

INSTITUTE
OF ECONOMICS



Scuola Superiore
Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM

WORKING PAPER SERIES

Far from the Madding Crowd: Collective Wisdom in Prediction Markets

Giulio Bottazzi [°]
Daniele Giachini [°]

[°] Institute of Economics, Scuola Superiore Sant'Anna, Pisa, Italy

2016/14

April 2016

ISSN(ONLINE) 2284-0400

Far from the Madding Crowd: Collective Wisdom in Prediction Markets*

Giulio Bottazzi[†] and Daniele Giachini[‡]

Istituto di Economia, Scuola Superiore Sant'Anna
Piazza Martiri della Libertà 33, 56127 Pisa - Italy

April 4, 2016

Abstract

We investigate market selection and bet pricing in a simple Arrow security economy which we show is equivalent to the repeated prediction market models studied in the literature. We derive the condition for long run survival of more than one agent (the crowd) and quantify the information content of prevailing prices in the case of two fractional Kelly traders with heterogeneous beliefs. It turns out that, apart some non-generic situations, prices do not converge, neither almost surely nor on average, to true probabilities. Nor are they always nearer to the truth than the beliefs of all surviving agents. Moreover, we show that by adapting their beliefs to past prices, agents further decrease the agreement between market prices and true probabilities.

JEL Classification: C60, D53, G11, G12

*If there is some good idea in this paper, it probably emerged in some discussion with Pietro Dindo. In any case, we retain the exclusive paternity of all the bad ones. We thank Filippo Massari and Remco Zwinkels for useful suggestions. We gratefully acknowledge the funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 640772 - DOLFINS.

[†]bottazzi@sssup.it

[‡]d.giachini@sssup.it

Far from the madding crowd's ignoble strife
Their sober wishes never learn'd to stray;

Elegy Written in a Country Churchyard
Thomas Gray

1 Introduction

In his 2004 best-seller book “*The Wisdom of Crowds*” James Surowiecki provides many examples and anecdotes of how, by aggregating the beliefs of independently deciding individuals, one gets rather accurate predictions about some uncertain events. Thus, a form of *collective intelligence* seems to emerge when different opinions are collected and aggregated, letting a Crowd be wiser than any of its members. The point raised by Surowiecki is not new in the economic literature: it can be traced back to the classical article by Hayek (1945) “*The Use of Knowledge in Society*”, where the famous economist supports the thesis that a decentralized economy is successful in aggregating the large amount of sparse pieces of information through prices. That is, the price that comes out from the market interaction of agents with heterogeneous and incomplete information condenses all the relevant knowledge and provides a means by which agents can take the right decisions.

The relevance of the informative role of markets is pervasive in the financial literature and practice, from the valuation of equity to the option pricing, and it is possibly epitomized in the so called prediction markets. These are markets where bets on binary events are exchanged.¹ The Iowa Electronic Market (IEM), operated by the University of Iowa, is the most prominent example: in that market people can bet on events like Presidential Elections or Congressional Control by means of future contracts that pay 1 dollar if the event is realized and 0 otherwise. In this frame it is common to interpret the price of the bet as the probability the Crowd assigns to the realization of the event. About this point, a recent article appeared on Science Magazine and coauthored by several famous economists (Arrow et al., 2008) supported the development of prediction markets exactly because of their ability to provide accurate evaluations of the likelihood of uncertain events.

However, the formal investigation of the precision and reliability of prediction markets predictions is rather recent. The early contributions investigated whether the price that comes out from a single round of betting matches the average belief of the population. Gjerstad (2005) finds that, when agents are risk averse and the distribution of beliefs is plausible, then the equilibrium price is very near

¹Trading in some financial instruments is basically equivalent to betting on the outcome of an uncertain event. For instance, the price of Credit Default Swaps is generally recognized as a prediction about the probability of default of the underlying bond issuer.

to population average belief. Wolfers and Zitzewitz (2006) confirm such result extending it to a broad class of models. Both papers agree on the fact that the estimation of the average beliefs provided by prediction market prices can be biased, but they quantify that bias as small. Conversely Manski (2006) shows that under risk neutrality the divergence between the price and the average belief can be large. He and Treich (2012) summarize and complement the previous results providing conditions for equilibrium price to match the average belief: for every possible beliefs distribution it is necessary and sufficient that agents share a logarithmic utility, or, alternatively, for every possible strictly concave utility, it is necessary and sufficient that the beliefs distribution is symmetric around one half.

In recent years, the attention has increasingly switched to the case of repeated prediction markets, which share more similarities with other financial market models. Beygelzimer et al. (2012) and Kets et al. (2014), in a dynamic version of the previous models, assume that in every time step a prediction market on a binary event is hold. The success probability of the binary event, constant over time, is supposed unknown. Traders have a certain amount of wealth and in every period they decide, according to their beliefs, how much to invest in the prediction market and what position they shall take. In every time step a central auctioneer collects the orders and establishes the price accordingly. At the end of the period the outcome of the event is revealed, agents' wealth is updated and the process is repeated. This dynamic framework has the advantage of providing an objective probability (the one that drives the Bernoulli trials) that acts as a benchmark for evaluating the correctness of market price and agents' beliefs. Beygelzimer et al. (2012) find that, under the assumption that agents bet according to the Kelly criterion (Kelly, 1956), the market price adapts to the success probability at optimal rate (i.e. in a Bayesian fashion) and provides a prediction that is only slightly worse than the best agent's one. Kets et al. (2014) go further and show that with Kelly investors only the agents with the most accurate beliefs survives in the long run and, consequently, market prices converge there. Based on extensive numerical simulations, they also suggest that if Kelly traders bet only a small fraction of their wealth, according to the so-called "fractional Kelly" rule (see for instance (MacLean et al., 1992, 2004, 2005)), then more than one agent will survive in the long run and the market is efficient, as the expected price seems to converge to the true probability of the binary event.

Exploiting the equivalence of prediction markets with Arrow security pure exchange economies² we use the general results in Bottazzi and Dindo (2014, 2015) to derive the conditions for the survival of multiple agents. In particular, and following Kets et al. (2014), we will focus on the case of two fractional Kelly traders.

²This is the origin of the similarity of the results in Kets et al. (2014) with those in Blume and Easley (1992)

The simple case of two agents might seem simplistic, but we will see that it is complicated enough to account for a sufficient range of possible outcomes and it already provides important hints on the expected behavior of the market also in more complicated situations. Fractional Kelly investing relates to the traditional framework of expected utility maximization if one assumes that agents myopically maximize a CRRA utility with price dependent relative risk aversion coefficient. We exploit this analogy to discuss the role of risk aversion in the survival of traders and in the efficiency of the market.

We generally find that a difference between the expected price and the success probability emerges and is persistent, so that the average price is neither a consistent nor an unbiased estimator of the success probability. It turns out that for reasonable levels of risk aversion, this bias is in fact quite big if compared to the observed price fluctuations. We also show that there are generic situations in which the average prevailing price is a worst estimate of the true probability than the belief of one of the trader. Finally, we show that also when some useful information is revealed in prices, if one agent tries to exploit it to obtain better predictions, the information disappear and the overall informative efficiency of the market is impaired, with a mechanism strongly reminiscent of the detrimental effect that *social influence* has on the *Wisdom of Crowds*, as suggested in Surowiecki (2004) and experimentally identified by Lorenz et al. (2011).

