Time Inconsistency and Uncertainty Aversion in Prediction Markets*

Ivan Corwin† and Abraham Othman‡

May 7, 2008

Abstract

Starting from first principles we derive a method for detecting price biases in prediction market data. Using this method on Tradesports contracts from the 2005-06 NBA season, we demonstrate that trades executed in the last hour of trading have a significant longshot price bias, while trades occurring earlier do not. We present a new theoretical model which uses uncertainty aversion to explain our findings.

1 Introduction

Models of price formation in prediction markets find their source in Ali's work examining the payouts of track betting [2]. That work was the first that demonstrated the *longshot bias*, in which likely events are underpriced and unlikely events are overpriced. The longshot bias has been demonstrated in many other betting markets, including soccer [4] and baseball [16].

In prediction markets, the longshot bias was first demonstrated theoretically by Manski [8]. In Manski's model, agents are risk-neutral price takers with independent beliefs that an event will occur. Taking the objective probability of the event as the mean of the distribution of beliefs, Manski's result suggested that equilibrium prices would exhibit the longshot bias. Further independent studies by Gjerstad [7] and by Wolfers and Zitzewitz [15] showed that Manski's results were a function of the risk-neutrality of the participating agents. As participating agents become more risk averse, equilibrium prices resemble objective probabilities, and for extreme

*This work is supported by the National Science Foundation under the Graduate Research Fellowship Program and ITR grant IIS-0427858. We thank John Delaney for the contract data.

†Department of Mathematics, Courant Institute of Mathematical Sciences, New York University. corwin@cims.nyu.edu.

‡Computer Science Department, Carnegie Mellon University. aothman@cs.cmu.edu.

values of risk aversion prices exhibit the *reverse* longshot bias, in which likely events are overpriced and unlikely events are underpriced. With regards to empirical studies, Wolfers and Zitzewitz provided the first study of the longshot bias within prediction markets, demonstrating the effect in the venerable Iowa Electronic Markets [14].

More recently, Borghesi [3] examined contracts from Tradesports to study how markets react to information shocks (touchdowns in NFL games). Note that unlike traditional wagering systems, Tradesports does not stop trading once an event begins. As a result, Tradesports datasets provide a new and unique perspective on how prices react to the filtration of information into the pool of traders.

There are many questions influencing the outcome of a game which over the course of play are answered, enriching the set of information accessible to traders. Will key players struggle? How will nagging injuries impact the outcome? Will one team's strategy work particularly well or poorly? With all of these questions undecided before a game one would imagine that market prices become more accurate predictors of game results as the game unfolds and the answers to these questions are revealed. In that sense our results are counter-intuitive: we show that prices actually become *less accurate* as the game proceeds. In the last hour of play prices reflect a longshot bias which was nonexistent earlier in and before the game.

In our empirical analysis we develop a robust method to analyze trading data with pronounced cluster correlation which is simpler than the variance estimation methods introduced by Woodruff [17], reintroduced by Williams [13] and used by Borghesi.

Following our surprising empirical results we describe and justify a theoretical model for trade prices based on uncertainty aversion. We argue that the multitude of uncertainties present before and in the early stages of the game serve to discourage participation, even by risk-neutral traders. We show that uncertainty aversion works to counter the longshot bias, but as the game unfolds and uncertainty aversion diminishes, the longshot bias emerges as observed.

2 Data Interpretation and Results

Our data is in the form of a list of 16177 time-stamped trades from 129 different games. Trades represent an exchange being made at a price $p \in \{.5, 1, 1.5, \ldots, 99.5\}$, where shares pay off at

100 if the event occurs, and 0 if it does not.

To provide motivation and to ground our study, we seek to answer the following hypothetical: Imagine that we arrive at the market and see a transaction occurring at price p. What is the probability that the event actually occurs?

We begin by making two simplifying assumptions. First, note that if a contract for team A to beat B is trading at price p, it is the same as if the contract for team B to beat A were trading at price $1 - p$. Observing this symmetry around a price of 50, we collapse our data into a range $\{50, \ldots, 99.5\}$ — if we observe a trade happening below 50, we simply reflect the trade into our range. A reflected trade pays out the opposite of an unreflected trade. It will be useful to define the *parity* of a trade to be 1 if the trade price is greater than 50, -1 if less than 50 and 0 if equal to 50. An important implicit assumption we make here is that the contracts, as created by Tradesports, are not meaningfully directed towards either team. Secondly, we simplify our analysis by bucketing trades into intervals of 2 cents. That is, trades with similar prices are grouped together. This is perhaps more controversial than our first simplification because the actual buckets are relatively small in size. Tetlock [12] uses large buckets with 20-cent spreads, while Wolfers and Zitzewitz, like us, use smaller buckets. We feel strongly that smaller buckets are more meaningful because the size of the bias is much larger relative to the size of the buckets. After reflecting prices through 50 cents and bucketing, our price space looks like $\{[50, 52), [52, 54), \ldots, [98, 100)\}$, where $[52, 54)$ means that a price of 52 would be thought of as being in the $[52, 54)$ interval but a price of 54 would not.

