Betting Strategies, Market Selection, and the Wisdom of Crowds

Willemien Kets
Northwestern University
w-kets@kellogg.northwestern.edu

David M. Pennock
Microsoft Research New York City
dpennock@microsoft.com

Rajiv Sethi
Barnard College, Columbia University
Santa Fe Institute
rs328@columbia.edu

Nisarg Shah
Carnegie Mellon University
nkshah@cs.cmu.edu

Abstract

We investigate the limiting behavior of trader wealth and prices in a simple prediction market with a finite set of participants having heterogeneous beliefs. Traders bet repeatedly on the outcome of a binary event with fixed Bernoulli success probability. A class of strategies, including (fractional) Kelly betting and constant relative risk aversion (CRRA) are considered. We show that when traders are willing to risk only a small fraction of their wealth in any period, belief heterogeneity can persist indefinitely; if bets are large in proportion to wealth then only the most accurate belief type survives. The market price is more accurate in the long run when traders with *less* accurate beliefs also survive. That is, the survival of traders with heterogeneous beliefs, some less accurate than others, allows the market price to better reflect the objective probability of the event in the long run.

Introduction

Prediction markets have attracted considerable attention in recent years, both as vehicles for the aggregation of diverse opinions, and as independent objects of research (Chen and Plott 2002; Cowgill, Wolfers, and Zitzewitz 2009; Ungar et al. 2012).

A central question in the analysis of such markets concerns the interpretation and accuracy of prices. In the case of a single event, with prices determined by a static market clearing condition, the answer depends crucially on the preferences of traders. When traders have log utility, the equilibrium price equals the budget-weighted average of trader beliefs (Pennock 1999). This remains approximately true for other preference specifications as long as traders are risk-averse (Gjerstad 2004; Wolfers and Zitzewitz 2006). Under risk-neutrality, however, the equilibrium price corresponds to a quantile of the belief distribution, which can be quite distant from the average belief (Manski 2006).

In markets with betting on a sequence of events, the interpretation of prices is more challenging, because the wealth of traders change endogenously as the outcomes of bets are realized. This is the setting with which the present paper is concerned. Specifically, we consider the limiting distributions of trader wealth and market prices when betting occurs over an infinite sequence of periods. In each period traders bet against each other on the outcome of a binary event with fixed (but unknown) Bernoulli probability. The sizes of their bets depend on their wealth and their beliefs relative to the market price, and the price is determined by market clearing. Trader preferences belong to a class that includes constant relative risk aversion (CRRA) utility, as well as (fractional) Kelly betting, both defined below.

We show that when preferences are such that traders bet only a small portion of their wealth in any given period, the limiting wealth distribution is non-degenerate in the sense that traders with beliefs of varying degrees of accuracy have positive expected wealth in the limit. In this case, the market price in the long run reflects the objective probability of the event on which bets are placed. In contrast, when traders place large bets (as a proportion of their wealth) then only the trader with the most accurate belief survives. In this case the market price comes to reflect this trader's belief, which may not be close to the objective probability of the event. Hence the survival of relatively *inaccurate* traders is necessary for the market to make accurate forecasts. This survival is ensured if traders are reluctant to place large bets relative to their income, for instance because of high levels of risk aversion. That is, high levels of risk aversion not only allow inaccurate beliefs to survive *longer*, but allow their survival *in the limit* (in cases where they previously did not survive) and cause the market forecast to become more accurate.

Our work is related to a literature on market selection that dates back at least to Kelly (1956), with important contributions by De Long et al. (1990), Blume and Easley (1992; 2006), Sandroni (2000), and Beygelzimer, Langford, and Pennock (2012), among many others. Blume and Easley (2006) and Sandroni (2000) show that when markets are complete and traders can optimally allocate wealth across periods, traders who survive have the most accurate beliefs (assuming that they all have a common discount factor). These results require that traders make optimal consumption and asset allocation plans over an infinite horizon, and do not apply to our setting, in which there is no consumption and asset allocation is myopic. De Long et al. (1990) show, in an overlapping-generations model, that traders with incorrect beliefs can survive, since they take on more risk, and this is rewarded with greater expected returns. By contrast, in our setting, the survival of traders with inaccurate beliefs requires them to take on *less* risk.

Preliminaries

We consider a prediction market with n traders who make a sequence of bets against each other on binary events. Let w_i^t denote the wealth of trader i at the end of period t, with $w_i^0 > 0$ denoting the initial wealth levels. In each period t, the outcome of the event X^t is an independent Bernoulli trial with success probability p^*; this probability is unknown to the traders and never revealed. Trader i's belief about the success probability is denoted p_i. These beliefs are exogenously given and are not updated over time, although traders with inaccurate beliefs may face loss of wealth and eventual disappearance from the market.