2 The Model

Consider a discrete time economy populated by N agents who repeatedly bet on a binary event. Agents can choose to gamble on the occurrence of the event or against it. The amount of wealth which is not bet is considered invested in a riskless security. Without loss of generality we assume that no interest is payed on the riskless investment. The total amount bet is redistributed among the winners proportionately to the amount they have bet, that is according to the procedure commonly know as *parimutuel*. The risky bet is based on an independent Bernoulli trial s_t with success probability π^* unknown to the agents, $s_t = 1$ meaning that at time t the event occurred and $s_t = 0$ that it did not. At every time t every agent i has to decide the fraction of wealth $b_t^i \in [0, 1]$ bet on the uncertain outcome and the side of the bet, $\sigma_t^i \in \{0, 1\}$. The fraction of wealth $1 - b_t^i$, considered invested in a riskless security, carries over to the next time step. Let w_t^i denote the wealth of agent i at time t , then given the parimutuel procedure, it is immediate to see that the evolution of individual wealth reads

$$w_t^i = (1 - b_t^i)w_{t-1}^i + \delta_{s_t, \sigma_t^i} w_{t-1}^i b_t^i \frac{\sum_{j=1}^N b_t^j w_{t-1}^j}{\sum_{j=1}^N \delta_{s_t, \sigma_t^j} b_t^j w_{t-1}^j}$$

where $\delta_{a,b}$ stands for the Kronecker delta: it is equal to 1 if $a = b$ and 0 otherwise. The parimutuel procedure simply redistributes the wealth among bettors, but does not change its total amount so that, without loss of generality, we can assume $\sum_{i=1}^N w_t^i = 1$.

This is the model typically considered in the contributions discussed before. We start by showing that this model is equivalent to a discrete time economy in which two Arrow securities are exchanged: the first security pays 1 if the event occur and zero otherwise, the second does the opposite. Let $p_t = \sum_{j=1}^N \delta_{1,\sigma_t^j} b_t^j w_{t-1}^j / \sum_{j=1}^N b_t^j w_{t-1}^j$. This quantity is the fraction of the amount bet on the occurrence of the event over the total amount bet, thus $p_t \in [0, 1]$. It can be thought as the “price” of the first Arrow security. The price of the second Arrow security is, consequently, $1 - p_t$. Moreover, define the investment function

$$\alpha_t^i = (1 - b_t^i) p_t + \sigma_t^i b_t^i .$$

Notice that, by definition, $\alpha_t^i \in [0, 1]$. The individual wealth evolution can then be rewritten as

$$w_t^i = \begin{cases} \frac{\alpha^i(p_t)}{p_t} w_{t-1}^i & \text{if } s_t = 1 , \\ \frac{1 - \alpha^i(p_t)}{1 - p_t} w_{t-1}^i & \text{if } s_t = 0 . \end{cases} \quad (1)$$

Thus, the evolution of wealth is equivalent to that of an agent investing a fraction α_t^i of his wealth in the first Arrow security and the remaining fraction $1 - \alpha_t^i$ in the second. Indeed, it is trivially verified that the price of the securities is set by the market clearing condition

$$p_t = \sum_{i=1}^N \alpha^i(p_t) w_{t-1}^i . \quad (2)$$

Following Gjerstad (2005) and Kets et al. (2014), we assume that agents bet according to constant rules that depend on the total wealth wagered on each outcome, that is on the contemporaneous price, $b_t^i = b^i(p_t)$ and $\sigma_t^i = \sigma^i(p_t)$. These assumptions imply that $\alpha_t^i = \alpha^i(p_t)$. With this further specification, (2) and (1) describe a stochastic system which is a particular case of the one analyzed in Bottazzi and Dindo (2014).

More specifically, the literature on repeated betting assumes that each agent possesses an individual constant belief π^i , in general different from the truth π^* , about the probability of occurrence of the event and invests according to a fractional Kelly rule, described by the simple function

$$\alpha^i(p_t) = c^i \pi^i + (1 - c^i) p_t , \quad (3)$$

with $c^i \in (0, 1]$. This investment rule is a generalization of the Kelly rule and is defined as a linear combination of the individual belief and the market price, with a “mixing” parameter c^i . It is rather common in the literature (MacLean et al., 1992, 2004, 2005, 2010; MacLean and Ziemba, 1999; Ziemba, 2003; Thorp, 2006; Beygelzimer et al., 2012; Kets et al., 2014) and can be considered a good approximation of the behavior of a risk averse agent. In particular (3) is the optimal investment rule for myopic expected utility maximizer with CRRA utility function and a price dependent risk aversion coefficient (see Appendix A for the derivation)

$$\gamma_t^i(p_t) = \frac{\log \frac{\pi^i}{1 - \pi^i} - \log \frac{p_t}{1 - p_t}}{\log \frac{c^i \pi^i + (1 - c^i) p_t}{1 - c^i \pi^i - (1 - c^i) p_t} - \log \frac{p_t}{1 - p_t}}. \quad (4)$$

Notice that $\partial \gamma_{i,t} / \partial c^i < 0$ and $c^i = 1$ implies $\gamma_t^i = 1$. Thus, if $c^i = 1$ we recover the case of a Kelly trader, which is equivalent to a log-utility investor, while considering $0 < c^i < 1$, one gets an agent with a relative risk aversion coefficient higher than 1 and increasing when the value of c^i decreases.

3 Pairwise Comparison

Let us focus on the two-agent model, $N = 2$. As we will see in a moment, the notion of “crowd” and “wisdom”, which are the core of the present analysis, can be effectively discussed also in this simple case. Given the conservation of aggregate wealth, with two agents the stochastic system is one dimensional and we simplify the notation setting $w_t = w_t^1$, such that $w_t^2 = 1 - w_t$. We consider investment rules as in (3) and without loss of generality we assume $\pi^1 < \pi^2$. Thus, in terms of (w_t, p_t) , the dynamics of the economy is described by the following³

$$w_t = \begin{cases} \frac{c^1 \pi^1 + (1 - c^1) p_t}{p_t} w_{t-1} & \text{if } s_t = 1 \\ \frac{1 - c^1 \pi^1 - (1 - c^1) p_t}{1 - p_t} w_{t-1} & \text{if } s_t = 0 \end{cases} \quad (5)$$

together with the market clearing condition

$$p_t = \frac{c^1 \pi^1 w_{t-1} + c^2 \pi^2 (1 - w_{t-1})}{c^1 w_{t-1} + c^2 (1 - w_{t-1})}. \quad (6)$$

³Under these assumptions, the model matches exactly the example provided in Bottazzi and Dindo (2014), Section 2.

Then we say that agent 1 dominates (or, equivalently, agent 2 vanishes) if

$$\lim_{t \rightarrow \infty} w_t = 1 \text{ almost surely (a.s.)}, \quad (7)$$

agent 2 dominates (or, equivalently, agent 1 vanishes) if

$$\lim_{t \rightarrow \infty} w_t = 0 \text{ a.s.} \quad (8)$$

and both agents survive if

$$\limsup_{t \rightarrow \infty} w_t > 0 \text{ a.s.} \quad \wedge \quad \liminf_{t \rightarrow \infty} w_t < 1 \text{ a.s.} \quad (9)$$

Thus, in this simple model, what we identify with the persistent existence of a crowd is nothing else than the asymptotic survival of both agents.

