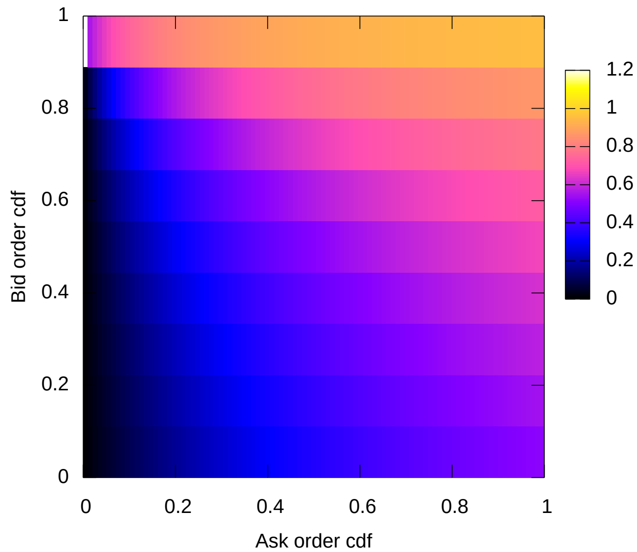

Figure 2: Heat map of Equation 1.

transactions. Let f_A, f_B, and f_T represent the probability density functions for asks, bids, and transactions respectively. By the quotient rule:

$$
\begin{aligned}
f_T &= \frac{f_A(F_A + 1 - F_B) - F_A(f_A - f_B)}{(F_A + 1 - F_B)^2} \\
&= \frac{f_B F_A + f_A(1 - F_B)}{(F_A + 1 - F_B)^2}
\end{aligned}
$$

The numerator here represents the probability density of a transaction taking place at a value - either a bid is placed at that value and is crossed by a lower ask (i.e. $f_B F_A$), or an ask is placed at that value and is crossed by a higher bid (i.e. $f_A(1 - F_B)$).

2.1.3 Implications

As Cliff and Bruten demonstrated in several examples, mean transaction prices in a ZI market can diverge sharply from efficiency. Their argument is that Gode and Sunder's markets demonstrated prices close to efficiency by design; they write that: "[Q]ualitatively at least, it would appear that near-equilibrium transaction prices are expected because of the shape of the PDF for valid deals: this [Zero-Intelligence] system is structured *a priori* to generate mean transaction prices close to the theoretical equilibrium price".

Our work shows that this argument can be rejected. The "shape of the PDF for valid deals", that is, the density of transactions, is much more complex than Cliff and Bruten had posited. Consequently, it is difficult to assert that a ZI system is structured *a priori* to arrive at any specific conclusion. In Section 3.4 we give the equations involved in the calculation of expected transaction prices in our prediction market model; it is difficult to arrive at any conclusion regarding the efficiency of those prices without numerical inquiry.

2.2 The Fallacy of "Budget Constraint"

We now turn to a second distinct critique of Zero-Intelligence agents. Gjerstad and Shachat have recently argued that ZI agents do not quite square with their moniker.

In particular, Gode and Sunder study two types of Zero-Intelligence agents. The first are *unconstrained* agents, which place orders arbitrarily, regardless of private valuations. The results of simulations with these agents do not at all resemble the results of human simulations - prices fluctuate wildly and without response to different supply and demand curves. The second type of agents are *budget constrained*: they are not allowed to place an order with directly negative proceeds, that is, to sell below cost or buy above value. These are the agents which produced allocative efficiency on par with human experiments and were sensitive to shifts in supply and demand curves.

Gode and Sunder fold this budget constraint into their market mechanism — it is the market's imposition that traders be able to settle accounts that forbids agents from making a trade which will run them a loss. The authors do not attribute this property to the agents themselves, and they write "The difference between the performance of the human markets and that of the ZI-C (i.e. budget constrained agents) is attributable to systematic characteristics of human traders ... Traders have no intelligence in either the ZI-U (i.e. the unconstrained agents) or the ZI-C market; the ZI-C market prevents the traders from engaging in transactions that they cannot settle".

