
Learning Performance of Prediction Markets with Kelly Bettors

Alina Beygelzimer
IBM Research
Hawthorne, NY
beygel @ us.ibm.com

John Langford, David M. Pennock
Yahoo! Research
New York, NY
{jl,pennockd} @ yahoo-inc.com

Abstract

Kelly betting is an optimal strategy for taking advantage of an information edge in a prediction market, and fractional Kelly is a common variant. We show several consequences that follow by assuming that every participant in a prediction market uses (fractional) Kelly betting. First, the market prediction is a wealth-weighted average of the individual participants' beliefs, where fractional Kelly bettors shift their beliefs toward the market price as if they've seen some fraction of observations. Second, if all fractions are one, the market *learns* at the optimal rate, the market prediction has low log regret to the best individual participant, and when an underlying true probability exists the market converges to the true objective frequency as if updating a Beta distribution. If fractions are less than one, the market converges to a time-discounted frequency. In the process, we provide a new justification for fractional Kelly betting, a strategy widely used in practice for ad-hoc reasons. We propose a method for an agent to learn her own optimal Kelly fraction.

1 Introduction

Consider a gamble on a binary event, say, that Obama will win the 2012 US Presidential election, where x dollars risked earns xb dollars net profit if the gamble pays off. How many dollars x of your wealth should you risk if you believe the probability is p ? The gamble is favorable if $bp - (1-p) > 0$, in which case betting your entire wealth w will maximize your expected profit. However, that's extraordinarily risky: a single stroke of bad luck loses everything. Over the course of many such gambles, the probability of bankruptcy approaches 1. On the other hand, betting a small fixed amount avoids bankruptcy but cannot take advantage of compounding growth.

The *Kelly criteria* prescribes choosing x to maximize the expected compounding growth rate of wealth, or equivalently to maximize the expected logarithm of wealth. Kelly betting is asymptotically optimal, meaning that in the limit over many gambles, a Kelly bettor will grow wealthier than an otherwise identical non-Kelly bettor with probability 1 [1,3,7,16,17].

Assume all agents in a market optimize according to the Kelly principle, where b is selected to clear the market. We consider the implications for the market as a whole and properties of the market odds b or, equivalently, the market probability $p_m = 1/(1+b)$. We show that the market prediction p_m is a wealth-weighted average of the agents' predictions p_i . Over time, the market itself—by reallocating wealth among participants—adapts at the optimal rate with bounded log regret to the best individual agent. When a true objective probability exists, the market converges to it as if properly updating a Beta distribution according to Bayes rule. These results illustrate that there is no “price of anarchy” associated with well-run prediction markets.

We also consider fractional Kelly betting, a lower-risk variant of Kelly betting that's popular in practice but with less theoretical grounding in the appendix. We provide a new justification for fractional

Kelly based on agent confidence. In this case, the market prediction is a confidence-and-wealth-weighted average that empirically converges to a time-discounted version of objective frequency. We propose a method for agents to learn their optimal fraction over time.

2 Kelly betting

When offered b -to-1 odds on an event with probability p , the Kelly-optimal amount to bet is f^*w , where

$$f^* = \frac{bp - (1 - p)}{b}$$

is the optimal fixed fraction of total wealth w to commit to the gamble.

If f^* is negative, Kelly says to avoid betting: expected profit is negative. If f^* is positive, you have an information edge; Kelly says to invest a fraction of your wealth proportional to how advantageous the bet is. In addition to maximizing the growth rate of wealth, Kelly betting maximizes the geometric mean of wealth and asymptotically minimizes the mean time to reach a given aspiration level of wealth [17].

Suppose fair odds of $1/b$ are simultaneously offered on the opposite outcome (e.g., Obama will *not* win the election). If $bp - (1 - p) < 0$, then betting on this opposite outcome is favorable; substituting $1/b$ for b and $1 - p$ for p , the optimal fraction of wealth to bet becomes $1 - p - bp$.

An equivalent way to think of a gamble with odds b is as a prediction market with price $p_m = 1/(1 + b)$. The volume of bet is specified by choosing a quantity q of *shares*, where each share is worth \$1 if the outcome occurs and nothing otherwise. The price represents the cost of one share: the amount needed to pay for a chance to win back \$1. In this interpretation, the Kelly formula becomes

$$f^* = \frac{p - p_m}{1 - p_m}$$

The optimal action for the agent is to trade $q^* = f^*w/p_m$ shares, where $q^* > 0$ is a buy order and $q^* < 0$ is a sell order, or a bet against the outcome.