Now we return to our motivating question. To best predict the results of a game based on a single observed price, we should look at all trades made at that price and determine the empirical probability of the contract paying out. Though this method should limit to the correct probability, it is not well suited for our data. Splitting the 129 contracts traded into quartiles based on the number of trades per contract we quickly notice a strong popularity bias. Table 1 shows that the 12 most popular contracts, as well as the 67 least popular contracts each comprised a quarter of the total trades. Since we are dealing with binary events, this popularity bias means that if we treat every trade equally, the results of those 12 popular contracts will heavily influence our empirical probability estimate. Moreover, standard error estimates would fail to correctly take into account the strong correlation structure in the data resulting from this

Quartile	Number of Contracts in Quartile
Q1	12
Q2	20
Q3	30
Q4	67

Table 1: The twelve most popular games received as many trades as the 67 least popular games. This over/under representation is resolved by introducing trading regimes.

popularity bias.

2.1 Trading Regimes

Rather than trying to tease out the effect of the cluster correlation between trades, we modify our motivating hypothetical and try to answer a very related and more tractable question: Imagine that we arrive at the market and see a *regime* of transactions occurring at price p. What is the probability that the event actually occurs?

Before defining exactly what a regime of transactions means, we identify how this questions differs from our original. In the original formulation, the implicit assumption in *arrive at the market and see a transaction* was that we arrive uniformly distribution over every trade. This is why we were initially led to average over every trade. The only different in this reformulation of our question is that rather than arriving uniformly over every trade, we arrive uniformly over every regime of trades. Thus our statistical analysis will involve regimes and not individual trades.

The point of splitting our data into regimes is to do away with issues of cluster correlation. We assume, without need for much justification, that the outcome of two distinct games (even if the same teams are playing) are independent. Furthermore, within a single game, we make the assumption that if trading on the contract hits 50 cents then the event that trades before that time pay out, and the event that trades after that time pay out are independent. This assumption deserves justification. Recall that by our earlier assumption we may use symmetry to view all trades as occurring at a price between 50 and 100. Consider now a trade occurring before the contract price hits 50 and then one after that event. Our assumption amounts to the independence of the parity of these two trades. This is justified since a price of 50 is strongly indicative of the market resetting its beliefs. After such a reset, there is no reason that market beliefs should be influenced by previous represented beliefs.

We now precisely define our regimes. Assume that we have grouped every trade based on contract and then within contracts based on time. The first trade in the first contract will be in our first regime. Hence forth we use the following rules to classify iteratively the remaining trades:

- A trade is in a new regime if
  - it is the first trade of a new contract, or
  - the parity of the price differs from the previous trade;
- otherwise the trade is in the same regime as the previous one.

Implementing this method of classification reduces our data to 744 regimes. Under the assumptions detailed above, trades in different regimes (even if they are in the same game) can be assumed to be on independent events. By removing the dependence structure inherent in counting every trade, we can now proceed with a more standard analysis of the probability of a binary event.

For every price bucket we count the number of regimes in which a trade is made in said bucket. Of those regimes we further count how many resulted in wins. If a regime was of trades with parity 1, and the contract expires to 100, this constitutes a win. Likewise for parity -1 regimes, a win is when the contract expires to 0. Lastly, if the regime had parity 0 (i.e. contained only bets at 50) then we randomly classify as either a win or a loss. The results of this tallying are recording in table 2. Taking the ratio of number of winning regimes to total number of regimes, we calculate and plot in figure 1 the empirical probability of winning, as well as the fair probability versus price.