Traders place bets in each period based on their beliefs, wealth, and the market price. If the market price is p_m, a buy order (a bet for the outcome) costs p_m, and pays off 1 if the outcome is positive, and 0 otherwise. Similarly, a sell order (a bet against the outcome) costs $1 - p_m$, and pays off 1 if the outcome is negative, and 0 otherwise. (In practice, a trade of one unit occurs when a buyer pays p_m and a seller pays $1 - p_m$ to the exchange; the entire sum is delivered to the buyer if the event occurs and to the seller otherwise.) We assume that orders can be of arbitrary size, including fractional units.

At any given price p_m, each trader has an optimal order size, which can be aggregated across traders to obtain a net demand to buy or sell. A competitive equilibrium price is one at which the buy and sell orders are in balance, so that the aggregate net demand is zero. At this price all orders can be executed. Let p_m^t denote the market clearing price in period t.

At the end of each period, the outcome is revealed, and the traders' bets pay off accordingly. Wealth is redistributed among traders, but remains constant in the aggregate. This process is then repeated indefinitely. We normalize the wealth of the traders such that $\sum_{i=1}^n w_i^0 = 1$. Since wealth is only redistributed, we have $\sum_{i=1}^n w_i^t = 1$ for all t.

We assume that each trader i uses a *betting strategy* represented by a function $b_i : [0,1] \to [0,1]$ that denotes the fraction of his wealth that the trader wishes to bet as a function of the market price. That is, with wealth w and market price p, the trader would bet $wb_i(p)$ of his wealth. This specification allows for a variety of common betting strategies, some of which are discussed below. We assume that each b_i is fixed over time, and satisfies the following mild conditions.

1. $0 \le b_i(p_m) \le 1$ for all $p_m \in [0,1]$, i.e., the trader cannot bet more than his wealth.
2. $b_i(p_m) = 0$ if and only if $p_m = p_i$, i.e., the trader does not bet anything if and only if the market price matches his belief.
3. $b_i(p_m) < 1$ if $p_m \in (0,1)$, i.e., the trader does not bet his entire wealth except when betting for or against the event is free. This rules out risk-neutrality.
4. The farther the market price from his belief, the higher fraction of his wealth the trader bets. That is, $b_i(p_a) \ge b_i(p_b)$ for $p_a \le p_b \le p_i$ and $b_i(p_a) \ge b_i(p_b)$ for $p_a \ge p_b \ge p_i$.

Special cases of popular betting strategies that satisfy our assumptions include fractional Kelly betting and myopic betting with constant relative risk aversion (CRRA) preferences. If p denotes the belief of the trader, Kelly betting, the optimizer of log utility, is given by the following formula:

$$b(p_m) = \begin{cases} \frac{p-p_m}{1-p_m} & \text{if } p_m \le p, \\ \frac{p_m - p}{p_m} & \text{if } p_m > p. \end{cases}$$

Fractional Kelly betting is similar, but with reduced exposure to risk. Specifically, a c-Kelly trader bets a fixed fraction c of the share of wealth that a Kelly trader with the same belief would bet. As shown by Beygelzimer, Langford, and Pennock (2012), a c-Kelly trader acts like a Kelly trader whose belief is a linear combination of his own belief and the market price, given by $cp + (1-c)p_m$. The higher the value of c, the lower the weight on the market price. Hence, c can be interpreted as the level of confidence in one's own belief.

The η-CRRA betting strategy places the optimum myopic bet according to the expected value of the power utility (isoelastic utility) function of wealth. Specifically, if w denotes the end-of-the-period wealth, the bettor maximizes the expected value of

$$u(w) = \begin{cases} \frac{w^{1-\eta}-1}{1-\eta} & \eta > 0, \eta \ne 1, \\ \log(w) & \eta = 1. \end{cases}$$

It is easily verified that this strategy satisfies our assumptions.[1] Clearly, η-CRRA with $\eta = 1$ corresponds to Kelly betting. Hence, both c-Kelly and η-CRRA can be seen as generalizations of Kelly betting, where the risk aversion or confidence of the trader can vary.