Yet as Gjerstad and Shachat show, the behavior of "budget constrained" agents is ontically identical to the behavior of agents which display individual rationality. An observer watching the behavior of a budget-constrained agent would be unable to assert that the agent was not individually rational. The budget-constrained agents do not place any orders which directly result in a negative payoff, regardless of whether they have "internalized" this constraint or whether the market "imposes" this constraint upon them. When all the structures in question, both in agent behavior and in market restrictions, are not physical but are instead implemented in software, this distinction becomes quite literally immaterial.

It was the imposition of the budget constraint that made ZI agents produce modeling results that were not just noise. But this budget constraint is, for all intents and purposes, equivalent to individual rationality. Thus, it can be regarded as an imposition of intelligence on the agents. ZI agents purport to be a model devoid of human qualities. Are ZI agents still useful as a modeling tool?

We argue that they are. The "intelligence" given to these agents is of a very prescribed sort, so restricted as to not compromise them as being fundamentally different from human intelligence. No one could confuse the behavior of a ZI agent with that of a human. Moreover, an important quality of ZI modeling is that we should expect human agents to only act in ways that are restrictions of the behavior of ZI agents — for instance, maximizing profits, or learning and tracking prices before participating. This is not the case for more descriptive agent-based models — consider a market model (as in [15]) in which agents are classified as chartists (trend-followers) or fundamentalists. It is easy to suggest plausible human behavior that this model does not reflect — for instance, people which alternate between being fundamentalists and chartists, or traders employing charting strategies which are not among the choices presented to agents. In contrast, the restriction that ZI agents perform no action which directly results in a negative payoff does not limit our understanding of the way people behave.

3. MODELING PREDICTION MARKETS

Gode and Sunder's work involved Zero-Intelligence agents participating in double exchanges. Though prediction markets are also double exchanges, different notions of "correct" prices are at play in the two modeling literatures. In a traditional double exchange, the equilibrium price is given by the intersection of the supply and demand schedules. In a prediction market, the objective price is thought of as the mean of the distribution of agent beliefs over the probability of the event occurring.

Our objective of this section is to build a ZI model that can be used to generate price estimates. This model allows us to answer questions like "If the objective probability of an event occurring is 25 percent, what price should we expect in a ZI market?". We begin by giving an overview of prediction market modeling, and then introduce our own ZI model.

In real prediction markets, the objective probability of an event that trades at a price can be estimated by observing many independent trials and averaging. In our ZI simulation, the objective probability is given by the mean of the distribution of agent beliefs. This makes the choice of belief distributions a crucial component of the model. We select a particular set of beta distributions as our belief distributions and justify our choice. Finally, we conclude our model by showing how to solve numerically for the expected transaction prices in our ZI model.

3.1 Overview of Previous Models

In Manski's model of prediction markets, each agent possesses a private belief that an event will occur. Prices are taken as exogenous and agents are price-takers, making investment decisions to maximize utility based on those exogenous prices and their individual beliefs over what the future holds. For example, a risk-neutral agent with a belief that an event will occur with probability one-half will wager their entire endowment if the price of such an asset is 45 cents.

Manski's model suggested an inherent structural inefficiency in prediction markets: equilibrium prices may not align with their objective probabilities, which he took to be the mean of the distribution of beliefs. In particular, Manski suggested that prediction markets were susceptible to the *longshot bias*, wherein unlikely events (longshots) are overpriced, and likely events (favorites) are underpriced. Further studies, by Gjerstad [8] and Wolfers and Zitzewitz [18], showed that the longshot bias was a function of the risk-neutrality of the participants. As participants became more risk-averse, prices as a whole would become closer to objective probabilities. When participants had logarithmic utility, prices would align precisely with objective probabilities, and when traders were even more risk averse, prices would demonstrate a *reverse* longshot bias — precisely the opposite effect.

The longshot bias in prediction markets was demonstrated empirically by Wolfers and Zitzewitz [17] by compiling years of data from the Iowa Electronic Markets and determining the relation between observed prices and expected payouts. The authors tie their finding into the broad history of economic work which has studied the longshot bias in sportsbooks, dating back to Ali's 1977 study of favorites and longshots at the racetrack [1].