In order to draw any meaningful statistical conclusions for this type of analysis we must correctly identity a robust method for estimating confidence intervals in small sample datasets. In this discussion we draw from the work of Sauro and Lewis [10] as well as Agresti and Coull [1] which both address the question of confidence intervals for binomial random variables. Sauro and Lewis explain that a 95% confidence interval is generally interpreted as meaning that out of 100 experiments which result in a particular confidence interval, 95 will have a true value within the

Range	Range Count	Wins in Range	Emp. Win Prob.	Fair Prob.	Adj. Wald Interval
$[50, 52)$	353	184	0.52	0.51	0.052
$[52, 54)$	176	91	0.52	0.53	0.073
$[54, 56)$	226	124	0.55	0.55	0.064
$[56, 58)$	134	75	0.56	0.57	0.083
$[58, 60)$	137	77	0.56	0.59	0.082
$[60, 62)$	201	126	0.63	0.61	0.066
$[62, 64)$	126	81	0.64	0.63	0.083
$[64, 66)$	167	112	0.67	0.65	0.071
$[66, 68)$	123	91	0.74	0.67	0.077
$[68, 70)$	128	93	0.73	0.69	0.077
$[70, 72)$	146	105	0.72	0.71	0.072
$[72, 74)$	92	71	0.77	0.73	0.085
$[74, 76)$	128	98	0.77	0.75	0.073
$[76, 78)$	76	56	0.74	0.77	0.098
$[78, 80)$	90	76	0.84	0.79	0.076
$[80, 82)$	112	93	0.83	0.81	0.070
$[82, 84)$	75	65	0.87	0.83	0.079
$[84, 86)$	106	95	0.90	0.85	0.060
$[86, 88)$	80	73	0.91	0.87	0.066
$[88, 90)$	82	76	0.93	0.89	0.061
$[90, 92)$	109	98	0.90	0.91	0.059
$[92, 94)$	73	70	0.96	0.93	0.055
$[94, 96)$	93	87	0.94	0.95	0.055
$[96, 98)$	78	76	0.97	0.97	0.047
$[98, 100)$	115	114	0.99	0.99	0.028

Table 2: Calculated empirical winning percentage and Adjusted Wald confidence intervals for trades occurring in different price buckets over entire contract period.

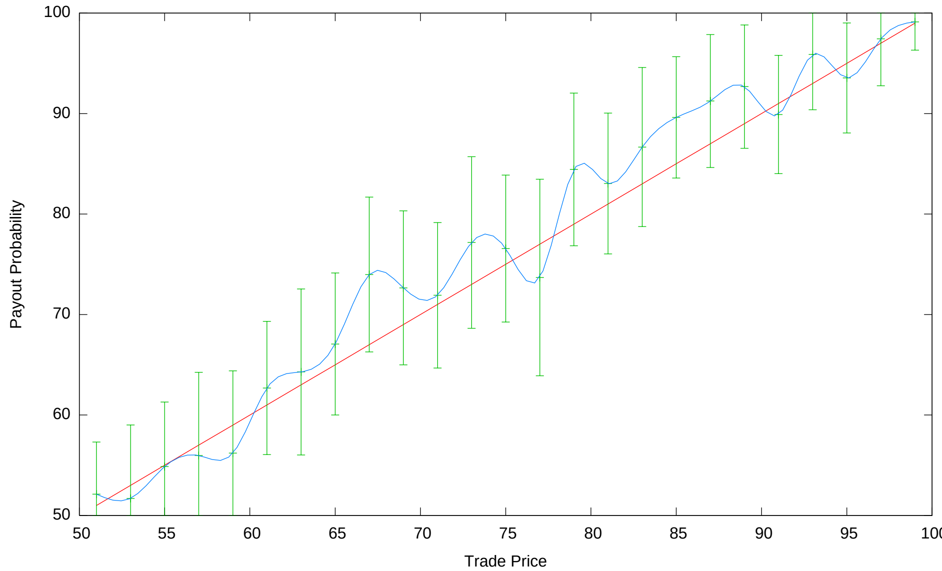

Figure 1: The straight line reflects the fair probability of payout for each given price bucket. The line with confidence intervals represents the empirically calculated probability of payout for each given price bucket gathered over all trades. The intervals represent the range of 95% confidence as calculated by the Adjusted Wald statistic.

interval. For large samples, when binomial distributions begin to approach normal distributions, standard statistical methods are applicable. However, as in our case, when datasets are smaller than or on the order of 100 samples per experiment (bucket), such standard techniques are not appropriate. The four techniques for small sample datasets discussed in Sauro and Lewis are the Exact, Score, Wald, and Adjusted Wald. To determine which of these is the best, SL study the *coverage* of each method, which records the actual percent of experiments which fall within the confidence interval. Target coverage is 95%, however some tests, such as the Exact actually have coverage close to 99% while the Wald has very bad covering in many cases. High coverage is not desirable since it implies that the confidence interval is excessively wide. Of the methods considered the Adjusted Wald was clearly the best as far as appropriate coverage and ease in computing.