3.1 Reconsidering the Crowd

Since the repeated prediction market considered here is a particular short-lived assets market, we can exploit the analysis put forward in Bottazzi and Dindo (2014) to investigate its long-run dynamics, providing an analytically background for, and extending, the numerical results in Kets et al. (2014). First of all, notice that $p_t \in [\pi^1, \pi^2]$ and that $p = \pi^i$ with $i = 1, 2$ are deterministic fixed points of the dynamics. That is, if $p_t = \pi^i$, then it will be $p_\tau = \pi^i$ for any $\tau > t$. In the language of Bottazzi and Dindo (2014), these are two Market Selection Equilibria (MSE). The discussion of this section relies on how the system behaves in the proximity of these two MSE. Define the relative entropy of strategy α^i at price p_t with respect to the true probability measure of the underlying Bernulli process as

$$I_{\pi^*}(\alpha^i(p_t)) = \pi^* \log \frac{\pi^*}{\alpha^i(p_t)} + (1 - \pi^*) \log \frac{1 - \pi^*}{1 - \alpha^i(p_t)}.$$

Then from (5) it is immediate to see that the conditional expect drift of the log-difference of the individual wealth

$$\begin{aligned} \mu(p_t) &= \mathbb{E} \left[\log \frac{w_t}{w_{t-1}} - \log \frac{1 - w_t}{1 - w_{t-1}} \middle| p_t \right] = \\ &\pi^* \log \frac{\alpha^1(p_t)}{\alpha^2(p_t)} + (1 - \pi^*) \log \frac{1 - \alpha^1(p_t)}{1 - \alpha^2(p_t)} = I_{\pi^*}(\alpha^2(p_t)) - I_{\pi^*}(\alpha^1(p_t)), \end{aligned} \quad (10)$$

is equal to the difference of the relative entropy of the two strategies. The wealth of the strategy with zero relative entropy, the one of a Kelly trader ($c = 1$) with correct beliefs, never decreases in expectation. This is the optimal strategy for an




Figure 1: Local Stability. Left panel: $\pi^* = 0.5$, $\pi^1 = 0.6$, $c^1 = 0.9$, $\pi^2 = 0.8$ and $c^2 = 0.3$. Right panel: $\pi^* = 0.5$, $\pi^1 = 0.2$, $c^1 = 0.5$, $\pi^2 = 0.4$ and $c^2 = 0.8$.

agent, irrespective of what the other agents does, but it unrealistically requires the precise knowledge of the probability π^* . More generally, if for some p_t it is $I_{\pi^*}(\alpha^1(p_t)) > I_{\pi^*}(\alpha^2(p_t))$, then, at that price, the wealth of agent 1 will decrease and the wealth of agent 2 will increase, in expectation. As a consequence, if the previous inequality is satisfied by any feasible price p_t , it is reasonable to expect that $\lim_{t \rightarrow \infty} w_t = 0$ and agents 1 vanish. Conversely, if the opposite inequality is persistently verified, it is agent 2 who is expected to vanish. With the language of stochastic dynamical systems, Bottazzi and Dindo (2014) prove that if $I_{\pi^*}(\alpha^1(\pi^1)) < I_{\pi^*}(\alpha^2(\pi^1))$ then the MSE in which agent 1 dominates ($w_1 = 1$) is asymptotically stable, while if $I_{\pi^*}(\alpha^1(\pi^1)) > I_{\pi^*}(\alpha^2(\pi^1))$ it is unstable. It works analogously for the MSE in which agent 2 dominates ($w_1 = 0$): just switch 1 and 2 in the previous inequalities. Before providing a formal proof of how the MSE stability conditions translate into information about the global behaviour of the system, it is useful to graphically inspect some cases.

Consider the situation in the left panel of Fig. 1. The inclined lines represent to the two strategies α^1 and α^2 as function of price. Their intersections with the diagonal (dotted line) represents the two MSE, $p = \pi^1$ and $p = \pi^2$, while the horizontal line is the true probability, $\pi^* = 0.5$. In this case the Euclidean distance is proportional to the “information” distance measured using the relative entropy. We can conclude thus that the MSE $p = \pi^1$ is stable while the MSE $p = \pi^2$ is unstable. More generally, for all realized prices in $[\pi^1, \pi^2]$, strategy α^1 is always nearer to the truth than strategy α^2 . As a consequence, along any trajectory, the expected growth rate of the wealth of agent 1 is positive. On the right panel the situation is the opposite. In this case the beliefs of both agents are below the truth, and is the wealth of agent 2 that constantly increases on average. In fact the MSE in which agent 2 dominates is the only stable MSE. These situations are in fact rather generic: when both π^1 and π^2 are greater or lower than π^* , the agent with the belief nearer to the truth always dominates, irrespective of the value of




Figure 2: Local stability. In both the plots we set $\pi^* = 0.5$, $\pi^1 = 0.3$, $\pi^2 = 0.8$, $c^1 = 0.9$. In the first $c^2 = 0.9$, in the second $c^2 = 0.4$.

the mixing coefficient c^1 and c^2 .

The case when $\pi^1 < \pi^* < \pi^2$ is more interesting. Consider the situation in Fig. 2 where the belief of agent 1 is the closest to the truth. Notice that the MSE $p = \pi^2$ is always unstable: when the price is near to π^2 , the strategy of the wealthier agent, agent 2, is always further from the truth than the strategy of agent 1. Conversely, the stability of the MSE $p = \pi^1$ depends on the value of c^2 : for large values of c^2 (right panel) it is asymptotically stable, while for low values of c^2 (left panel) it is unstable.

We can observe that there exists a value of c^2 , call it \bar{c}^2 , such that the distance between π^1 and π^* and the distance between $\alpha^2(\pi^1)$ and π^* is the same. Then, if $c^2 < \bar{c}^2$ the MSE $p = \pi^1$ is unstable, while if $c^2 > \bar{c}^2$ it is asymptotically stable. In the case in which the belief of agent 2 is the closest to the truth, the reverse holds (see the right panel of Fig 2).

The analysis so far concerned the local stability of the two MSE of the system. While in a generic stochastic dynamical system local results do not translate into global predictions, in the present model, using the results in Bottazzi and Dindo (2015), one can prove the following

Proposition 3.1. *Given the system defined in (5) and (6), and with the definition in (10), one has*

- i) if the MSE $p = \pi^1$ is stable ($\mu(\pi^1) > 0$) then the MSE $p = \pi^2$ is unstable ($\mu(\pi^2) > 0$), agent 1 dominates, agent 2 vanishes and $p_t \rightarrow \pi^1$ almost surely,*
- ii) if the MSE $p = \pi^2$ is stable ($\mu(\pi^2) < 0$) then the MSE $p = \pi^1$ is unstable ($\mu(\pi^1) < 0$), agent 2 dominates, agent 1 vanishes and $p_t \rightarrow \pi^2$ almost surely,*
- iii) if both MSE are unstable ($\mu(\pi^1) \leq 0$ and $\mu(\pi^2) \geq 0$) then both agents survive, $\limsup_{t \rightarrow \infty} p_t = \pi^2$ and $\liminf_{t \rightarrow \infty} p_t = \pi^1$ almost surely.*

Proof. See Appendix B. □

The previous theorem derives global predictions about the asymptotic dynamics of the system from the local stability analysis of MSE. The theorem rules out the possibility that both MSE are asymptotically stable⁴

It is easy to see that the last condition requires $\pi^1 < \pi^* < \pi^2$. That is, a persistent crowd requires not only a certain degree of divergence of opinions, but also the contemporaneous presence of a “pessimist” and an “optimist” belief, that is beliefs respectively lower and higher than the true success probability. Moreover, agents can both survive only if the wealth share of each agent grows, in expectation, when the other agent is wealthier. This entails the seemingly counter-intuitive condition that when the prevailing market price is near to one agent’s belief, that agent, on average, loses wealth in favor of the opponent. This is made possible by the presence of the mixing coefficient. When agent 1 dominates the market, agent 2 invests according to a strategy that mixes his belief with the belief of the other agent, producing a result that might have a lower relative entropy with respect to the success probability than the belief of agent 1. It is immediate to see that this can be the case only when the “mixing” coefficient of the agent with the worst belief is small enough (c.f. Fig 2).⁵ The combinations of beliefs that allow for the survival of both agents covers a substantial part of the parameter space, see Fig 3, and is greater the lower the value of c . One can interpret this conclusion as saying that, not knowing the true probability, the agent with the lowest c has *ex-ante* the highest chances to survive. Under the CRRA interpretation, this translate into the fact that the chances of survival are higher for the more risk averse agent.⁶ Hence, consistently with what done in Kets et al. (2014), and to avoid favoring, so to speak, one agent over the other, in what follows we will assume $c_1 = c_2 = c$.