Without loss of generality, assume that the traders are sorted by their beliefs, i.e., $p_i \le p_{i+1}$ for all i, and that $0 < p_1 < p_n < 1$. That is, no trader is subjectively certain of the outcome, and not all beliefs are identical. It is easily seen that the market clearing price p_m^t in all periods must lie strictly between p_1 and p_n as long as both these traders have positive wealth. Therefore, trader n consistently bets that the event will occur, and trader 1 consistently bets against it.

Theoretical Results

In this section, we present a few theoretical results that will provide support for our empirical observations in the next section.

Expected Wealth and Prices

Our first result shows that the survival of traders with heterogeneous beliefs is beneficial for the accuracy of the market as a whole.

Theorem 1. *Consider two traders with beliefs p_1 and p_2, and wealth w_1^{t-1} and w_2^{t-1} respectively at the end of period $t-1$. If $\mathbb{E}[w_1^t | w_1^{t-1}] = w_1^{t-1}$ and $\mathbb{E}[w_2^t | w_2^{t-1}] = w_2^{t-1}$, then one of the following conditions holds.*

[1] See Wolfers and Zitzewitz (2006) for an explicit expression for bet size given η-CRRA preferences.

1. *Only trader* 1 *survives, i.e.* $w_1^{t-1} = 1$ *and* $w_2^{t-1} = 0$.
2. *Only trader* 2 *survives, i.e.* $w_2^{t-1} = 1$ *and* $w_1^{t-1} = 0$.
3. *Both traders survive with unique wealth in* $(0, 1)$*, and the market price in the next period is* $p_m^t = p^*$.

Proof. Note that $\mathbb{E}[w_1^t | w_1^{t-1}] = w_1^{t-1}$ and $\mathbb{E}[w_2^t | w_2^{t-1}] = w_2^{t-1}$ are equivalent because $w_1^{t-1} + w_2^{t-1} = w_1^t + w_2^t = 1$. In period t, trader 2 bets $w_2^{t-1} \cdot b_2(p_m^t)$ for the outcome (because $p_1 \le p_m^t \le p_2$). Due to the Markov property of his wealth and that the success probability of the event is p^*, we have

$$\begin{aligned} E\left[w_2^t \,\middle|\, w_2^{t-1}\right] &= w_2^{t-1} \Bigg[p^* \left(1 - b_2(p_m^t) + \frac{b_2(p_m^t)}{p_m^t} \right) \\ &\qquad + (1 - p^*) \left(1 - b_2(p_m^t)\right) \Bigg] \\ &= w_2^{t-1} \left[1 - b_2(p_m^t) \cdot \frac{p_m^t - p^*}{p_m^t} \right]. \qquad (1) \end{aligned}$$

Now, $E\left[w_2^t \,\middle|\, w_2^{t-1}\right] = w_2^{t-1}$ holds if $w_2^{t-1} = 0$ (thus $w_1^{t-1} = 1$). It also holds if $b_2(p_m^t) = 0$, which happens if and only if $p_m^t = p_2$ due to our assumptions on the betting strategies. However, the latter must imply $w_2^{t-1} = 1$ and $w_1^{t-1} = 0$, otherwise trader 1 would bet a positive amount while trader 2 will bet nothing, violating the fact that p_m^t is the market clearing price.

Finally, if neither of the above cases hold (i.e., both traders survive), then we must have $p_m^t = p^*$ in Equation (1). In this case, the market clearing equation at price p^*, namely

$$\frac{w_1^{t-1} \cdot b_1(p^*)}{1 - p^*} = \frac{w_2^{t-1} \cdot b_2(p^*)}{p^*}, \qquad (2)$$

and the wealth normalization equation $w_1^{t-1} + w_2^{t-1} = 1$ yield unique values for w_1^{t-1} and w_2^{t-1}. These values further belong to $(0, 1)$, because the numerators on both sides of Equation (2) are non-zero as both traders survive in this case. □

We remark that Theorem 1 computes the fixed points of the wealth and price updates, and not their stationary distributions. Hence, it only provides weak evidence that survival of both traders would lead the expected market price to converge to p^* (under the condition that expected future wealth equals current wealth).

We now show that with two traders, in the extreme cases of $p^* \le p_1$ and $p^* \ge p_2$, only trader 1 and only trader 2 survive, respectively. This realization of conditions 1 and 2 of Theorem 1 comes with a stronger notion of convergence — *almost sure convergence*.