The Adjusted Wald confidence interval is based on the calculation of the Wald confidence interval. Given an empirical probability $\hat{p}$ coming from n samples, the Wald interval equals

$$Z_{\alpha/2}\sqrt{\frac{\hat{p} * (1-\hat{p})}{n}},$$

where α is the failure probability, in our case $100\%-95\%$ and the corresponding $Z_{\alpha/2} = 1.96 \simeq 2$. This gives the interval above and below $\hat{p}$. To find the Adjusted Wald confidence interval, we "add two successes and two failures" and then calculate the Wald interval. That means that

$$\hat{p} = \frac{\text{successes} + 2}{\text{experiments} + 4}.$$

2.2 Longshot Bias Temporal Dependence

Since Tradesports trading is unique in that trading continues through the game as information grows, we further refine our motivating question by asking how the relationship between price and probability behaves as a function of time. We discriminate into two time periods of trades, those more than one hour before the end of the game (overtime included) and those less than one hour before the end of the game. When tallying up the number of regimes as well as the number of winning regimes for each bucket, we filter by whether the bucket contains trades in each of the two time periods. For each time period we conduct the same statistical analysis and form confidence intervals based on the Adjusted Wald method. These results are shown in figures 2 and 3.

The plot 2 of trades occurring before the last hour of trading does not show any systematic bias and with the exception of the regime $[84, 86)$, all confidence intervals overlap the fair price line. This implies that systematically prices before and early in the game accurately reflect objective game result probabilities. Turning attention to the graph 3 of trades occurring in the last hour of trading we see a pronounced shift, especially in the three bucket ranges of $[76, 78)$, $[78, 80)$, and $[80, 82)$. In these ranges the probability of a typical contract traded is significantly higher than the price of that contract. This is what we have previously referred to as the longshot bias. Surprising as it may be, the longshot bias does not develop in our data until the last hour of play.

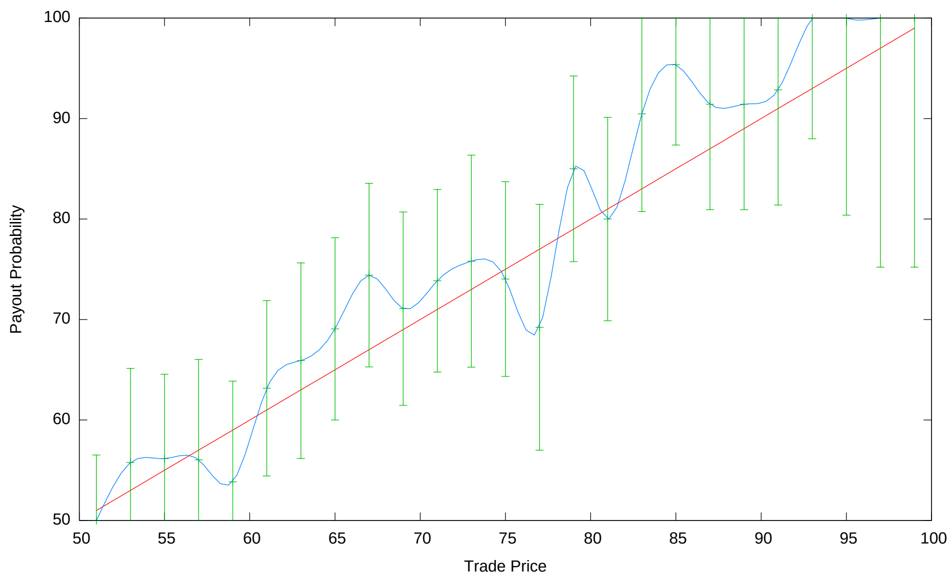

Figure 2: The straight line reflects the fair probability of payout for each given price bucket. The line with confidence intervals represents the empirically calculated probability of payout for each given price bucket gathered over all trades occurring earlier than the last hour of trading. The intervals represent the range of 95% confidence as calculated by the Adjusted Wald statistic.