From the point of view of a community interested in agents and learning, however, this strictly economic work is not particularly insightful. The studies do not yield any insight into the role of learning within markets and contain an intrinsic contradiction in the way agents use their intelligence. On one hand, the private-value model places great weight in the role of actors to be perfectly rational decision makers — the ability to take optimal, utility-maximizing action when faced with a price and a probability. On the other hand, those same agents myopically assume that *their* valuation is the correct one, even if market prices do not reflect this belief.

One way to resolve this tension is to replace a private-value model with a common-value model, where the value an agent has for the share of an event is dependent on the private signals received by all agents. Such a model was recently discussed by Serrano-Padial [13] and looks promising, but complex. Our own model resolves this tension in the opposite way, by abandoning the notion of utility maximization by agents.

Though our private-value model is adapted from previous work, our model differs in one very important respect, that of the price formation mechanism. Because our model is based around the behavior of agents, prices no longer need to be set exogenously — they are formed through the direct actions of individuals. Previous, non-agent-based models have relied on a Walrasian auctioneer or market-maker to adjust prices to equilibrium levels. We feel this is a particularly stringent requirement to impose on a model of prediction markets, which are less structured than traditional markets and generally do not have agents designated to determine prices.

3.2 A ZI Model

As we concluded in Section 2.2, the only intelligence associated with ZI agents is that they cannot place an order with directly negative consequences. In all other matters, Gode and Sunder's agents acted on the principle of uniformity - given a choice between alternatives, they select arbitrarily with equal probability to all actions. Additionally, agents should approach the market as a blank slate: with no knowledge of their past actions, no recognition of present conditions, and no expectation over future outcomes. ZI agents cannot learn, because they have no history. Given these three guidelines, we can now create our own ZI agents to function in a prediction market setting.

There exist a pool of homogenous agents, each of which has a belief $b_i \in [0, 1]$ that an event will occur. These beliefs are drawn independently from a distribution with density function f. Agents are called to market with uniform probability, where they decide to either place a buy (bid) or sell (ask) order with equal probability. The agent then places a bid order for one share of the event at a price $\sim [0, b_i]$, or an ask order for one share at a price $\sim [b_i, 1]$. Agents do not observe current market prices, and do not react to any result of their actions; they keep no record of previous orders or if those orders resulted in trades.

3.3 Belief Distributions

As both Gode and Sunder's as well as Manski's studies relied fundamentally on the distribution of agent valuations, it should be no surprise that our study also hinges on the distribution of private information, in our case, individual beliefs. We produce four requirements for our belief distributions, and show how we can use families of Beta distributions to meet these requirements.

For a given belief distribution over the probability that an event will occur, we follow Manski's convention that the objective probability (and, therefore, the objective trading price for risk-neutral agents) that the event occurs is given by the expectation of that distribution. Our objective is to produce a linkage between these objective probabilities and the prices that emerge in a market of ZI agents with the specified beliefs, so we will need to calculate expected prices using a multiplicity of distributions, each with a different mean. We should thus favor distributions that are simple and fast to work with numerically. Additionally, even though the distributions we use with different means must necessarily be different, they should ideally be related in an intrinsic way.

We can add two further desiderata. Because events cannot occur with probability less than zero or greater than one, our belief distributions should be double-bounded. Finally, our distributional choice should reflect the inherent symmetry of prediction market shares — that a market for an event occurring with probability p is identical to a market for an event not occurring with probability $1-p$. To formalize this requirement, let f_z represent the probability density function of our distribution with mean z. The symmetry we require is that:

$$f_z(x) = f_{1-z}(1-x)$$

Specific families of Beta distributions provide these properties for us. Recall that Beta distributions rely on two parameters, α and β, and have mean $\alpha/(\alpha+\beta)$. Beta distributions can also be parameterized by the sum Σ and mean μ, such that $\Sigma = \alpha+\beta$ and $\mu = \alpha/(\alpha+\beta)$. Our scheme for generating distributions of a specified mean is to use the beta distribution that has the specified mean and whose parameters sum to a constant value.