3.2 Reconsidering the Wisdom

In the analysis of one shot prediction markets the contributions mentioned in Section 1 ascribe two possible degrees of wisdom to the crowd. The prevailing

⁴This is not a general properties of this kind of market models, however. This result depends on the assumption that agents adopt fractional Kelly rules. What might happen with different CRRA strategies, for instance, is discussed in Bottazzi and Dindo (2013)

⁵Conversely, the agent’s mixing parameter does not affect the ability of a trader to dominate the market. However, as MacLean et al. (1992) argue, there is a trade-off between risk and expected growth. Nonetheless here we are interested in asymptotic outcomes, thus, assuming that we are in a situation in which an agent dominates, the fact that her c is small only implies that her wealth will converge to 1 slower than a case in which her c is large.

⁶This is generally true because we are considering risk aversion coefficients not lower than one. For a more general analysis of the relation between risk aversion and survival in short lived security markets see Bottazzi and Dindo (2013)




Figure 3: Combinations of beliefs that let both agents survive. Shaded areas of increasing dark color correspond to increasing values of $c_1 = c_2 = c$: the lightest is 0.01, then 0.1, 0.25 and 0.5 which is the darkest. Every shaded area contains those darker than itself.

price could replicate (almost) exactly the true value of the asset, a sort of “strong” wisdom, or it could be just more precise than (almost) any individual evaluation, a sort of “weak” wisdom. When we extend the analysis to repeated markets, the question arises if these properties are to be considered (almost) always satisfied or, conversely, satisfied only on average. In the present setting, we formally advance the following

Definition 3.1. Let $E_t[p]$ stands for the unconditional price average computed after the initial transient of t periods

$$E_t[p] = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{\tau=t}^T p_\tau,$$

then when $t \rightarrow \infty$

1. the crowd is *almost surely strongly wise* if $\text{Prob}\{p_t = \pi^*\} \rightarrow 1$;
2. the crowd is *almost surely weakly wise* if $\text{Prob}\{|p_t - \pi^*| < |\pi^i - \pi^*|\} \rightarrow 1$ for any agent i ;
3. the crowd is *on average strongly wise* if $E_t[p] \rightarrow \pi^*$;
4. the crowd is *on average weakly wise* if $|E_t[p] - \pi^*| < |\pi^i - \pi^*|$ for any agent i .

By expressing the above properties asymptotically, we allow for a possible initial transient phase and focus on the long run dynamics of the system. Since we




Figure 4: $E[p] - \pi^*$ for $\pi^1 = 0.3$, $\pi^2 = 0.8$, $c = 0.01$ and $\pi^* \in \{0.35, 0.4, \dots, 0.75\}$. Confidence intervals are set to three standard errors.

are interested in investigating the properties of the crowd, the analysis is obviously limited to the region of the parameters space in which both agents survive, that is in which the conditions of point *iii* in Proposition 3.1 are satisfied. The implication of that Proposition is, however, that in this case the price fluctuates between π^1 and π^2 . Consequently, condition 1 in Definition 3.1 is never satisfied while condition 2 occurs only when agents' beliefs are at the same distance from the truth. One has the following

Corollary 3.1. *In the economy with two fractional Kelly traders as defined in (5) and (6), when both traders survive, the crowd is never almost surely strongly wise while it is almost surely weakly wise if and only if $\pi^* - \pi^1 = \pi^2 - \pi^*$.*

Thus, the asymptotic probability that the prevailing price is equal to the success probability is zero, irrespective of agent's belief and the mixing parameter they adopt.⁷

Consider now two beliefs π^1 and π^2 whose relative entropy with respect to the success probability is opposite to their Euclidean distance, that is such that $I_{\pi^*}(\pi^1) > I_{\pi^*}(\pi^2)$ but $\pi^* - \pi^1 < \pi^2 - \pi^*$. Due to the convexity of the relative entropy function this is possible and generic. Let \tilde{c} be the value of the mixing parameter for which $\mu(\pi^2) = 0$. Then when $c > \tilde{c}$, agents 2 dominates and it

⁷Following the same reasoning, it is easy to show that this also applies to the more general case when $c^1 \neq c^2$.

is $\lim_{t \rightarrow \infty} E_t[p] = \pi^2$. Conversely if $c < \tilde{c}$ both agents survive. However, for continuity argument, if c is sufficiently near to \tilde{c} , the expected price will be near to π^2 , thus not only different from π^* , but also farther away from π^* than π^1 .⁸ We can then conclude that

Corollary 3.2. *In the economy with two fractionally Kelly traders as defined in (5) and (6), generically there exist individual beliefs and mixing coefficients for which the crowd is neither on average strongly wise nor on average weakly wise.*

Since this seems to contradict the finding in Kets et al. (2014), we repeat their numerical experiment to analyze the source of disagreement. We set $\pi^1 = 0.3$, $\pi^2 = 0.8$, $c = 0.01$ and for each value of $\pi^* \in \{0.35, 0.4, \dots, 0.75\}$ we set the initial price equal to π^* . We iterate the stochastic map generated by (5) and (6) for a sufficiently high number of steps. The number of steps depends on the parameter values and it is chosen high enough for the price history to distribute according to the invariant distribution of the process. The procedure is independently repeated for $N = 1000000$ times and from the N independent runs the average and standard deviation are computed. Figure 4 reports the difference of the average over final prices and the true value π^* , together with $\sim 99\%$ confidence intervals obtained adding and subtracting from the difference three times the standard deviation over \sqrt{N} . The hypothesis that the crowd is on average strongly wise is rejected for a large set of π^* values. The difference with respect to the results presented in Kets et al. (2014) lies in the computation of confidence intervals. In Kets et al. (2014) the authors compute 1000 time averages over the last 10000 observations of runs consisting of 100000 periods. They estimate the expected price as the average across the 1000 time averages, while they use the 5-th and the 95-th percentiles of the time averages' empirical distribution as confidence interval. The price dynamics is however strongly autocorrelated and this can cause biases in the estimation of the time averages which result in too wide confidence intervals.

As a robustness check we also compute the expected price using the Fokker-Planck approximation of the invariant distribution of agents relative wealth. Specifically, we consider the dynamics of the log wealth difference $z_t = \log w_t / (1 - w_t)$. We embed the discrete dynamics for z_t in continuous time using a homogeneous Poisson process. Then we truncate the Kramer-Moyal expansion of the resulting continuous time Chapman-Kolmogorov equation at the second order, to obtain the diffusive approximation of the invariant probability distribution F_z .⁹ From this distribution, one can simply obtain an estimate of the asymptotic expected

⁸As a specific example, take $\pi^* = 0.2$, $\pi^1 = 0.1$, $\pi^2 = 0.32$ and $c = 0.96$. Numerical simulations show that in this case it is $E[p] = 0.3191 \pm 0.00003$.