Theorem 2. *For two traders with beliefs* p_1 *and* p_2*, if* $p^* \le p_1$ *(resp.* $p^* \ge p_2$*), then as* $t \to \infty$*,*

1. *almost surely,* $w_1^t \to 1$ *and* $w_2^t \to 0$ *(resp.* $w_2^t \to 1$ *and* $w_1^t \to 0$*), and*
2. *almost surely,* $p_m^t \to p_1$ *(resp.* $p_m^t \to p_2$*).*

Proof. Without loss of generality, we prove the case of $p^* \le p_1$. First, we show that $w_2^t \to 0$ almost surely. Consider Equation (1) from the proof of Theorem 1. It is easy to check that $p_m^t > p_1 \ge p^*$ implies $E[w_2^t | w_2^{t-1}] < w_2^{t-1}$. That is, $\{w_2^t\}_{t \ge 0}$ is a bounded super-martingale. By Doob's super-martingale convergence theorem, w_2^t converges almost surely. It remains to show that the limit is 0.

Almost sure convergence of w_2^t implies almost sure convergence of w_1^t, thus of p_m^t. If the limit of w_2^t were greater than 0, the limit of p_m^t would be strictly between p_1 and p_2. Hence, the ratio $E\left[w_2^t \,\middle|\, w_2^{t-1}\right] / w_2^{t-1}$ would converge to a constant less than 1, which is a contradiction. Thus, almost surely, w_2^t converges to 0, and w_1^t and p_m^t converge to 1 and p_1, respectively. □

Kelly traders

Beygelzimer, Langford, and Pennock (2012) analyzed markets where all traders use Kelly betting. They showed that in such a market, the market clearing price is the wealth-weighted average of the trader beliefs, and gave an explicit formula for wealth updates over time.

Proposition 1. *When all traders use Kelly betting, for every period* t*,* $p_m^t = \sum_{i=1}^n w_i^{t-1} \cdot p_i$*.*

Proposition 2. *When all traders use Kelly betting, for every period* t *and every trader* i*,*

$$w_i^t = w_i^0 \cdot \prod_{l=1}^t \left(\frac{p_i}{p_m^l} \right)^{X^l} \cdot \left(\frac{1 - p_i}{1 - p_m^l} \right)^{1 - X^l} .$$

We can now derive the following result.

Theorem 3. *If all traders use Kelly betting and a unique trader* j *maximizes* $p^* \log(p_j) + (1 - p^*) \log(1 - p_j)$*, then as* $t \to \infty$*,*

1. *almost surely,* $w_j^t \to 1$ *and* $w_i^t \to 0$ *for all* $i \ne j$*, and*
2. *almost surely,* $p_m^t \to p_j$*.*

Proof. Consider the most accurate trader j and any other trader i. From Proposition 2,

$$\frac{w_j^t}{w_i^t} = \alpha \cdot \frac{\prod_{l=1}^t p_j^{X^l} \cdot (1 - p_j)^{1 - X^l}}{\prod_{l=1}^t p_i^{X^l} \cdot (1 - p_i)^{1 - X^l}},$$

where $\alpha = w_j^0 / w_i^0$ is a constant. Next, we take logarithm on both sides and divide by t. The term $\log(\alpha)/t$ goes to 0 as $t \to \infty$ because we assume non-zero initial wealth. Using the strong law of large numbers, the logarithms of the numerator and the denominator converge to $p^* \cdot \log(p_j) + (1 - p^*) \cdot \log(1 - p_j)$ and $p^* \cdot \log(p_i) + (1 - p^*) \cdot \log(1 - p_i)$ respectively. Since the former is greater than the latter by our assumption, we have that $\lim_{t \to \infty} (\log(w_j^t) - \log(w_i^t))/t > 0$ almost surely. This can happen only when $\lim_{t \to \infty} w_i^t = 0$ almost surely. Since this holds for all $i \ne j$, we have $\lim_{t \to \infty} w_j^t = 1$ almost surely. Finally, using Proposition 1, we have that $\lim_{t \to \infty} p_m^t = p_j$ almost surely. □

Blume and Easley (1992) obtained a similar result in a somewhat different context, with assets in positive net supply, traders investing a fixed fraction of their wealth in each period, and the total wealth varying based on the bets placed and the state realized.

Fractional Kelly Traders

As noted above, a fractional Kelly trader with belief p acts like a Kelly trader whose belief is a linear combination of p and the market price p_m. Using the methods in Beygelzimer, Langford, and Pennock (2012), Propositions 1 and 2 can be generalized to a market where all traders use c-Kelly betting with the same c.