Note that q^* is the optimum of expected log utility $p \ln((1 - p_m)q + w) + (1 - p) \ln(-p_m q + w)$. This is not a coincidence: Kelly betting is identical to maximizing expected log utility.

3 Market prediction

In order to define the prediction market's performance, we must define its prediction b , or the equilibrium payoff odds reached when all agents are optimizing, and supply and demand are precisely balanced. Recall that the market's probability implied by the odds of b is $p_m = 1/(1 + b)$. We will show that p_m is $\sum_i w_i p_i$.

3.1 Payout balance

The first approach we'll use is payout balance: the amount of money at risk must be the same as the amount payed out.

To see this, recall that $f_i^* = (p_i - p_m)/(1 - p_m)$ for $p_i > p_m$. For $p_i < p_m$, Kelly betting prescribes taking the other side of the bet, with fraction

$$\frac{(1 - p_i) - (1 - p_m)}{1 - (1 - p_m)} = \frac{p_m - p_i}{p_m}$$

So the market equilibrium occurs at the point p_m where the payout is equal to the payin. If the event occurs, the payin is

$$(1 + b) \sum_{i:p_i > p_m} \frac{p_i - p_m}{1 - p_m} w_i = \frac{1}{p_m} \sum_{i:p_i > p_m} \frac{p_i - p_m}{1 - p_m} w_i.$$

Thus we want

$$\begin{aligned} \frac{1}{p_m} \sum_{i:p_i > p_m} \frac{p_i - p_m}{1 - p_m} w_i &= \sum_{i:p_i > p_m} \frac{p_i - p_m}{1 - p_m} w_i + \sum_{i:p_i < p_m} \frac{p_m - p_i}{p_m} w_i, \quad \text{or} \\ \frac{1 - p_m}{p_m} \sum_{i:p_i > p_m} \frac{p_i - p_m}{1 - p_m} w_i &= \sum_{i:p_i < p_m} \frac{p_m - p_i}{p_m} w_i, \quad \text{or} \\ \sum_{i:p_i > p_m} (p_i - p_m) w_i &= \sum_{i:p_i < p_m} (p_m - p_i) w_i, \quad \text{or equivalently} \\ \sum_i p_i w_i &= \sum_i p_m w_i. \end{aligned}$$

Using $\sum_i w_i = 1$, we get

$$p_m = \sum_i p_i w_i.$$

In other words, the market clearing price is the wealth-weighted average prediction of participants.

3.2 Log utility maximization

An alternate derivation of the market prediction utilizes the fact that Kelly betting is equivalent to maximizing expected log utility. Let $q = x(b + 1)$ be the gross profit of an agent who risks x dollars, or in prediction market language the number of shares purchased. Then expected log utility is

$$E[U(q)] = p \ln((1 - p_m)q + w) + (1 - p) \ln(-p_m q + w).$$

The optimal q that maximizes $E[U(q)]$ is

$$q(p_m) = \frac{w}{p_m} \cdot \frac{p - p_m}{1 - p_m}. \quad (1)$$

Proposition 1. *In a market of agents each with log utility and initial wealth w , the competitive equilibrium price is*

$$p_m = \sum_i w_i p_i \quad (2)$$

where we assume $\sum_i w_i = 1$, or w is normalized wealth not absolute wealth.

Proof. These prices satisfy $\sum_i q_i = 0$, the condition for competitive equilibrium (supply equals demand), by substitution. \square

This result can be seen as a simplified derivation of that by Rubinstein [13,14,15] and is also discussed by Pennock and Wellman [11,10] and Wolfers and Zitzewitz [18].

4 Learning Prediction Markets

Individual participants may have varying prediction qualities and individual markets may have varying odds of payoff. What happens to the wealth distribution and hence the quality of the market prediction over time? We show next that the market *learns* optimally for two well understood senses of optimal.

4.1 Wealth redistributed according to Bayes Law

In an individual round, if an agent's belief is $p_i > p_m$, then they bet $\frac{p_i - p_m}{1 - p_m} w_i$ and have a total wealth afterward dependent on outcome y according to:

1. If $y = 1$, $(\frac{1}{p_m} - 1) \frac{p_i - p_m}{1 - p_m} w_i + w_i = \frac{p_i}{p_m} w_i$
2. If $y = 0$, $(-1) \frac{p_i - p_m}{1 - p_m} w_i + w_i = \frac{1 - p_i}{1 - p_m} w_i$

Similarly if $p_i < p_m$, we get:

1. If $y = 1$, $(-1)^{\frac{p_m - p_i}{p_m}} w_i + w_i = \frac{p_i}{p_m} w_i$
2. If $y = 0$, $(\frac{1}{1 - p_m} - 1)^{\frac{p_m - p_i}{p_m}} w_i + w_i = \frac{1 - p_i}{1 - p_m} w_i$

which is identical.