Beta distributions are common and are implemented well in major software packages. They possess our symmetric property, and they are double-bounded. It is immediate that they meet all of our desiderata except for the "innately related" quality.

Why should we group together Beta distributions whose parameters have the same sum? To answer this, consider that the the Beta distribution is the conjugate prior for observing Bernoulli trials. Imagine that we have current beliefs that the probability an event occurs is uniform on $[0,1]$. This is the Beta distribution with parameters $\alpha = 1$, $\beta = 1$. After observing N Bernoulli trials and seeing s successes and $N-s$ failures, our beliefs over the likelihood of the event are given by a Beta distribution with parameters $\alpha = 1+s$, $\beta = 1+N-s$. Note that regardless of the realized numbers of successes, the resulting Beta posterior distribution will be given by parameters that always sum to the same value. Beta distributions with the same parameter sum therefore can be thought of as representing responses to the same amount of information. As the sum gets larger, the distributions become more specific.

We refer to families of belief distributions with the same sum as *belief regimes*, and we explore three different regimes in our computational study:

Belief Regime	Sum of Parameters
Low Accuracy	5
Middle Accuracy	25
High Accuracy	100

Some graphical intuition for these regimes is given in Figure 3, which depicts the distributions of the different belief regimes at two means. As the figure shows, the low accuracy belief regime is diffuse, while the high accuracy belief regime is quite tight around its mean. Note that we are not suggesting that human agents in real markets have belief distributions corresponding to one of the regimes, instead, we are using them as devices to gain perspective on what happens to market prices under different belief distributions.

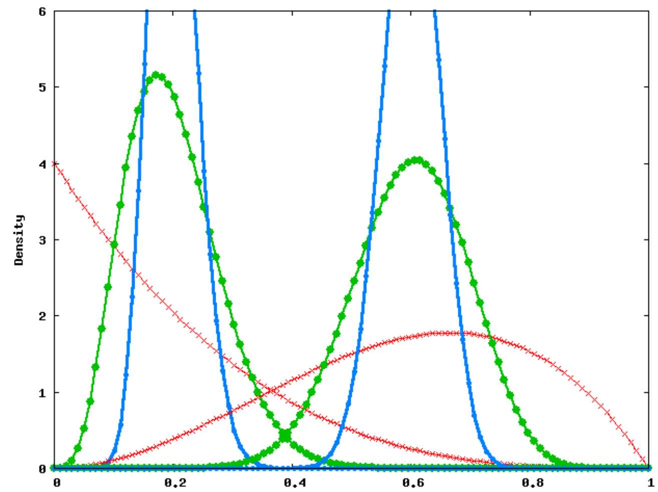

Figure 3: Illustration of the three different belief regimes discussed in the text, at distributional means .2 and .6.

3.4 Solving for Expected Transaction Prices

Previous studies of ZI agents have generally involved computational simulation. With our increased understanding of how transactions are distributed, from Equation 1, we can solve for expected prices accurately without simulating a market in software and running thousands of trials. As a result, we are able to calculate equilibria for many more configurations than in previous studies.

Note that Equation 1 depends on only two terms: $F_A(x)$ and $1-F_B(x)$. The latter of these is the probability of a bid order being placed at a price greater than x, and the former is the probability of placing an ask order at a price less than x. By the continuous analog to exhaustive partitioning:

$$F_A(x) = \int_0^x \left(\frac{x-b}{1-b}\right) f(b)\, db$$

$$1 - F_B(x) = \int_x^1 \left(\frac{b-x}{b}\right) f(b)\, db$$

If the belief distribution is relatively smooth then numerical quadrature methods will have no problem evaluating these integrals. Finally, we can calculate the expected transaction price using the tail-sum theorem:

$$\mathbb{E}(T) = \int_0^1 1 - F_T(x)\, dx = \int_0^1 \frac{1-F_B(x)}{F_A(x)+1-F_B(x)}\, dx$$

Once again, numerical quadrature will be able to produce solutions to this second level of integral. The fact that we are chaining together two levels of numerical integration should not be alarming, because the beta distributions we use are suitably well-behaved.