Proposition 3. *When all traders use c-Kelly betting with the same c, we have that for all t,*

$$p_m^t = \sum_{j=1}^{n} w_j^{t-1} \cdot (c \cdot p_j + (1-c) \cdot p_m^t) \Rightarrow p_m^t = \sum_{j=1}^{n} w_j^{t-1} \cdot p_j.$$

Proposition 4. *When all traders use c-Kelly betting with the same c, for every period t and every trader i,*

$$w_i^t = w_i^0 \cdot \prod_{l=1}^{t} \left(\frac{c \cdot p_i + (1-c) \cdot p_m^t}{p_m^l} \right)^{X^l} \cdot \left(\frac{1 - (c \cdot p_i + (1-c) \cdot p_m^t)}{1 - p_m^l} \right)^{1-X^l}.$$

Although we were not able to solve for the limiting wealth analytically, one can derive an update rule for the wealth distribution using Propositions 3 and 4. Consider the simple case of two traders with beliefs p_1 and p_2. Let $F^{(t)}$ denote the CDF of the distribution of the wealth of trader 1 at the end of period t, i.e., that of w_1^t. Then,

$$F^{(t)}(x) = p^* F^{(t-1)}(f_1^{-1}(x)) + (1-p^*) F^{(t-1)}(f_0^{-1}(x)),$$

where f_1 and f_0 defined below are the wealth update functions in case the outcome is positive and negative respectively.

$$f_1(x) = x \cdot \frac{c \cdot p_1 + (1-c) \cdot p_m}{p_m},$$
$$f_0(x) = x \cdot \frac{c \cdot (1-p_1) + (1-c) \cdot (1-p_m)}{1-p_m},$$
$$p_m = x \cdot p_1 + (1-x) \cdot p_2.$$

In the experiments described below, we observe that both traders survive for lower values of c. In fact, we conjecture (as verified by our simulations) that as $c \to 0$,

1. $\mathbb{E}[w_1^t] \to (p_2 - p^*)/(p_2 - p_1)$, and
2. $\mathbb{E}[w_2^t] \to (p^* - p_1)/(p_2 - p_1)$.

Using Proposition 3 and the linearity of expectation, we then have $\mathbb{E}[p_m^t] \to p^*$, as $t \to \infty$.

Simulation Results

We performed extensive simulations using markets having different combinations of traders with heterogeneous beliefs and strategies, and with different values of the Bernoulli success probability p^*. We are interested in the limiting (expected) trader wealth and market price. In each of the following simulations, the market is run for 10^5 iterations. The average wealth of each trader and the average market price over the last 10000 iterations is used as an estimate of the expected values from the corresponding stationary distributions. These estimates are further averaged over 1000 independent runs, and the 5th and the 95th percentiles are used for confidence intervals.

In our simulations, we are interested in the stationary distributions of the market price p_m^t and the trader wealth w_i^t. Unless specified otherwise, we say that a trader *survives* if his expected limiting wealth is positive, and *vanishes* if it is zero; the reader should distinguish this from stronger notions of survival, e.g., in which the limit infimum of the wealth must converge to a positive value almost surely.

Kelly and fractional Kelly betting

First, we analyze markets with Kelly and fractional Kelly traders. We begin with a market that has two c-Kelly traders: trader 1 with belief $p_1 = 0.3$ and trader 2 with belief $p_2 = 0.8$. We are interested in the limiting behavior of trader wealth and market price. We present three experiments: First, we vary c from 0 to 1, then we analyze the limiting case of $c = 1$, followed by the limiting case of $c = 0.01$. We choose $c = 0.01$ as a proxy for $c = 0$ because our assumptions do not allow $c = 0$.

Varying c. To vary the Kelly fraction c, we fix the objective probability of the event to $p^* = 0.6$. Thus, trader 2 is more accurate than trader 1. Figures 1(a) and 1(b) show the limiting expected wealth and market price respectively in this case. Note that at $c = 1$ (the right end of the figure), only the most accurate trader (trader 2) survives and the market price converges to his belief, as predicted by Theorem 3. At $c = 0.01$, however, both traders survive, and the limiting market price coincides with the objective event probability p^*, which matches our conjecture in the previous section. In fact, the survival of both traders is observed for a large range of values of c.

These observations also generalize to the case of more than two traders. Due to lack of space, we only provide graphs for the case of 3 traders. Figures 1(c) and 1(d) show the limiting expected wealth and market price respectively when an additional trader with belief $p_3 = 0.5$ is added to the market. It can be seen that all three traders survive at $c = 0.01$, but the least accurate trader 1 quickly vanishes as c increases. As c further approaches 1, only the most accurate trader 3 survives. The market price also gradually shifts from p^* to the belief of the most accurate trader.