⁹The derivation is sketched in Appendix D. For more details see Bottazzi and Giachini (2016)




Figure 5: $E[p] - \pi^*$ for $\pi^1 = 0.3$, $\pi^2 = 0.8$, $c = 0.01$ and $\pi^* \in \{0.35, 0.4, \dots, 0.75\}$. Confidence intervals are set to three standard errors.

price using (6)

$$E[p] = \int_{-\infty}^{+\infty} dF_z \frac{\pi^1 e^z + \pi^2}{e^z + 1}. \quad (11)$$

In Fig. 5 we report the difference between the average price computed with (11) and π^* , together with the result of the previous numerical exercise. As can be seen, the result obtained through the Fokker-Planck approximation is never significantly different from the numerical simulations and provide more robust evidence of the discrepancy between the average price and the true probability. Moreover, since F_z approximates the invariant distribution, we can conclude that the average price is not only a biased estimator of π^* , but, in general, it is also inconsistent. That is, the lack of wisdom is not a matter of how long one averages the price: even letting the number of observation become arbitrary large, the bias does not disappear. Notice that the bias goes in the direction of the agent that is, in terms of relative entropy, closest to the truth, suggesting that the market rewards the agent that has the best belief with an higher average wealth.

If instead of looking at the distance between prices and true probabilities one looks at their relative entropy, the efficiency of the market mechanism seems improved, as one can easily prove the following

Proposition 3.2. *In the economy with two fractionally Kelly traders as defined in (5) and (6), if both agents survive, then the relative entropy of the average price*




Figure 6: $(E[p] - \pi^*)/\sigma(p)$ for $\pi^1 = 0.3$, $\pi^2 = 0.8$, $\pi^* \in (0.35, 0.75)$ and different values of c .

with respect to the true probability is always lower than the relative entropy of their beliefs.

Proof. See Appendix C. □

This is not surprising. As already suggested elsewhere (Blume and Easley, 2009; Evstigneev et al., 2009; Bottazzi and Dindo, 2014), the information processed by the market and revealed by prices is more conveniently compared using the relative entropy than a metric distance. If one adopts the former kind of measure instead of the latter, then the market is always on average weakly efficient.

After having established the existence of a persistent difference between the true probability π^* and the prevailing prices p_t , we try to quantify it. A natural way to do it is to compare the bias with the magnitude of price fluctuations, by considering the ratio $(E[p] - \pi^*)/\sigma(p)$. If this ratio is small, the estimate obtained by taking the price average is better than just picking one realized price at random.

Figure 6 shows the values of $(E[p] - \pi^*)/\sigma(p)$ for $\pi^* \in (0.35, 0.75)$ and different values of c computed using the approximate invariant distribution in (14). Notice that for all the c considered, it is $E[p] = \pi^*$ for $\pi^* \simeq 0.5656$ which is close, but not equal, to the point, $\pi^* \simeq 0.5609$, in which the relative entropy of agents' beliefs with respect to the truth is equal. With lower values of c , one gets lower values of $(E[p] - \pi^*)/\sigma(p)$. About this point, Kets et al. (2014) conjecture that




Figure 7: $(E[p] - \pi^*)/\sigma(p)$ for $\pi^1 = 0.3$, $\pi^2 = 0.8$, $\pi^* \in \{0.35, 0.4, \dots, 0.75\}$ and $c = 0.25, 0.5$.

$\lim_{c \rightarrow 0} E[p] = \pi^*$. Using the diffusive approximation of F_z we prove that, as $c \rightarrow 0$, one has $E[p] \rightarrow \pi^*$ while the variance of p goes to zero, confirming the conjecture (See Appendix D for the derivation). However, notice that a smaller value of c implies a higher risk aversion. Thus, in principle, we should allow agents to become infinitely risk averse in order to have the convergence of the expected price to the success probability. Referring to the price-dependent coefficient in (4), $c = 0.05$ corresponds to a γ between 15 and 25, $c = 0.01$ corresponds to a γ between 75 and 120 and $c = 0.005$ corresponds to a γ between 150 and 220. According to the classical experiment of Holt and Laury (2002), just a small fraction of their sample (between 1% and 6%, depending on the treatment) show a relative risk aversion larger than 1.37, hence such large values of risk aversion seem quite unrealistic. Indeed, MacLean et al. (1992) take as realistic examples of fractional Kelly betting the half and the quarter Kelly betting (respectively $c = 0.5$ and $c = 0.25$). If we adopt his proposal, we obtain the situation shown in Fig. 7. The bias shows the same patterns discussed in advance, but here the magnitude of the error is much larger.

Summarizing, a trade-off emerges between the size of the bias (compared to standard deviation and in absolute terms) and the implied risk aversion of the agents. In fact, realistic risk attitudes imply large bias and in order to keep the bias low, one should allow for unrealistically high levels of risk aversion.

4 More than two agents

The analysis so far concerned the seemingly special case of just two traders. Although this small number might appear peculiar, we have seen that also in this simple case, we are able to discuss, measure and explain the presence of wisdom, or madness, in prevailing market prices. The aim of this section is to show that when more agents are considered, the qualitative dynamics of the market and its efficiency are not expected to change much. Assume $N > 2$ agents are trading in the market, with individual beliefs $\pi^1 < \dots < \pi^k < \pi^* < \pi^{k+1} < \dots < \pi^N$. Thus agent k and $k + 1$ have the most correct beliefs, while agent 1 and N are characterized by somehow extreme beliefs, and are farthest from the truth. Consider $j < k$ and notice that

$$I_{\pi^*}(\alpha^k(\pi^j)) = I_{\pi^*}(c\pi^k + (1-c)\pi^j) \leq cI_{\pi^*}(\pi^k) + (1-c)I_{\pi^*}(\pi^j) < I_{\pi^*}(\pi^j),$$

which means that when an agent in $\{1, \dots, k-1\}$ owns almost all wealth, agent k grows more than him. As a consequence, none of these agents will ever dominate. An identical argument shows that also the agents in $\{k+2, \dots, N\}$ never dominate.

But will the agents with beliefs farther from the truth survive in the long run? The fate of these agents depends on the value of the mixing parameter c . When this value is sufficiently large, for $j < k$ it is

$$I_{\pi^*}(\alpha^k(p)) \sim I_{\pi^*}(\pi^k) < I_{\pi^*}(\pi^j) \sim I_{\pi^*}(\alpha^j(p)), \forall p.$$

In this case, the drift of the relative wealth of agents k with respect to agent j is always positive so that $\lim_{t \rightarrow \infty} \log(w_t^k/w_t^j) = +\infty$ almost surely and agent j vanishes. Following the same argument, it is $\lim_{t \rightarrow \infty} \log(w_t^{k+1}/w_t^j) = +\infty$ for $j > k+1$. Thus, all agents different from k and $k+1$ vanish when c is sufficiently large.

Conversely, using the Taylor expansion of the logarithm, one simply has

$$\lim_{c \rightarrow 0} I_{\pi^*}(\alpha^j(p)) - I_{\pi^*}(\alpha^i(p)) \sim c(\pi^i - \pi^j) \frac{\pi^* - p}{p(1-p)}.$$

Thus, when $c \rightarrow 0$, the price keeps fluctuating in the interval (π^1, π^N) . When $p > \pi^*$, the wealth of all agents with beliefs below π^* increases in expectation with respect to the wealth of any agent with beliefs above it. This brings the price below π^* and now the reverse happens. This dynamics is qualitatively similar to what would happen if only agents 1 and N were present in the market.

Summarizing, when c is sufficiently close to zero, all agents in the market survive. As c grows, the agents with the less accurate beliefs gradually disappear. When the mixing parameter is sufficiently large, only the two agents with the most

accurate beliefs, one above and one below the truth, survive. In all these cases, the dynamics of the market is similar to a pairwise comparison, in which the beliefs to be compared belong to the surviving agents with extremal beliefs.