4. RESULTS

Figure 4 (on the next page) shows the results of numerically calculating the expected transaction price under the three different belief regimes. All three regimes show the longshot bias. Consider, for instance, when the distributional mean (which, we argue, is the objective probability of the event) is .8, the expected prices for the high, middle, and low accuracy belief regimes are .781, .766, and .742 respectively. Similarly, when the objective probability of an event is .2, prices will be .219, .234, and .258. Likely events are underpriced, while unlikely events are overpriced.

It is intuitive that in the limit of accuracy ZI markets would be unbiased: If all agents had identical b_i, the only price at which transactions could occur would be that value. In Section 3.4, we solved for the equations corresponding to the expected transaction prices. Given those complex equations, it is not immediately evident that the amount of bias in prices should increase with the diffuseness of beliefs. Furthermore, it is not obvious that any difference between objective prices and expected transaction prices should be in the direction implied by longshot bias.

Figure 5 depicts the longshot bias as found by Wolfers and Zitzewitz [17] in their study of prediction market data from the IEM. To facilitate comparisons with our computational results, we have structured our plot on equivalent axes. Observe the significant longshot bias for shares priced around 25 cents, which virtually never pay off, and shares priced around 75 cents, which pay off with near certainty. By comparing the two plots, it is clear that the expected transaction price of ZI traders under the low accuracy belief regime closely matches the observed data from real prediction markets populated by human traders.

5. CONCLUSIONS

Agent-based modeling of the real world is necessarily dubious. Attempting to model the rich tapestry of human behavior within economic structures — both the outstandingly bad and the terrifically complex — is a futile task. Worse, from the perspective of the modeler, is that descriptive agent-based models of human behavior are perpetually subject to the trenchant complaint: "That is not the way people behave".

Zero-Intelligence agents provide a way out of this trap. The power of ZI agents is not that they serve to provide an accurate or essentialist model of human behavior; they do not. Instead, they play precisely the opposite role — ZI agents represent players that are distinctly non-human. The insight from ZI agents comes not from the results of their simulation, but rather from the difference between those results and the results of actual, human, markets. This difference is a product of what is not in the ZI model: distinctly human qualities — of reasoning, learning, strategizing.

In the case of prediction markets, our ZI model of agents with diffuse beliefs produces results that closely match those found in real markets. We can think of three explanations for this phenomenon:

- The belief distributions of participants in a real market are actually more diffuse than in our low accuracy belief regime, indicating that human qualities drive markets further towards efficiency.
- The belief distributions of participants in a real market are approximated well by our low accuracy belief regime, and so human qualities do not play a significant role in price formation.
- The belief distributions of participants in a real market are more exact than those of the low accuracy belief regime. Thus, human qualities actually drive prices *away* from efficiency.

Of these options, we feel that the first can be discarded. The low accuracy belief regime was selected because it would be difficult to imagine actual beliefs being any more diffuse. As Figure 3 depicts, the beliefs in the low accuracy model are quite spread out already, with significant density far away from the mean. Furthermore, if human qualities like reasoning or learning are driving prices towards efficiency, then why should they stop at an inefficient result? The fact that empirical results from prediction markets shows the longshot bias is a strong argument against this explanation.

Concluding that the wealth of qualities that separate human traders from ZI agents either have no effect or drive markets towards inefficiency is quite provocative in and of itself, but because of previous studies of prediction markets it is almost heretical. The team of researchers responsible for the Iowa Electronic Markets first put forth the Marginal Trader Hypothesis in the context of the market run on the 1988 presidential election [6]. That market, even though it was quite small, was a terrific success: it predicted the outcome of the election remarkably well.

Yet it is important to note that the 1988 market operated on rules that were fundamentally different from the rules that the prediction markets of today operate upon: that market was a *vote-share market*, where shares expire at the share of the popular vote a candidate receives, rather than a *winner-take-all* market in which shares of the winning candidate expire at a dollar and losing shares expire worthless. Because they have universal interpretation outside of elections, modern prediction markets are run virtually exclusively as winner-take-all markets.