These experiments are consistent with our hypothesis that when the traders are willing to bet relatively low fractions of their wealth, even the less accurate traders survive indefinitely, and in turn help improve the accuracy of the market

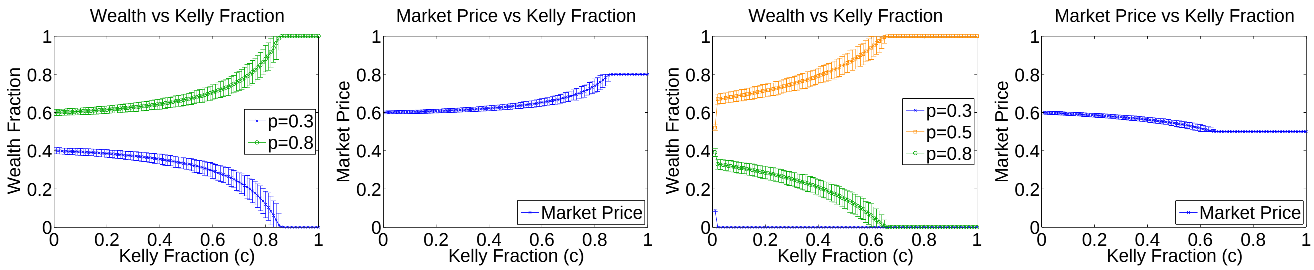

(a) $n = 2$, trader wealth (b) $n = 2$, market price (c) $n = 3$, trader wealth (d) $n = 3$, market price

Figure 1: Varying the Kelly fraction c from 0.01 to 1.

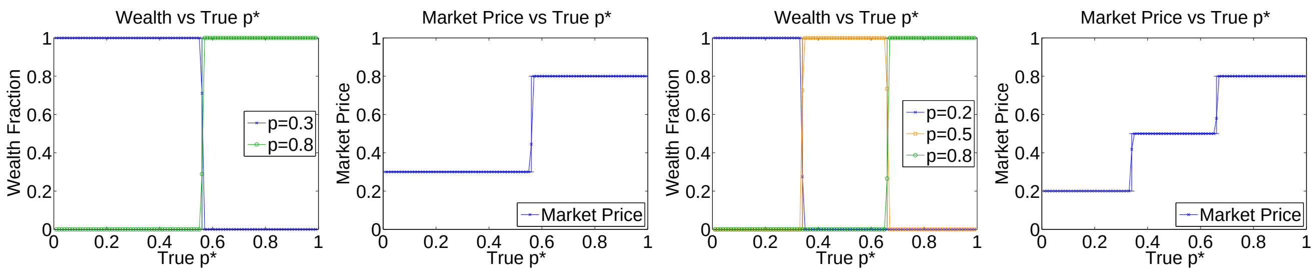

(a) $n = 2$, trader wealth (b) $n = 2$, market price (c) $n = 3$, trader wealth (d) $n = 3$, market price

Figure 2: The extreme case of $c = 1$ (full Kelly betting).

prediction by bringing the market price closer to the objective event probability p^*.

The $c = 1$ case. Next, we investigate the extreme cases in detail. When both traders use Kelly betting, Figures 2(a) and 2(b) show the limiting expected wealth and market price respectively as a function of the objective event probability p^*. We see that only the most accurate trader survives and the market price converges to his belief, as predicted by Theorem 3. In fact, the transition point between trader 1's dominance and trader 2's dominance is observed at $\hat{p} \approx 0.56$, which is very close to the theoretical transition point $p^* \approx 0.560874$ predicted by Theorem 3. The latter is obtained by solving $p^* \log p_1 + (1 - p^*) \log(1 - p_1) = p^* \log p_2 + (1 - p^*) \log(1 - p_2)$. All the above observations also hold for markets with more than two Kelly traders; see, e.g., the limiting expected trader wealth and market price in Figures 2(c) and 2(d) respectively when a trader with belief $p_3 = 0.5$ is added. The endpoints of the region of p^* in which the new trader has the most accurate belief can be checked to coincide with the transition points predicted by Theorem 3.

The $c \approx 0$ case. On the other end of the spectrum, we explore the case of $c = 0.01$ in detail. Figures 3(a) and 3(b) show the limiting expected wealth and market price respectively as a function of the objective event probability p^*. We see that in the extreme cases of $p^* \le p_1$ and $p^* \ge p_2$, only the most accurate trader survives, and the market price converges to his belief, as predicted by Theorem 2. For the remaining case of $p^* \in (p_1, p_2)$, we observe that both traders survive, and the market price consistently coincides with p^*. This matches our conjecture in the previous section accurately.[2] Further, this is in contrast to the case of $c = 1$ where only the most accurate trader survives, and the market price converges to his belief, away from p^*.