5 Exploiting the Wisdom

Even if the wisdom seems to be present only in modest amounts in the crowd we are studying, and a sizeable difference does generally emerge between prevailing market prices and underlying probabilities, still we have seen that the former convey some information about the latter. The question we start to address in this section is whether this information can be actually exploited by a trader. For instance, one might advance the argument that recording past prices and using them to guess π^* , might improve the betting behavior of an agent. In order to investigate this issue we consider the simplest possible situation in which a bettor employs as belief the last realized market price¹⁰.

We start with the “exogenous” case in which the imitating fractional Kelly bettor exists outside the market and bets on the outcomes of the process without influencing the price. The evolution of her wealth w'_t reads

$$w'_t = \begin{cases} \frac{cp_{t-1} + (1-c)p_t}{p_t} w'_{t-1} & \text{if } s_t = 1, \\ \frac{1 - cp_{t-1} - (1-c)p_t}{1 - p_t} w'_{t-1} & \text{if } s_t = 0. \end{cases} \quad (12)$$

Once the state of the word at time t is realized, the relation between p_t and p_{t-1} is known, specifically

$$p_t^1 = p_t(s_t = 1) = p_{t-1} + \frac{c(p_{t-1} - \pi^1)(\pi^2 - p_{t-1})}{p_{t-1}}$$

$$p_t^0 = p_t(s_t = 0) = p_{t-1} - \frac{c(p_{t-1} - \pi^1)(\pi^2 - p_{t-1})}{1 - p_{t-1}}.$$

Thus we can compute the expected growth rate of this external agent conditional

¹⁰Notice that, in the present model, if agents can upgrade their beliefs, it is smarter to learn the success probability recording the sequence of states of the world than looking at prices. However, here the point is not to analyze learning issues, but to investigate what happens to the agent who adopt some form of imitative behavior based exclusively on price information. This is reminiscent of what happens in real markets, in which agents lack a reliable description of the underlying process, the world economy, and try to build efficient portfolios based on the past realized prices

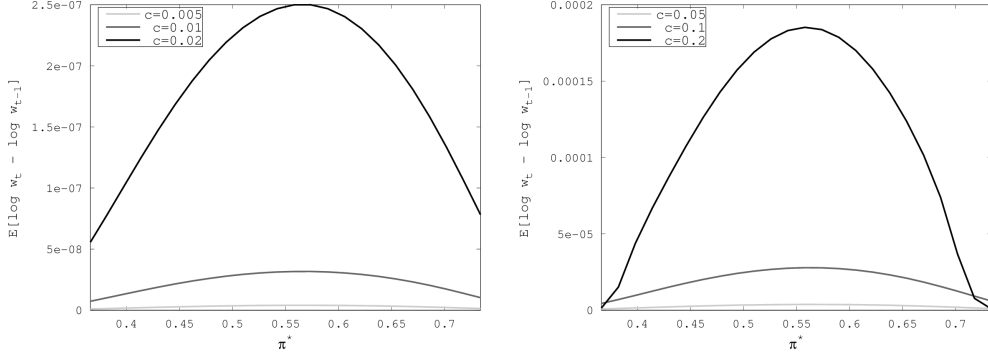


Figure 8: $E[\log w'_t - \log w'_{t-1}]$ for $\pi^1 = 0.3$, $\pi^2 = 0.8$, and different values of π^* and c .

on p_{t-1} , namely

$$\begin{aligned}
E \left[\log w'_t - \log w'_{t-1} \middle| p_{t-1} \right] &= \\
&= (\pi^*)^2 \log \left(\frac{cp_{t-1} + (1-c)p_t^1}{p_t^1} \right) + \pi^*(1-\pi^*) \log \left(\frac{1 - cp_{t-1} - (1-c)p_t^1}{1 - p_t^1} \right) \\
&+ \pi^*(1-\pi^*) \log \left(\frac{cp_{t-1} + (1-c)p_t^0}{p_t^0} \right) + (1-\pi^*)^2 \log \left(\frac{1 - cp_{t-1} - (1-c)p_t^0}{1 - p_t^0} \right).
\end{aligned}$$

and finally we can use the approximate expression of the price distribution derived from (14) to compute the unconditional growth rate.

Figure 8 reports the result of this computation as a function of π^* for one specific set of parameters. The log-growth rates are positive for all the values of π^* and c considered. Our numerical investigation suggests that this is always the case.¹¹ Hence, using as belief the last realized price is a winning strategy for an agent who does not influence the price. This consideration is also informative of the performance of an imitating agent that participates in the market but owns a negligible share of the total wealth. In that case, the wealth of the imitator will grow in expectation. In other terms, adopting the last observed price as belief, seems to be a viable survival strategy, at least.

What happens, instead, if we assume that the imitating agent does actually trade in the market? To answer that question, consider the model of Section 3 and assume that agent 1 starts at $t = 1$ with a fixed belief π^1 , but in each subsequent time step $t > 1$, adopts as belief the last observed price. Remember that at each time step the market price is a convex combination of agents beliefs, thus $p_1 \in (\pi^1, \pi^2)$. Since the next time the belief of agent 1 becomes p_1 , it is $p_2 \in (p_1, \pi^2)$.

¹¹Details available upon request.

In other terms, $p_t \in (p_{1-1}, \pi^2)$ and the sequence of prices p_t is strictly increasing on any realization of the underlying process. Since there are no other fixed points apart the MSE discussed in Section 3, this implies that $\lim_{t \rightarrow \infty} p_t = \pi^2$, that is, the price converges to π^2 in the long-run¹². Notice how, in terms of price’s informative power, the situation is now worse than in the case of fixed beliefs, in particular the distance between $E[p]$ and π^* becomes larger. Even assuming that agent 1 initially knew the truth, $\pi^1 = \pi^*$, with the updating procedure the price converges to the belief of agent 2. We pass from a situation which is (potentially) the most beneficial for market efficiency, agent 1 dominates and sets $p_t = \pi^*$ in the long-run, to a situation in which the asymptotic price that emerges is far away from the truth.

This simple example shows that there exist cases in which having agents that try to take advantage of the information embedded in the price can be noxious for the efficiency of the market. In fact, adopting as belief the last realized price, the agent does not convey any piece of information that is new with respect to what already is in the market. On the contrary, in a sort of *social influence*, by setting as belief the past price, she gradually incorporates the piece of information owned by the other and discards her own. As time goes by, the updating agent ends up having the same piece of information of the opponent (that is, she becomes completely “*influenced*”) and the price reveals it. In fact, this argument is not new. The detrimental effect of social influence for the *Wisdom of Crowds* was already mentioned by the same Surowiecki (2004) and showed in the experiments of Lorenz et al. (2011).

6 Conclusion

We presented a very simple model in which two fractional Kelly traders repeatedly bet on a uncertain binary outcome, occurring with fixed and unknown probability. If agent beliefs about the probability of success are one bigger and one lower than the truth, and if their risk aversion is sufficiently high, then they both survive and none ends up winning all the money. This is the effect of the fractional Kelly strategies, which mixes *a priori* beliefs with contemporary prices, and leads to a situation in which the wealthier agents is, persistently, the one who gains less (or loses more). Some previous contributions in the literature suggest that in this case the “crowd” of the two survivors becomes “wise”, in the sense that the price history emerging from their interaction carries more information about the true probability than the individual beliefs of the agents. We show that this is not always true. Specifically, we show that, in general, the price does not converge

¹²This example can be easily generalized to a case in which agent 1 adopt a fixed belief for T periods and then starts updating.

almost surely, nor in expectation, to the true unknown probability. Moreover, there are situations in which the average prevailing price, while having a lower relative entropy with respect to the truth than the beliefs of agents, is farther from it. This consideration, while stressing the necessity to adopt the right quantity to compare strategies and outcomes (along the lines discussed in Blume and Easley (2009); Evstigneev et al. (2009); Bottazzi and Dindo (2014)), gets rid of the naive presumption that market prices, emerging from the interaction of heterogeneous traders, are intrinsically “better” than their individual opinions.