The difference in the outcome (and the interpretation of that outcome) between the two market settings is profound. In a winner-take-all market of the 1988 election, shares of Michael Dukakis, the losing candidate, would have expired worthless, but in the vote-share market they expired at more than 45 cents on the dollar. Had the 1988 market been winner-take-all, it would have only been a couple data points — of a share trading at a price expiring at a dollar, for instance. Instead, in a vote-share market, it became an instance of a prediction market correctly guessing the result of the popular vote far better than polls. This result prompted the search for an explanation which culminated in the development of the Marginal Trader Hypothesis.

Once one accepts that prediction markets are efficient and that marginal traders are the cause of that efficiency, it is easy to construct narratives that make the pair mutually reinforcing. For instance: prediction markets are efficient because the desire to profit motivates better-informed agents to take advantage of the misjudgments of less-informed agents. By the time enough data about prediction markets in the winner-take-all context could be gathered to draw conclusions, this narrative of efficiency had already been ensconced for more than a decade. But the greatest proponent for a hypothesis should not be inertia.

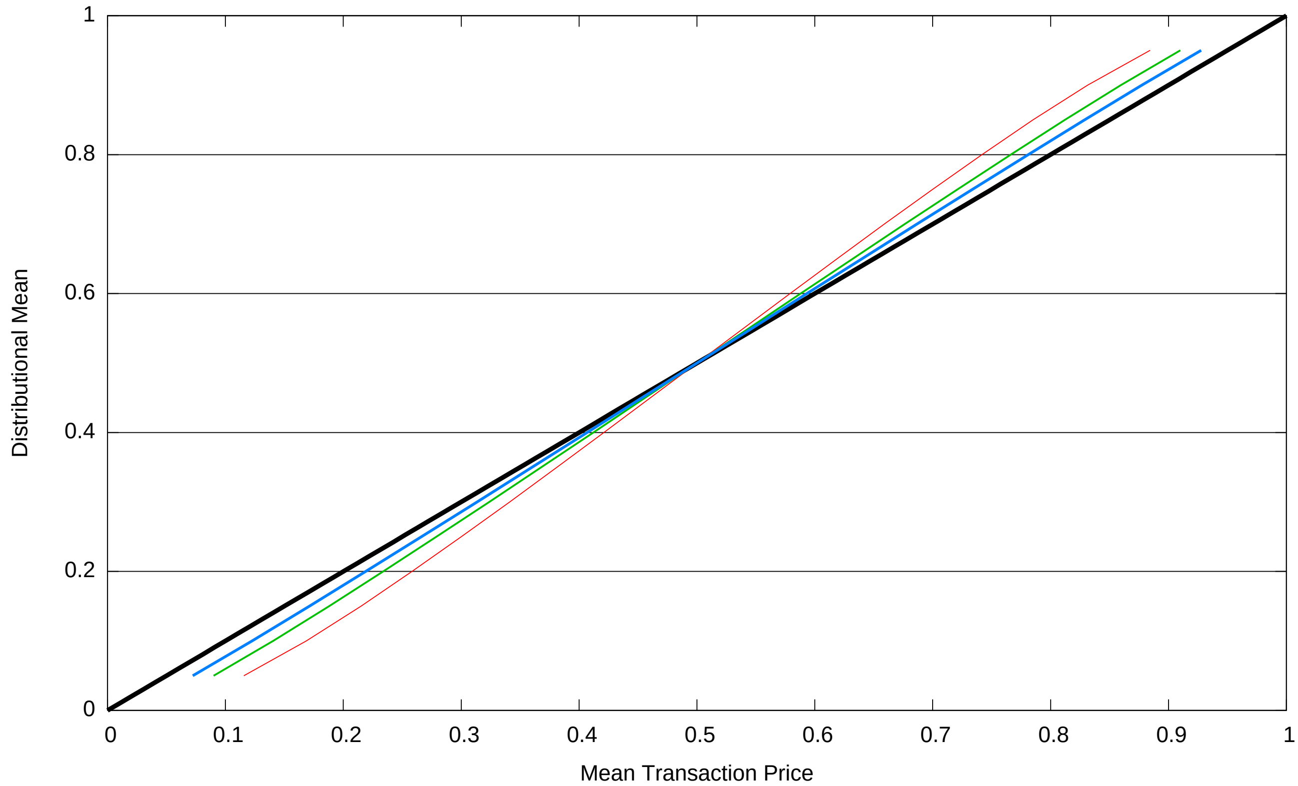

Figure 4: Expectation of transaction prices in a ZI market under different belief regimes. From left to right across the top right of the plot, the lines show mean transaction prices under a low, middle, and high accuracy belief regimes, respectively. The fourth, bold, line represents perfect accuracy.

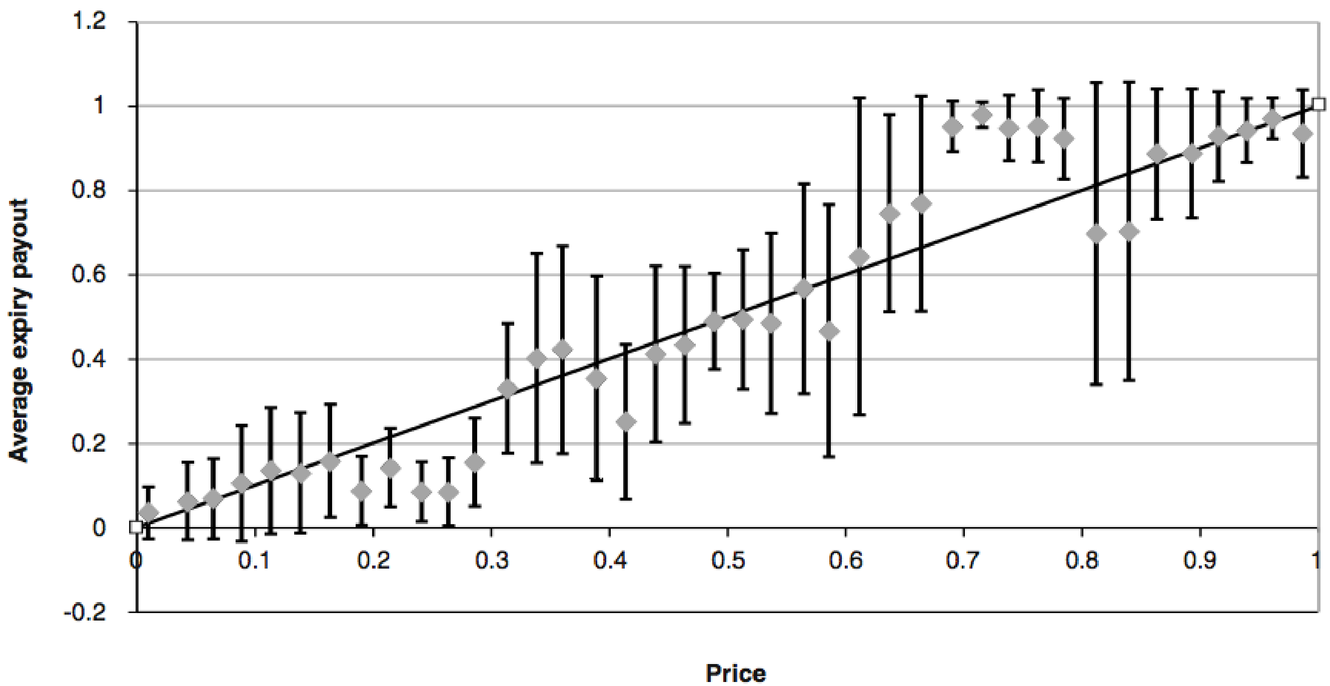

Figure 5: The longshot bias in real markets. Gray circles represent the average payout from a share trading at a specific price, and black bars are 95th percentile confidence intervals. Taken from Wolfers and Zitzewitz [17].