If we treat the prior probability that agent i is correct as w_i , Bayes' law states that the posterior probability of choosing agent i is

$$P(i | y = 1) = \frac{P(y = 1 | i)P(i)}{P(y = 1)} = \frac{p_i w_i}{p_m} = \frac{p_i w_i}{\sum_i p_i w_i},$$

which is precisely the wealth computed above for the $y = 1$ outcome. The same holds true when $y = 0$, and so Kelly bettors redistribute wealth according to Bayes' law.

4.2 Market Sequences

It is well known that Bayes' law is the correct approach for integrating evidence into a belief distribution, which shows that Kelly betting agents optimally summarize all past information if the true behavior of the world was drawn from the prior distribution of wealth.

Often these assumptions are too strong—the world does not behave according to the prior on wealth, and it may act in a manner completely different from any one single expert. In that case, a standard analysis from learning theory shows that the market place has low *regret*, performing almost as well as the best market participant.

For any particular sequence of markets we have a sequence p_t of market predictions and $y_t \in \{0, 1\}$ of market outcomes. We measure the accuracy of a market according to log loss as

$$L \equiv \sum_{t=1}^T I(y_t = 1) \log \frac{1}{p_t} + I(y_t = 0) \log \frac{1}{1 - p_t}.$$

Similarly, we measure the quality of market participant making prediction p_{it} as

$$L_i \equiv \sum_{t=1}^T I(y_t = 1) \log \frac{1}{p_{it}} + I(y_t = 0) \log \frac{1}{1 - p_{it}}.$$

So after T rounds, the total wealth of player i is

$$w_i \prod_{t=1}^T \left(\frac{p_{it}}{p_t} \right)^{y_t} \left(\frac{1 - p_{it}}{1 - p_t} \right)^{1 - y_t},$$

where w_i is the starting wealth. We next prove a well-known theorem for learning in the present context (see for example [4]).

Theorem 1. *For all sequences of participant predictions p_{it} and all sequences of revealed outcomes y_t ,*

$$L \leq \min_i L_i + \ln \frac{1}{w_i}.$$

This theorem is extraordinarily general, as it applies to *all* market participants and *all* outcome sequences, even when these are chosen adversarially. It states that even in this worst case situation, the market performs only $\ln 1/w_i$ worse than the best market participant i .

Proof. Initially, we have that $\sum_i w_i = 1$. After T rounds, the total wealth of any participant i is given by

$$w_i \prod_{t=1}^T \left(\frac{p_{it}}{p_t} \right)^{y_t} \left(\frac{1 - p_{it}}{1 - p_t} \right)^{1 - y_t} = w_i e^{L - L_i} \leq 1,$$

where the last inequality follows from wealth being conserved. Thus $\ln w_i + L - L_i \leq 0$, yielding