Notice that our agents are not able to learn the probability of success. Absurd as it seems, this is done in search of a higher realism. Indeed, if we endow the agents with the ability to learn the truth, not only the realized prices will converge there, but the crowd essentially disappears as agents becomes asymptotically identical. Since the progressive convergence of opinions is clearly not something that we observe in real markets, we can deduce that traders there are not able to learn the truth, whatever it might be. We decided to shape our agents on their example.

In addition, we simply illustrate the detrimental effects that social influence might have on the wisdom of the crowds, as argued by some scholars: if an agent in our model tries to exploit market prices to build more accurate predictions and invest accordingly, the informative content of prices is further spoiled and prices may end up revealing the worst, that is less accurate, opinion, among those initially available.

References

- Arrow, K. J., R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J. O. Ledyard, S. Levmore, R. Litan, P. Milgrom, F. D. Nelson, et al. (2008). The promise of prediction markets. *SCIENCE-NEW YORK THEN WASHINGTON*- 320(5878), 877.
- Beygelzimer, A., J. Langford, and D. M. Pennock (2012). Learning performance of prediction markets with kelly bettors. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*, pp. 1317–1318. International Foundation for Autonomous Agents and Multiagent Systems.
- Blume, L. and D. Easley (1992). Evolution and market behavior. *Journal of Economic Theory* 58, 9–40.
- Blume, L. and D. Easley (2009). Market selection and asset pricing. In T. Hens and K. Schenk-Hoppé (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. North-Holland (Handbooks in Economics Series).
- Bottazzi, G. and P. Dindo (2013). Selection in asset markets: the good, the bad, and the unknown. *Journal of Evolutionary Economics* 23(3), 641–661.
- Bottazzi, G. and P. Dindo (2014). Evolution and market behavior with endogenous investment rules. *Journal of Economic Dynamics and Control* 48, 121–146.
- Bottazzi, G. and P. Dindo (2015). Drift criteria for persistence of discrete stochastic processes on the line. Technical report, Laboratory of Economics and Management (LEM), Sant’Anna School of Advanced Studies, Pisa, Italy.
- Bottazzi, G. and D. Giachini (2016). Wealth and price distribution by diffusive approximation in a repeated prediction market. LEM Papers Series 2016/13, Laboratory of Economics and Management (LEM), Sant’Anna School of Advanced Studies, Pisa, Italy.
- Evstigneev, I., T. Hens, and K. Schenk-Hoppé (2009). Evolutionary finance. In T. Hens and K. Schenk-Hoppé (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. North-Holland (Handbooks in Economics Series).
- Gjerstad, S. (2005). Risk aversion, beliefs, and prediction market equilibrium. Technical report, Economic Science Laboratory, University of Arizona.
- Hayek, F. A. (1945). The use of knowledge in society. *The American economic review* 35(4), 519–530.

- He, X.-Z. and N. Treich (2012). Heterogeneous beliefs and prediction market accuracy. Technical report, Citeseer.
- Holt, C. and S. Laury (2002). Risk aversion and incentive effects. *American Economic Review* 92(5), 1644–1655.
- Kelly, J. (1956). A new interpretation of information rates. *Bell System Technical Journal* 35, 917–926.
- Kets, W., D. M. Pennock, R. Sethi, and N. Shah (2014). Betting strategies, market selection, and the wisdom of crowds. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- Lorenz, J., H. Rauhut, F. Schweitzer, and D. Helbing (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences* 108(22), 9020–9025.
- MacLean, L., E. Thorp, and W. Ziemba (2010). Good and bad properties of the kelly criterion. *Risk* 20(2), 1–11.
- MacLean, L., W. T. Ziemba, and G. Blazenko (1992). Growth versus security in dynamic investment analysis. *Management Science* 38(11), 1562–1585.
- MacLean, L. C., R. Sanegre, Y. Zhao, and W. T. Ziemba (2004). Capital growth with security. *Journal of Economic Dynamics and Control* 28(5), 937–954.
- MacLean, L. C. and W. T. Ziemba (1999). Growth versus security tradeoffs in dynamic investment analysis. *Annals of Operations Research* 85, 193–225.
- MacLean, L. C., W. T. Ziemba, and Y. Li (2005). Time to wealth goals in capital accumulation. *Quantitative Finance* 5(4), 343–355.
- Manski, C. F. (2006). Interpreting the predictions of prediction markets. *economics letters* 91(3), 425–429.
- Surowiecki, J. (2004). *The wisdom of crowds*. Anchor.
- Thorp, E. O. (2006). The kelly criterion in blackjack, sports betting, and the stock market. *Handbook of asset and liability management* 1, 385.
- Wolfers, J. and E. Zitzewitz (2006). Interpreting prediction market prices as probabilities. Technical report, National Bureau of Economic Research.
- Ziemba, W. T. (2003). The stochastic programming approach to asset, liability, and wealth management.

A Fractional Kelly and expected utility maximization

The fractional Kelly investing can be mapped in the optimal rule of an agents that maximize next period utility under CRRA preferences with price dependent risk aversion coefficient $\gamma_{i,t}$ (see also MacLean et al. (2010)). Consider the utility function

$$u_{i,t}(w_t^i) = \begin{cases} \frac{(w_t^i)^{1-\gamma_{i,t}} - 1}{1 - \gamma_{i,t}} & \text{if } \gamma_{i,t} \in (0, +\infty), \gamma_{i,t} \neq 1 \\ \log(w_t^i) & \text{if } \gamma_{i,t} = 1 \end{cases}$$

Solving the expected utility maximization problem under the agent specific belief π^i , one gets

$$\alpha^i(p_t) = \frac{(1-p_t)^{\frac{1-\gamma_{i,t}}{\gamma_{i,t}}} (\pi^i)^{\frac{1}{\gamma_{i,t}}}}{(1-p_t)^{\frac{1-\gamma_{i,t}}{\gamma_{i,t}}} (\pi^i)^{\frac{1}{\gamma_{i,t}}} + p_t^{\frac{1-\gamma_{i,t}}{\gamma_{i,t}}} (1-\pi^i)^{\frac{1}{\gamma_{i,t}}}}. \quad (13)$$

Setting (13) equal to (3) and solving for $\gamma_{i,t}$, one obtains (4).

One can also show that fractional Kelly investing emerges as investing behavior of a fully rational agent who intertemporally maximizes the geometrically discounted expected logarithmic utility of wealth under price dependent beliefs $\pi_t^i = c^i \pi^i + (1-c^i)p_t$.

B Proof of Proposition 3.1

Let us define

$$\Delta_{\pi^*}(\pi^1 || \pi^2) = \pi^* \log \frac{\pi^1}{\pi^2} + (1-\pi^*) \log \frac{1-\pi^1}{1-\pi^2}.$$

Using the strict concavity of the logarithmic function it is immediate to see that $\mu(\pi^1) < c^2 \Delta_{\pi^*}(\pi^1 || \pi^2)$ and $\mu(\pi^2) > c^1 \Delta_{\pi^*}(\pi^1 || \pi^2)$.