2.3 Methods Analysis

It is important to point out a slight inconsistency in our method of temporally dividing the trading regimes into those occurring before and those occurring during the last hour of trading. Observe that an trader sees the amount of time left in regular play. There are two consideration relating to that observation. First, as far as game results are concerned, the relevant measure of time is game minutes, rather than real world hours. We estimate roughly that the one hour cutoff corresponds to halftime. However due to fouls, injuries and timeouts, there is some variance in the actual amount of game play corresponding to one hour of real world time. From experience this variance is small and since we believe that the longshot bias gradually emerges as the game is played, this variance will have a minor effect on our empirical results. The second, more important consideration is that our one hour cutoff does not take care of overtime, or double overtime. To improve our statistical analysis we should compensate for overtime since agents participating in the market well before the end of the game do not know that there will be extra

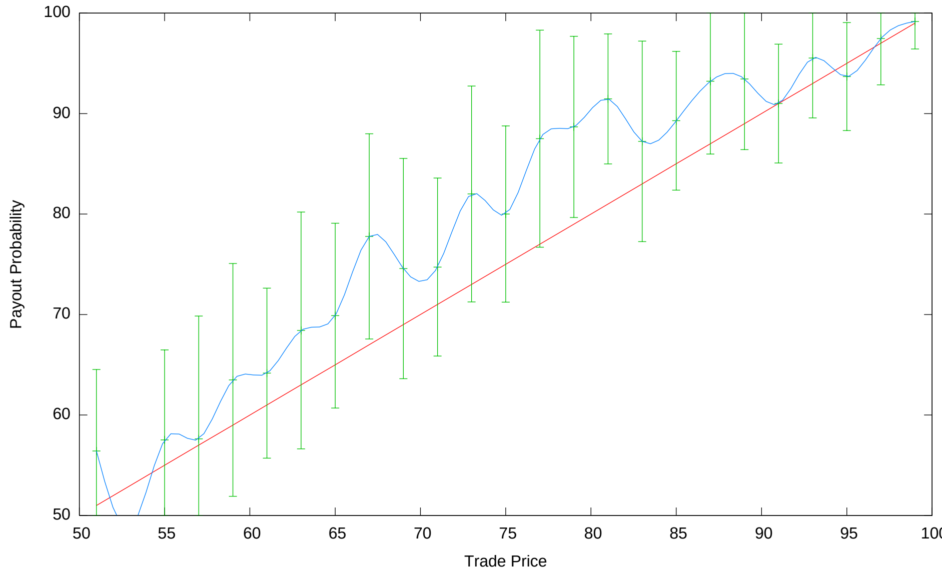

Figure 3: The straight line reflects the fair probability of payout for each given price bucket. The line with confidence intervals represents the empirically calculated probability of payout for each given price bucket gathered over all trades occurring in the last hour of trading. The intervals represent the range of 95% confidence as calculated by the Adjusted Wald statistic.

minutes of play.

A second methods consideration to note is our decision not to use cluster correlation variance analysis methods. Since our assumption is that the payoff of trades within a regime are completely correlated, and that payoffs between regimes are completely uncorrelated, cluster correlation methods seem overly complicated and general. Rather, we believe it is both more informative and truer to the nature of the data to use the simple correlation structure detailed above.

3 Theoretical Model

We now develop a theoretical model of price formation based on the concept of uncertainty aversion. We show that uncertainty aversion counteracts the longshot bias, and we argue that as the game unfolds uncertainty aversion disappears. This behavior captures our empirical results and explains the sharp rise in trade volume at the end of the games we see in the data.

3.1 Risk Neutrality and Uncertainty Aversion

It is by now well known that a pool of risk-neutral traders will display the longshot bias. The relation between this bias and the filtration of information to agents dates back to the original proposal of a risk-neutral model by Manski. As he points out, risk-neutral traders act strictly on the *expectation* of their private beliefs.

This observation eliminates perhaps the most natural model for decreasing uncertainty — agents having private belief distributions which tighten in response to information. A risk-neutral agent who believes that an event will occur with probability uniform on $[50, 90]$ will act in exactly the same way as a risk-neutral agent with beliefs uniform on $[65, 75]$. A model of this form would therefore be unable to explain our empirical observations.

Uncertainty aversion, however, does yield a model which corresponds to observed behavior. Uncertainty aversion was first developed in response to experimental studies which demonstrably violated tenets of orthodox expected utility theory [6]. The key to uncertainty aversion is that people prefer known risks to unknown risks, even if a distribution over the super-parameters controlling the unknown risk may imply a different course of action.

We hypothesize that sporting events function in the same manner. Before a game is played and in its earliest stages, there are an assortment of variables which one can imagine as super-parameters over the space of future outcomes. As the game plays out, more of these variables are revealed, and the risks become better known. A set of vague questions — Will a team hustle on defense? Will a key player have a good game? — collapses into a single, direct question: given the game so far, what's the chance a team maintains its lead?