$$L \leq L_i + \ln \frac{1}{w_i}.$$

□

References

1. L. Breiman. Optimal gambling systems for favorable games. In *Berkeley Symposium on Probability and Statistics, I*, pages 65–78, 1961.
2. N. Cesa-Bianchi, Y. Freund, D. Helbold, D. Haussler, R. Schapire, and M. Warmuth. How to use expert advice. *Journal of the ACM*, 44(3):427–485, 1997.
3. Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory, Second Edition*. Wiley-Interscience, New Jersey, 2006.
4. Y. Freund, R. Schapire, Y. Singer, and M. Warmuth. Using and combining predictors that specialize. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing*, pages 334–343, 1997.
5. John D. Geanakoplos and Heracles M. Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982.
6. Sanford J. Grossman. An introduction to the theory of rational expectations under asymmetric information. *Review of Economic Studies*, 48(4):541–559, 1981.
7. J.L. Kelly. A new interpretation of information rate. *Bell System Technical Journal*, 35:917–926, 1956.
8. Paul Milgrom and Nancy L. Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.
9. Peter A. Morris. An axiomatic approach to expert resolution. *Management Science*, 29(1):24–32, 1983.
10. David M. Pennock. *Aggregating Probabilistic Beliefs: Market Mechanisms and Graphical Representations*. PhD thesis, University of Michigan, 1999.
11. David M. Pennock and Michael P. Wellman. A market framework for pooling opinions. Technical Report 2001-081, NEC Research Institute, 2001.
12. Emilio Rosenblueth and Mario Ordaz. Combination of expert opinions. *Journal of Scientific and Industrial Research*, 51:572–580, 1992.
13. Mark Rubinstein. An aggregation theorem for securities markets. *Journal of Financial Economics*, 1(3):225–244, 1974.
14. Mark Rubinstein. Securities market efficiency in an Arrow-Debreu economy. *American Economic Review*, 65(5):812–824, 1975.
15. Mark Rubinstein. The strong case for the generalized logarithmic utility model as the premier model of financial markets. *Journal of Finance*, 31(2):551–571, 1976.
16. Edward O. Thorp. Optimal gambling systems for favorable games. *Review of the International Statistical Institute*, 37:273–293, 1969.
17. Edward O. Thorp. The Kelly criterion in blackjack, sports betting, and the stock market. In *International Conference on Gambling and Risk Taking*, Montreal, Canada, 1997.
18. Justin Wolfers and Eric Zitzewitz. Interpreting prediction market prices as probabilities. Technical Report 12200, NBER, 2006.

A Fractional Kelly Betting

Fractional Kelly betting says to invest a smaller fraction λf^* of wealth for $\lambda < 1$. Fractional Kelly is usually justified on an ad-hoc basis as either (1) a risk-reduction strategy, since practitioners often view full Kelly as too volatile, or (2) a way to protect against an inaccurate belief p , or both [17]. Here we derive an alternate interpretation of fractional Kelly. In prediction market terms, the fractional Kelly formula is

$$\lambda \frac{p - p_m}{1 - p_m}.$$

With some algebra, fractional Kelly can be rewritten as

$$\frac{p' - p_m}{1 - p_m}$$

where

$$p' = \lambda p + (1 - \lambda)p_m. \quad (3)$$

In other words, λ -fractional Kelly is precisely equivalent to full Kelly with revised belief $\lambda p + (1 - \lambda)p_m$, or a weighted average of the agent's original belief and the market's belief. In this

light, fractional Kelly is a form of confidence weighting where the agent mixes between remaining steadfast with its own belief ($\lambda = 1$) and acceding to the crowd and taking the market price as the true probability ($\lambda = 0$). The weighted average form has a Bayesian justification if the agent has a Beta prior over p and has seen t independent Bernoulli trials to arrive at its current belief. If the agent envisions that the market has seen t' trials, then she will update her belief to $\lambda p + (1 - \lambda)p_m$, where $\lambda = t/(t + t')$ [9,12]. The agent's posterior probability given the price is a weighted average of its prior and the price, where the weighting term captures her perception of her own confidence, expressed in terms of the independent observation count seen as compared to the market.

B Market prediction with fractional Kelly

When agents play fractional Kelly, the competitive equilibrium price is

$$p_m = \frac{\sum_i \lambda_i w_i p_i}{\sum_i \lambda_i w_i}. \quad (4)$$

Prices retain the form of a weighted average, but with weights proportional to the product of wealth and self-assessed confidence.

C Market model

Suppose that we have a prediction market with N participants, where participant i has a starting wealth w_i with $\sum_i w_i = 1$. Each participant i uses Kelly betting to determine the fraction f_i^* of their wealth bet, depending on their predicted probability p_i .

We model the market as an auctioneer matching supply and demand, taking no profit and absorbing no loss. We adopt a competitive equilibrium concept, meaning that agents are “price takers”, or do not consider their own effect on prices if any. Agents optimize according to the current price and do not reason further about what the price might reveal about the other agents' information. An exception of sorts is the fractional Kelly setting, where agents do consider the market price as information and weight it along with their own.

A market is in competitive equilibrium at price p_m if all agents are optimizing and $\sum_i q_i^* = 0$, or every buy order and sell order are matched.

D Example: Market dynamics with stationary objective frequency

The worst-case bounds above hold even if event outcomes are chosen by a malicious adversary. In this section, we examine how the market performs when the objective frequency of outcomes is unknown though stationary.