5.1 Future Directions

We base our conclusions regarding the inefficiency on the results of only a single study showing the bias in human prediction markets. Wolfers and Zitzewitz analyzed more than a dozen years of results from a well-respected prediction market, and grounded their finding of inefficiency in the much longer history of the longshot bias appearing in racetrack betting and sportsbooks. It would be interesting, however, to explore the longshot bias in other prediction markets. InTrade, a Dublin-based exchange, operates political, sports, and financial prediction markets. Do these data sets also display the longshot bias? More or less, and why?

Another promising direction is to increase the level of agent sophistication within the simulation. What kind of transaction prices are observed in a market populated by the simple learning ZIP agents? These more sophisticated agents can also be used to model the time dimension of prediction markets. Perhaps news shocks occur over time and are received in different strengths by different agents. Will a market with these simple learners over- or under-react to these new pieces of knowledge?

6. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under ITR grant IIS-0427858.

7. REFERENCES

[1] M. Ali. Probability and Utility Estimates for Racetrack Bettors. *The Journal of Political Economy*, 85(4):803–815, 1977.

[2] A. Bagnall and I. Toft. Zero Intelligence Plus and Gjerstad-Dickhaut Agents for Sealed Bid Auctions. *Workshop on Trading Agent Design and Analysis, part of international conference on autonomous agents and multiagent systems (AAMAS-2004)*, pages 59–64, 2004.

[3] J. Berg and T. Rietz. The Iowa Electronic Markets: Stylized Facts and Open Issues. *Information Markets: A New Way of Making Decisions, edited by Robert W. Hahn and Paul C. Tetlock. AEI-Brookings Joint Center for Regulatory Studies*, 2006.

[4] D. Cliff and J. Bruten. Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Technical Report HPL-97-141, Hewlett-Packard Laboratories Bristol, 1997.

[5] R. Das, J. Hanson, J. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI-2001)*, pages 4–10, 2001.

[6] R. Forsythe, F. Nelson, G. Neumann, and J. Wright. Anatomy of an Experimental Political Stock Market. *The American Economic Review*, 82(5):1142–1161, 1992.

[7] R. Forsythe, T. Rietz, and T. Ross. Wishes, expectations and actions: a survey on price formation in election stock markets. *Journal of Economic Behavior and Organization*, 39(1):83–110, 1999.

[8] S. Gjerstad. Risk Aversion, Beliefs, and Prediction Market Equilibrium. Technical report, Economic Science Laboratory, University of Arizona, 2005.

[9] S. Gjerstad and J. M. Shachat. Individual rationality and market efficiency. Purdue University Economics Working Papers 1204, Purdue University, Department of Economics, Aug. 2007.

[10] 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*, 101(1):119–37, 1993.

[11] C. F. Manski. Interpreting the predictions of prediction markets. NBER Working Papers 10359, National Bureau of Economic Research, Inc, Mar. 2004.

[12] D. M. Pennock and R. Sami. Computational Aspects of Prediction Markets. In *Algorithmic Game Theory*, chapter 26, pages 651–674. Cambridge University Press, 2007.

[13] R. Serrano-Padial. Strategic Foundations of Prediction Markets and the Efficient Markets Hypothesis. *Second workshop on prediction markets, part of ACM conference on Electronic commerce (EC2007)*, 2007.

[14] V. Smith. Markets as Economizers of Information: Experimental Examination of the "Hayek Hypothesis". *Economic Inquiry*, 20(2):165–79, 1982.

[15] K. Steiglitz and L. O'Calligan. Microsimulation of markets and endogenous price bubbles. *Third International Conference on Computing in Economics and Finance*, 1997.

[16] V. Walia, A. Byde, and D. Cliff. Evolving market design in zero-intelligence trader markets. *IEEE International Conference on E-Commerce (CEC 2003)*, pages 157–164, 2003.

[17] J. Wolfers and E. Zitzewitz. Five open questions about prediction markets. NBER Working Papers 12060, National Bureau of Economic Research, Inc, Feb. 2006.

[18] J. Wolfers and E. Zitzewitz. Interpreting prediction market prices as probabilities. NBER Working Papers 12200, National Bureau of Economic Research, Inc, May 2006.