Uncertainty aversion was formalized mathematically by Schmiedler [11], who modeled the phenomenon as a non-additive probability measure. More constructively, Dow and Werlang modeled the investment decisions of an uncertainty-averse agent and showed the existence of a pricing gap under which the agent, even if risk neutral, would not invest [5]. Our model owes the most debt to their work.

3.2 Modeling with Uncertainty Aversion

Imagine that teams A and B are playing a basketball game and let π_A represent an agent's belief that team A will win, and π_B represent an agent's belief that team B will win. Uncertainty aversion indicates that $\pi_A + \pi_B < 1$ and in particular measures the defect $1 - \pi_A - \pi_B$. Let the price of a contract paying one dollar when team A wins the game be p_A and the price of a contract which pays one dollar when team B wins be $p_B = 1 - p_A$. A risk-neutral agent holding a position (x_A, x_B) then has expected utility:

$$\begin{aligned} EU(x_A, x_B) &= \pi_A(p_B x_A - p_B x_B) + \pi_B(p_A x_B - p_A x_A) - (1 - \pi_A - \pi_B)(p_A x_A + p_B x_B) \\ &= x_A(\pi_A - p_A) + x_B(\pi_B - p_B). \end{aligned}$$

Examining $\frac{\partial EU}{\partial x_A}$, we see that

$$\frac{\partial EU}{\partial x_A} = \pi_A - p_A.$$

This is greater than zero if and only if:

$$\pi_A > p_A.$$

Symmetric equations hold with respect to x_B. Note that these values do not depend on x and so risk-neutral agents will invest their entire endowment if this expression is positive. We fix each agents' endowment to be a constant w.

Note that since, under uncertainty aversion, an agent's values for π_A and π_B sum to less than 1, there exist belief vectors (π_A, π_B) where agents do not invest (this result is roughly analogous to the one in Dow and Werlang). Let $\pi_A + \pi_B = k$, and fix k such that all agents have the same k. For example, let $p_A = .4$ and $k = .9$. Then only agents with $\pi_A > .4$ or $\pi_B > .6$ (equivalently $(\pi_A < .3)$) will have positive expectation and invest their endowment. Agents with $.3 < \pi_A < .4$ $(.5 < \pi_B < .6)$ will not participate in the market.

As Dow and Werlang noted, the behavior induced by uncertainty aversion resembles that seen by agents who operate with a bid/ask spread, or under transaction costs. As we will show, the effect of this behavior on prices counters the longshot bias. This is an intriguing observation, because Wolfers and Zitzewitz have suggested a bid/ask spread or transaction costs as being a *cause* of the longshot bias [14].

To assess the effect of uncertainty aversion on equilibrium prices, let the distribution of agent beliefs be given by f. Without uncertainty aversion and with risk-neutral traders and uniform endowments, price p^* is an equilibrium price if:

$$\frac{p^*}{1-p^*}\int_0^{p^*} f(b)\,db = \int_{p^*}^1 f(b)\,db.$$

Or as Manski writes, p^* is an equilibrium if p^* represents the $1-p^*$-st percentile of beliefs.

Without uncertainty aversion we typically model beliefs by specifying the distribution over a single parameter, say π_A, for market participants, which implicitly also defines a distribution over agent beliefs on π_B, because $\pi_A + \pi_B = 1$. We want to generalize the concept of belief distribution to the sub-additive probability measure setting of uncertainty aversion. We propose the following model: a distribution over beliefs is generated when agents are not uncertainty averse, and then a value for k is specified. This value of k represents a multiplicative scaling of the domain space, such that on an unadjusted belief density of $f(\pi_A, \pi_B) = f(\pi_A, 1-\pi_A) = f(\pi_A)$ becomes $f(k\pi_A, k\pi_B) = f(k\pi_A)$. Put another way, the belief density function f represents a function over the *share of total belief* an agent has over a team winning. For instance, without uncertainty aversion, ($k = 1$) the density of agents who believe that the teams have the same chance of winning ($\pi_A = \pi_B = .5$) is equivalent to, with uncertainty aversion ($k = .8$), the density of agents who believe that the teams have the same chance of winning ($\pi_A = \pi_B = .4$).

Under these specifications, a market composed of risk-neutral, homogeneously endowed traders reaches equilibrium at the price p such that:

$$p\int_0^{pk} f(b)\,db = (1-p)\int_{1-(1-p)k}^1 f(b)\,db.$$

Note that this expression collapses back to Manski's familiar equilibrium equation when $k = 1$.