The market consists of a single bet repeated over the course of T periods. Unbeknown to the agents, each event unfolds as an independent Bernoulli trial with probability of success π . At the beginning of time period t , the realization of event E_t is unknown and agents trade until equilibrium. Then the outcome is revealed, and the agents' holdings pay off accordingly. As time period $t + 1$ begins, the outcome of E_{t+1} is uncertain. Agents bet on the $t + 1$ period event until equilibrium, the outcome is revealed, payoffs are collected, and the process repeats.

In an economy of Kelly bettors, the equilibrium price is a wealth-weighted average (2). Thus, as an agent accrues relatively more earnings than the others, its influence on price increases. In the next two subsections, we examine how this adaptive process unfolds; first, with full-Kelly agents and second, with fractional Kelly agents. In the former case, prices react exactly as if the market were a single agent updating a Beta distribution according to Bayes' rule.

Market dynamics with full-Kelly agents

Figure 1.a plots the price over fifty time periods, in a market composed of one hundred Kelly agents, initial wealth $w_i = 1$, and p_i generated randomly and uniformly on $(0, 1)$. In this simulation $\pi = 0.5$. For comparison, the figure also shows the *observed frequency*, or the number of times that E

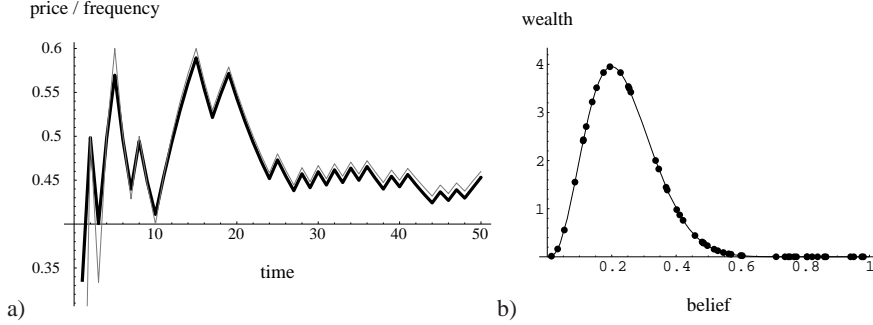


Fig. 1. a) Price (black line) versus the observed frequency (gray line) of the event over fifty time periods. The market consists of one hundred full-Kelly agents with initial wealth $w_i = 1$. b) Wealth after fifteen time periods versus belief for fifty Kelly agents. The event has occurred in three of the fifteen trials. The solid line is the posterior Beta distribution consistent with the observation of three successes in fifteen independent Bernoulli trials.

has occurred divided by the number of periods. The market price tracks the observed frequency extremely closely. Note that price changes are due entirely to a transfer of wealth from inaccurate agents to accurate agents, who then wield more power in the market; individual beliefs remain fixed.

Figure 1.b illustrates the nature of this wealth transfer. The graph provides a snapshot of agents' wealth versus their belief p_i after period 15. In this run, E has occurred in three out of the fifteen trials. The maximum in wealth is near $3/15$. The solid line in the figure is a Beta distribution with parameters $3 + 1$ and $12 + 1$, scaled only by a constant. This distribution is precisely the posterior probability of success that results from the observation of 3 successes out of 15 independent Bernoulli trials, when the prior probability of success is uniform on $(0,1)$. The fit is essentially perfect, and can be proved in the limit since the Beta distribution is conjugate to the Binomial distribution under Bayes Law.

Although individual agents are not adaptive, the market's composite agent computes a proper Bayesian update. Specifically, wealth is reallocated proportionally to a Beta distribution corresponding to the observed number of successes and trials, and price is approximately the expected value of this Beta distribution.¹ Moreover, this correspondence holds regardless of the number of successes or failures, or the temporal order of their occurrence. A kind of collective Bayesianity *emerges* from the interactions of the group.

We also find empirically that, even if not all agents are Kelly bettors, among those that are, wealth is still redistributed according to Bayes rule.

Market dynamics with fractional Kelly agents

In this section, we consider fractional Kelly agents who, as we saw in Section 2, behave like full Kelly agents with belief $\lambda p + (1 - \lambda)p_m$. Figure 2.a graphs the dynamics of price in an economy of such agents, along with the observed frequency. Over time, the price remains significantly more volatile than the frequency, which converges toward $\pi = 0.5$. Below, we characterize the transfer of wealth that precipitates this added volatility; for now concentrate on the price signal itself. Inspecting Figure 2.a, price changes still exhibit a marked dependence on event outcomes, though at any given period the effect of recent history appears magnified, and the past discounted, as compared with the observed frequency. Working from this intuition, we attempt to fit the data to an appropriately modified measure of frequency. Define the *discounted frequency* at period n as

$$d_n = \frac{\sum_{t=1}^n \gamma^{n-t} (1_{E(t)})}{\sum_{t=1}^n \gamma^{n-t} (1_{E(t)}) + \sum_{t=1}^n \gamma^{n-t} (1_{\overline{E}(t)})}, \quad (5)$$

where $1_{E(t)}$ is the indicator function for the event at period t , and γ is the *discount factor*. Note that $\gamma = 1$ recovers the standard observed frequency.