For more than two 0.01-Kelly traders, we observe that the market price still matches p^* so far as p^* is strictly between the lowest and the highest of the beliefs; otherwise, Theorem 2 applies. However, this does not allow prediction of unique trader wealth from Proposition 3 anymore. See, e.g., the limiting expected trader wealth and market price in Figures 3(c) and 3(d) respectively when a trader with belief $p_3 = 0.5$ is added. As p^* goes from 0.3 to 0.8,[3] the wealth of the new trader first increases starting from 0, achieves its peak near $p_3 = 0.5$, and then decreases back to 0.

CRRA betting

Recall that the parameter η in CRRA betting strategy denotes the risk aversion of the trader. Hence, increasing η lowers the sizes of the bets that maximize the traders' utility. This corresponds to decreasing c in the c-Kelly betting strategy. We analyze a market with two CRRA traders, trader 1 with belief $p_1 = 0.3$ and trader 2 with belief $p_2 = 0.8$. In this case, we are interested in the limiting expected trader wealth and limiting expected market price as the risk aversion parameter η varies.

[2] If the market price coincides with p^*, then the trader wealth must match our conjecture due to Proposition 3.

[3] If $p^* \notin (0.3, 0.8)$, then all beliefs lie on one side (either above or below) the objective probability and hence, from Theorem 2, only the most accurate belief type survives.

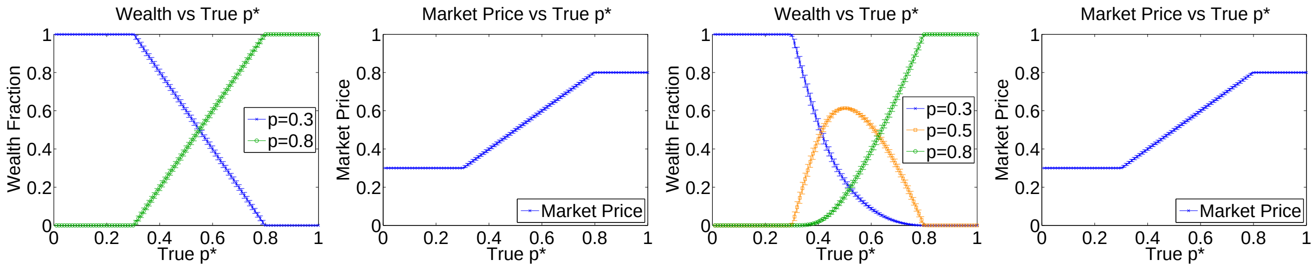

(a) $n = 2$, trader wealth (b) $n = 2$, market price (c) $n = 3$, trader wealth (d) $n = 3$, market price

Figure 3: The extreme case of $c = 0.01$ (fractional Kelly betting with almost zero fraction).

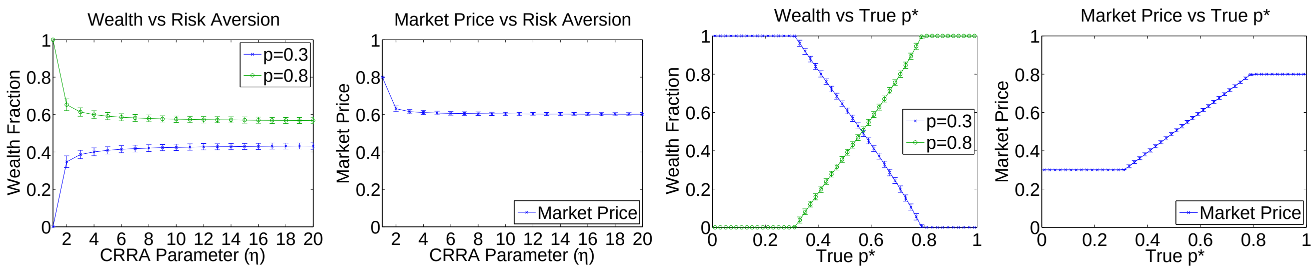

(a) Varying η, trader wealth (b) Varying η, market price (c) $\eta = 20$, trader wealth (d) $\eta = 20$, market price

Figure 4: Market with two CRRA traders.