We now define the excess supply function Φ.

$$\Phi(k,p) \equiv p \int_0^{pk} f(b)\,db - (1-p) \int_{1-(1-p)k}^{1} f(b)\,db.$$

$\Phi(\cdot,\cdot)$ indicates the imbalance between the two sides of the market. If $\Phi(k,p) > 0$, the "sell side" is over-represented and the equilibrium price will be smaller than p. If $\Phi(k,p) < 0$ the "buy side" is over-represented and the equilibrium price will be found at a price higher than p. By definition, $\Phi(1,p^*) = 0$.

To investigate how uncertainty aversion affects equilibrium prices, we examine $\frac{\partial \Phi}{\partial k}$.

$$\frac{\partial \Phi(k,p)}{\partial k} = p^2 f(pk) - (1-p)^2 f(1-(1-p)k).$$

Now evaluating at $k = 1$ and $p = p^*$:

$$\begin{aligned} \frac{\partial \Phi(1,p^*)}{\partial k} &= p^{*2} f(p^*) - (1-p^*)^2 f(p^*) \\ &= f(p^*)(2p^* - 1). \end{aligned}$$

Now note that under the assumption that $f(p^*) > 0$, the sign of this equation is the same as the sign of $2p^* - 1$. This function has a root at .5 and is negative below that value and positive above.

As a result, equilibrium prices are pushed downwards below .5, and pushed upwards above .5. Thus, uncertainty aversion counters the longshot bias, because in the longshot bias equilibrium prices under .5 are overpriced, while equilibrium prices above .5 are underpriced.

Our result holds only at a single (but most interesting) point for k, and short of placing stringent requirements on the shape of belief distributions, we are unable to make a more general statement regarding the sign of this derivative.

However, note that we are only interested in relatively large values of k, which makes our inability to make claims about values of k further away from 1 less of an issue. This is because even a relatively small gap between k and 1 can result in large, unrealistic changes to the state space. For instance, with $k = .9$ traders with beliefs falling in at least 10% of the applicable domain are unable to participate (beliefs in a 10 cent gap around the equilibrium price). Since we should expect, speaking broadly, that there would be a relatively significant share of agent beliefs around the equilibrium price, this means that a much larger percentage of the trading population may be unwilling to participate.

One practical, testable effect based on our model of uncertainty aversion is that the trading volume should rise dramatically as the game nears its conclusion. As we have argued, this is because agents become more comfortable (that is, less uncertainty averse) as the gambles affecting their game become more direct. Figure 4 shows the dramatic rise in trade volume near the end of trading.