¹ As t grows, this expected value rapidly approaches the observed frequency plotted in Figure 1.

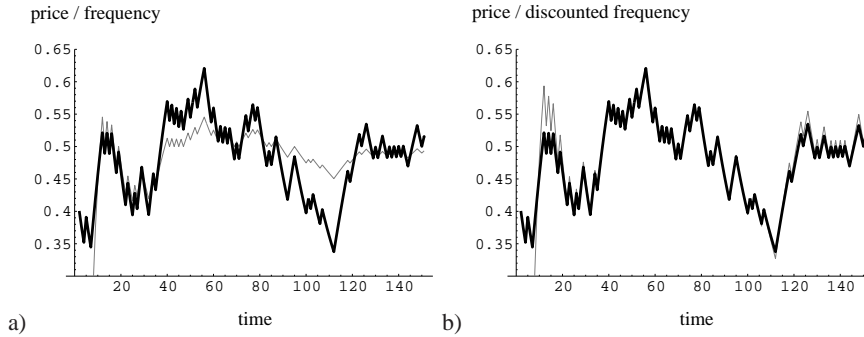


Fig. 2. a) Price (black line) versus observed frequency (gray line) over one hundred fifty time periods for one hundred agents with Kelly fraction $\lambda = 0.2$. As the frequency converges to $\pi = 0.5$, the price remains volatile. b) Price (black line) versus discounted frequency (gray line), with discount factor $\gamma = 0.96$, for the same experiment as (a).

Figure 2.b illustrates a very close correlation between discounted frequency, with $\gamma = 0.96$ (hand tuned), and the same price curve of Figure 2.a. While standard frequency provides a provably good model of price dynamics in an economy of full-Kelly agents, discounted frequency (5) appears a better model for fractional Kelly agents.

E Learning the Kelly fraction

In theory, a rational agent playing against rational opponents should set their Kelly fraction to $\lambda = 0$, since, in a rational expectations equilibrium [6], the market price is by definition at least as informative as any agent’s belief. This is the crux of the no-trade theorems [8]. Despite the theory [5], people do agree to disagree in practice and, simply put, trade happens. Still, placing substantial weight on the market price is often prudent. For example, in an online prediction contest called ProbabilitySports, 99.7% of participants were outperformed by the unweighted average predictor, a typical result.

In this light, fractional Kelly can be seen as an experts algorithm [2] with two experts: yourself and the market. We propose dynamically updating λ according to standard experts algorithm logic: When you’re right, you increase λ appropriately; when you’re wrong, you decrease λ . This gives a long-term procedure for updating λ that guarantees:

- You won’t do too much worse than the market (which by definition earns 0)
- You won’t do too much worse than Kelly betting using your original prior p

For example, if you allocate an initial weight of 0.5 to your predictions and 0.5 to the market’s prediction, then the regret guarantee of section 4.2 implies that at most half of all wealth is lost.

F Speculations

We’ve shown something intuitively appealing here: that self-interested agents with log wealth utility create markets which learn to have small regret according to log loss. There are two distinct “log”s in this statement, and it’s appealing to consider what happens when we vary these. When agents have some utility other than log wealth utility, can we alter the structure of a market so that the market dynamics make the market price have low log loss regret? And similarly if we care about some other loss—such as squared loss, 0/1 loss, or a quantile loss, can we craft a marketplace such that log wealth utility agents achieve small regret with respect to these other losses?

What happens in a market place without Kelly bettors? This can’t be described in general, although one special case is relevant: when one of the bettors acts according to Kelly and the others in some more irrational fashion. In this case, the basic Kelly guarantee implies that the Kelly bettor will come to dominate non-Kelly bettors with equivalent or worse log loss. If non-Kelly agents have a better

log loss, the behavior can vary, possibly imposing greater regret on the marketplace if the Kelly bettor accrues the wealth despite a worse prediction record. For this reason, it may be desirable to make Kelly betting an explicit option in prediction markets.