Varying η. Figures 4(a) and 4(b) show the limiting expected wealth and market price when η varies from 1 to 20 and the objective event probability is fixed at $p^* = 0.6$. We see that this market performs very similarly to the market with fractional Kelly traders. At $\eta = 1$, we recover the Kelly behavior where only the most accurate trader survives and the market price converges to his belief (away from p^*). As η increases, we begin to observe survival of the less accurate trader as well, and the expected market price quickly converges to p^*.

Once again, we conclude that when traders lower the sizes of their bets, the less accurate traders survive, and the market price more accurately reflects the objective event probability. We do not need to analyze the boundary case of $\eta = 1$, because it is identical to Kelly betting. Next, we next analyze the other boundary case of very high risk aversion. Consistent with the previous experiment, we choose $\eta = 20$, which turns out to be high enough to match the behaviour of 0.01-Kelly betting perfectly.

The $\eta = 20$ case. Figures 4(c) and 4(d) show the limiting expected wealth and market price respectively as a function of the objective probability p^*. The behavior is remarkably identical to the market where both traders used 0.01-Kelly betting. In the extreme cases, only the most accurate trader survives, and the expected market price converges to his belief. As p^* varies from p_1 to p_2, the wealth of the two traders decrease/increase linearly while the expected market price consistently coincides with p^*. Again, both traders survive and the accuracy of the market price is astonishingly high because the traders lowered their bet sizes.

Discussion

Our main conceptual contribution is that when traders decrease their bet sizes, the less accurate traders also survive along with the more accurate ones, and this helps improve the accuracy of the market as a whole; in the long run, the market price reflects the objective probability of the event on which bets are placed.

Our findings suggest a number of directions for future research. It would be useful to explore, using more fine-grained experiments, the dependence of limiting trader wealth and price behavior on factors such as the number of participants in the market, the level of heterogeneity in their beliefs, the overall accuracy of beliefs, and the composition of betting strategies. Further theoretical work, including a proof of our conjecture that high levels of risk-aversion imply the survival of heterogeneous belief types, might uncover interesting connections between the risk-aversion an belief of a trader and the survival of the trader.

We also view our model as a stepping stone to more realistic models of prediction markets, e.g., to a model where in every iteration, the objective event probability p^* is sampled from a distribution, and the beliefs of the traders are sampled from a distribution concentrated around p^*. Because the belief of a trader is a fresh noisy estimate of the objective event probability in each iteration, only the betting strategy plays a role in determining the survival of the trader in this case. We believe that our theoretical analysis could provide some insight into analyzing such more complex models. This could deepen our understanding of the conditions under which prediction markets provide accurate forecasts.

Acknowledgments

Sethi and Shah were visiting and partially supported by Microsoft Research New York City during the course of this work.

References

Beygelzimer, A.; Langford, J.; and Pennock, D. M. 2012. Learning performance of prediction markets with Kelly bettors. In *Proceedings of the 11th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1317–1318.

Blume, L., and Easley, D. 1992. Evolution and market behavior. *Journal of Economic Theory* 58(1):9–40.

Blume, L., and Easley, D. 2006. If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets. *Econometrica* 74:929–966.

Chen, K. Y., and Plott, C. 2002. Information aggregation mechanisms: Concept, design and field implementation. Social Science Working Paper no. 1131. Pasadena: CalTech.

Cowgill, B.; Wolfers, J.; and Zitzewitz, E. 2009. Using prediction markets to track information flows: Evidence from Google. Working Paper.

DeLong, B.; Schleifer, A.; Summers, L.; and Waldmann, R. 1990. Noise trader risk in financial markets. *Journal of Political Economy* 98:703–738.

Gjerstad, S. 2004. Risk aversion, beliefs, and prediction market equilibrium. Working Paper, Economic Science Laboratory, University of Arizona.

Kelly, J. 1956. A new interpretation of information rate. *Bell System Technical Journal* 35(4):917–926.

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

Pennock, D. M. 1999. *Aggregating Probabilistic Beliefs: Market Mechanisms and Graphical Representations*. Ph.D. Dissertation, The University of Michigan.

Sandroni, A. 2000. Do markets favor traders able to make accurate predictions? *Econometrica* 68:1303–1342.

Ungar, L.; Mellors, B.; Satopää, V.; Baron, J.; Tetlock, P.; Ramos, J.; and Swift, S. 2012. The good judgment project: A large scale test of different methods of combining expert predictions. AAAI Technical Report FS-12-06.

Wolfers, J., and Zitzewitz, E. 2006. Interpreting prediction market prices as probabilities. NBER Working Paper No. 12200.