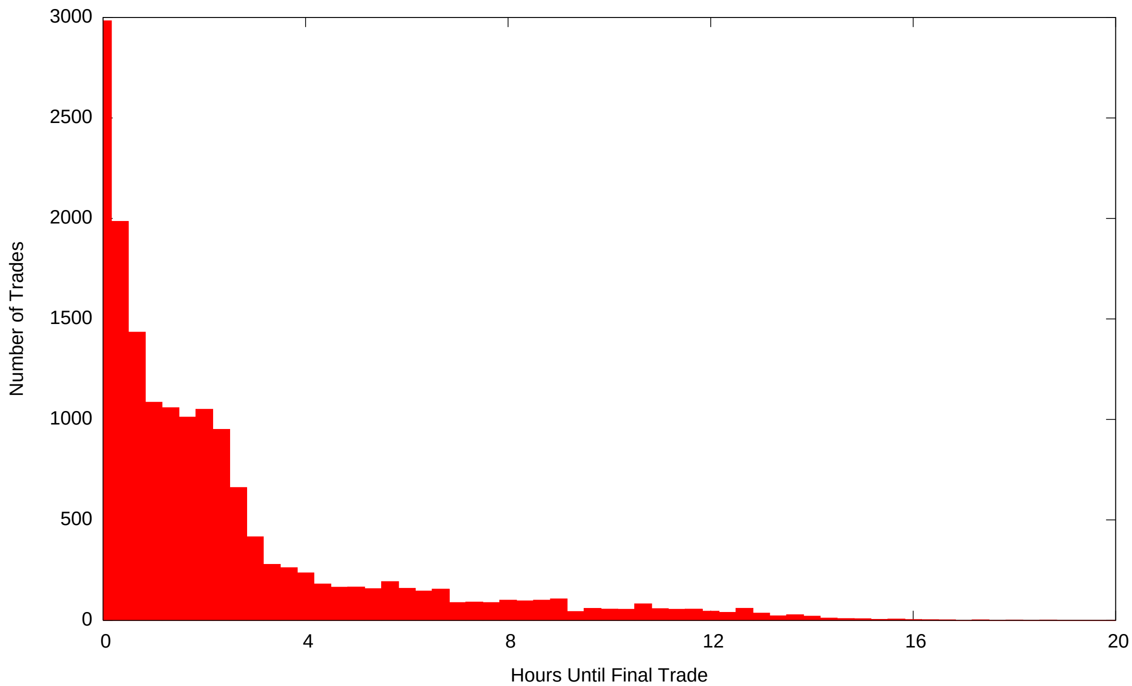

Figure 4: Trading volume increases near the end of trading.

It is interesting to compare this figure to the last-minute bidding behavior seen in eBay auctions [9]. In particular, both market formats show a dramatic rise in activity at their conclusion. Yet, intuitively, we should expect incentives to participate in the formats to be wholly opposite — shrewd and informed eBay bidders should avoid sending "high type" signals to other market participants, while better-informed traders in a prediction market should take advantage of the difference between market prices and their information and participate as soon as possible. A decreasing level of uncertainty aversion over time for agents in a prediction market provides a plausible explanation for the rise in trading volume at the end of trading.

4 Conclusion

Using the extensive dataset of all Tradesports trades for the 2005-06 NBA season, we determine the expected payouts of trades occurring in every price range. We analyze the effect of time until contract expiration on the expected payout. Using only first principles we determine confidence intervals which take into account the cluster correlation structure of trades occurring in the same game. Our analysis shows that in the last hour of trading, but not prior to that, there is a pronounce and statistically significant longshot bias. That is to say, in the last hour, prices in a range around 80 cents are lower than the expected payout, and pricing in the range around 20 cents are higher than the expected payout.

The methods used in this paper could be applied to many other sports such as football, baseball, hockey or soccer. By just looking at a single season of basketball it was not statistically reasonable to split trading into more than two time periods (more than one hour, and less than one hour left of betting). However, with a larger dataset, we could further refine our time periods to see whether the longshot bias emerges slowly as trading time diminishes.

We also presented a new theory of prediction market participation based around uncertainty aversion. Conventional risk-neutral agents always invest their endowment in one side of the market or another, and as Manski pointed out, do not respond to "learning" in the market as long as the agent does not flip sides of the market. The risk-neutral model, then, circumscribes what we can analyze with regards to learning and information absorption. As a practical matter, these restrictions manifest themselves in the inability of risk-neutral agents to produce behavior of the kind we observed in the NBA contract data. Uncertainty aversion, a framework based around acting on subjective beliefs, provides an intuitive extension to models of equilibrium which already involve agents making investment decisions based on their independent, subjective beliefs.

References

[1] A. Agresti and B. Coull. Approximate Is Better Than "Exact" for Interval Estimation of Binomial Proportions. *The American Statistician*, 52(2), 1998.

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

[3] R. Borghesi. Price biases in a prediction market: NFL contracts on Tradesports. *The Journal of Prediction Markets*, 1:233–253, December 2007.

[4] M. Cain, D. Law, and D. Peel. The Favourite-Longshot Bias and Market Efficiency in UK Football betting. *Scottish Journal of Political Economy*, 47(1):25–36, 2000.

[5] J. Dow and S. Werlang. Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio. *Econometrica*, 60(1):197–204, 1992.

[6] D. Ellsberg. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75(4):643–669, 1961.

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

[8] C. Manski. Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425–429, 2006.

[9] A. Roth and A. Ockenfels. Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet. *American Economic Review*, 92(4):1093–1103, 2002.

[10] J. Sauro and J. R. Lewis. Estimating completion rates from small samples using binomial confidence intervals: Comparisons and recommendations. *Human Factors and Ergonomics Society Annual Meeting Proceedings*, 49:2100–2104, 2005.

[11] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989.

[12] P. Tetlock. Liquidity and prediction market efficiency. Technical report, Yale University, March 2008.

[13] R. Williams. A Note on Robust Variance Estimation for Cluster-Correlated Data. *Biometrics*, 56(2):645–646, 2000.

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

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

[16] L. Woodland and B. Woodland. Market Efficiency and the Favorite-Longshot Bias: The Baseball Betting Market. *The Journal of Finance*, 49(1):269–279, 1994.

[17] R. Woodruff. A Simple Method for Approximating the Variance of a Complicated Estimate. *Journal of the American Statistical Association*, 66(334):411–414, 1971.