Then $\mu(\pi^1) > 0$ implies $\Delta_{\pi^*}(\pi^1 || \pi^2) > 0$ and, in turn, $\mu(\pi^2) > 0$. Analogously, $\mu(\pi^2) < 0$ implies $\Delta_{\pi^*}(\pi^1 || \pi^2) < 0$ and $\mu(\pi^1) < 0$.

Consider the dynamics of the ratio of individual wealth $z_t = \log w_t / (1-w_t)$. If $\pi^1 < p_t < \pi^2$, then $z_t - z_{t-1} < 0$ if $s_t = 1$ and $z_t - z_{t-1} > 0$ if $s_t = 0$. This excludes the possibility of a finite deterministic fixed point or of deterministic drift for z_t . Moreover, since

$$\log \frac{\pi^1}{\pi^2} \leq z_t - z_{t-1} \leq \log \frac{1-\pi^1}{1-\pi^2},$$

the process z_t has bounded increments. Finally notice that

$$\lim_{z \rightarrow +\infty} \mathbb{E}[z_t - z_{t-1} | z_{t-1} = z] = \mu(\pi^2)$$

and

$$\lim_{z \rightarrow -\infty} \mathbb{E}[z_t - z_{t-1} | z_{t-1} = z] = \mu(\pi^1) .$$

Then if $\mu(\pi^1) > 0$ and $\mu(\pi^2) > 0$, according to (Bottazzi and Dindo, 2015, Theorem 3.1), z_t diverges to $+\infty$ a.s. and the first statement of the proposition follows.

If $\mu(\pi^2) < 0$ and $\mu(\pi^1) < 0$, according to (Bottazzi and Dindo, 2015, Corollary 3.1), z_t diverges to $-\infty$ a.s. and the second statement of the proposition follows.

Finally if $\mu(\pi^2) > 0$ and $\mu(\pi^1) < 0$, according to (Bottazzi and Dindo, 2015, Theorem 2.2), the process of z_t is persistent. Since z_t increases with probability $1 - \pi^*$ and decreases with probability π^* , one has that, as long as $\pi^* \in (0, 1)$, it is $\limsup_{t \rightarrow \infty} z_t = +\infty$ and $\liminf_{t \rightarrow \infty} z_t = -\infty$ a.s., so that the third statement follows.

C Proof of Proposition 3.2

Without loss of generality consider the wealth of the first agent. If both agents survive it must be $\mathbb{E}[\log w_t/w_{t-1}] = 0$, which implies

$$\begin{aligned} c \mathbb{E} [I_{\pi^*}(p) - I_{\pi^*}(\pi^1)] &= c \mathbb{E} \left[\pi^* \log \frac{\pi^1}{\pi^*} + (1 - \pi^*) \log \frac{1 - \pi^1}{1 - \pi^*} \right] \leq \\ &\mathbb{E} \left[\pi^* \log \left(c \frac{\pi^1}{p} + 1 - c \right) + (1 - \pi^*) \log \left(c \frac{1 - \pi^1}{1 - p} + 1 - c \right) \right] = 0 , \end{aligned}$$

so that

$$\mathbb{E} [I_{\pi^*}(p)] \leq I_{\pi^*}(\pi^1) .$$

Moreover, from the convexity of the relative entropy and the Jensen inequality, it is

$$I_{\pi^*}(\mathbb{E}[p]) \leq \mathbb{E} [I_{\pi^*}(p)] ,$$

which proves the assertion.

D Diffusive Approximation

The dynamics of $z_t = \log w_t/(1 - w_t)$ reads

$$z_{t+1} = z_t + \begin{cases} \log \frac{c\pi^1 + (1-c)p(z_t)}{c\pi^2 + (1-c)p(z_t)} & \text{if } s_t = 1, \\ \log \frac{1 - c\pi^1 - (1-c)p(z_t)}{1 - c\pi^2 - (1-c)p(z_t)} & \text{if } s_t = 0. \end{cases}$$

with $p(z) = (\pi^1 e^z + \pi^2)/(e^z + 1)$. By embedding the process in continuous time through a Poisson arrival with homogeneous intensity λ and truncating at the second order the Kramer-Moyal expansion of the resulting Champan-Kolmogorov equation we obtain the diffusive limit

$$\frac{\partial}{\partial t} f(z, t) = -\frac{\partial}{\partial z} \left(\lambda m(z) f(z, t) - \frac{\partial}{\partial z} (\lambda v(z) f(z, t)) \right)$$

where

$$m(z) = \pi^* \log \frac{c\pi^1 + (1-c)p(z)}{c\pi^2 + (1-c)p(z)} + (1 - \pi^*) \log \frac{1 - c\pi^1 - (1-c)p(z)}{1 - c\pi^2 - (1-c)p(z)},$$

$$v(z) = \frac{\pi^*}{2} \left(\log \frac{c\pi^1 + (1-c)p(z)}{c\pi^2 + (1-c)p(z)} \right)^2 + \frac{1 - \pi^*}{2} \left(\log \frac{1 - c\pi^1 - (1-c)p(z)}{1 - c\pi^2 - (1-c)p(z)} \right)^2.$$

are obtained by the first central cumulants of the increment, $m(z) = E[z_{t+1} - z_t | z_t = z]$ and $v(z) = E[(z_{t+1} - z_t)^2 | z_t = z]/2$. Since the process for z_t is unbounded, we can safely assume that the distribution is asymptotically vanishing so that we can solve the previous equation to obtain the density

$$f_z(x) = F'_z(x) = \frac{f_0}{v(x)} e^{\int_{z_0}^x dy \frac{m(y)}{v(y)}}, \quad (14)$$

where z_0 is the unique value for which $m(z_0) = 0$. This values does in general depend on all the parameters of the model. Bottazzi and Giachini (2016) show that (14) constitutes a good approximation of the actual distribution in all the points of the parameter space. Moreover, the approximation is better the lower the value of c . One can rewrite (14) as

$$f_z(x) = \frac{f_0}{c^2 V(x)} e^{-\frac{1}{c} \int_{z_0}^x dy \frac{M(y)}{V(y)}},$$

where $M(x) = -m(x)/c$ and $V(x) = v(x)/c^2$. Let

$$W(x) = \int_{z_0}^x dy \frac{M(y)}{V(y)}$$

and note that $W'(z_0) = 0$ and $W(z) \geq 0$ for any z . Then, when c is small, one can use the asymptotic expansion of Laplace-like integrals to obtain

$$\mathcal{M}_n(c) = \int_{-\infty}^{+\infty} dx f_z(x) x^n = f_0 \sqrt{\frac{2\pi}{c^3 W''(z_0)}} \frac{z_0^n}{V(z_0)} e^{-W(z_0)} + o(c^{-3/2}).$$

so that

$$\lim_{c \rightarrow 0} \mathbb{E}[z] = \lim_{c \rightarrow 0} \frac{\mathcal{M}_1(c)}{\mathcal{M}_0(c)} = \lim_{c \rightarrow 0} z_0$$

and

$$\lim_{c \rightarrow 0} \mathbb{V}[z] = \lim_{c \rightarrow 0} \frac{\mathcal{M}_2(c)}{\mathcal{M}_0(c)} - \left(\frac{\mathcal{M}_1(c)}{\mathcal{M}_0(c)} \right)^2 = 0.$$

In the limit $c \rightarrow 0$ the distribution of z becomes atomic and peaked around the point z_0 . Since $\lim_{c \rightarrow 0} p(z_0) = \pi^*$, the distribution of prices is, in the limit, atomic and peaked around π^* .