THE
ECONOMIC
JOURNAL

The Economic Journal, **123** (*May*), 491–513. Doi: 10.1111/j.1468-0297.2012.02561.x. Published by John Wiley & Sons, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

DO PREDICTION MARKETS PRODUCE WELL-CALIBRATED PROBABILITY FORECASTS?*

Lionel Page and Robert T. Clemen

This article presents new theoretical and empirical evidence on the forecasting ability of prediction markets. We develop a model that predicts that the time until expiration of a prediction market should negatively affect the accuracy of prices as a forecasting tool in the direction of a 'favourite/longshot bias'. That is, high-likelihood events are underpriced, and low-likelihood events are overpriced. We confirm this result using a large data set of prediction market transaction prices. Prediction markets are reasonably well calibrated when time to expiration is relatively short, but prices are significantly biased for events farther in the future. When time value of money is considered, the miscalibration can be exploited to earn excess returns only when the trader has a relatively low discount rate.

Prediction markets are a relatively new form of financial market (Oliven and Rietz, 2004), in which traders buy and sell contracts that are bets on specific future events. The price of a contract is fully determined by supply and demand in the market. Prediction market prices are viewed by some as a mechanism to retrieve and aggregate all private and relevant information for estimating the underlying event's probability, which in turn can be a useful input for a decision maker (Wolfers and Zitzewitz, 2006*a*).

For example, a contract on an election may consist of a bet on the following event: 'Candidate A will win the election.' Buying a contract is betting on the occurrence of the event, whereas selling the contract is betting against it. Traders buy and sell contracts at market price $p \in [0, 1]$. If the event occurs, the buyer receives one from the seller, and zero otherwise.[1]

Interest in prediction markets stems from the fact that the market price of a contract can be viewed as an estimate of the probability of the underlying event. For example, Gjerstad (2005) and Wolfers and Zitzewitz (2006*b*) show that if traders have conventional utility functions with limited risk aversion (i.e. the utility functions are roughly linear over the range of the values considered), the equilibrium price is a good approximation of the mean of the belief distribution of traders on the likelihood of the event. When traders are risk averse with logarithmic utility functions, a general result from the research in finance on asset pricing with heterogeneous beliefs is that the asset price should be the arithmetic mean of the distribution of belief of the traders (Rubinstein, 1976).

If the traders' beliefs about the event probability are well calibrated on average, then a market price may be considered an estimate of the event probability. By 'calibrated'

* Corresponding author: Lionel Page, School of Economics and Finance, Queensland University of Technology, GPO Box 2434, Brisbane, QLD 4001, Australia. Email: lionel.page@qut.edu.au.

[1] The 'currency' used in the transactions can be thought of as tokens. For example, Intrade's currency unit is a point; one point equals \$0.10. On betting exchanges like Betfair, the price of assets is presented in the form of odds. In spite of the difference in presentation, the underlying mechanisms of traditional prediction markets and betting exchanges are similar. The analysis we develop here for prediction markets applies similarly to betting exchanges.

we mean that on average, when the trader's belief is p, the expected frequency of the corresponding outcome occurring is p. The same definition will be used to define calibrated market prices; a market price is calibrated when the expected frequency of occurrence equals the price.

In this article, we look at the ability of prediction markets to estimate the probability of a future event outcome that will not be known for one month or more. Such longer term prediction markets are most relevant for decision makers who must make decisions well before the uncertainties in the underlying events are resolved. In large organisations, decisions and their implementation take time. To be most useful in an organisational setting, prediction markets should be able to provide calibrated probabilities well in advance of the future event.

This article is the first to take into account the time dimension when modelling the equilibrium price in a prediction market. Until now, prediction market prices have been considered (at least implicitly) to be as accurate in the short term as in the long term. In this article, we show that the time dimension can play an important role in the calibration of the market price. When traders who have time discounting preferences receive no interest on the funds committed to a prediction-market contract, a cost is induced, with the result that traders with beliefs near the market price abstain from participation in the market. This abstention is more pronounced for the favourite because the higher price of a favourite contract requires a larger money commitment from the trader and hence a larger cost due to the trader's preference for the present. Under general conditions on the distribution of beliefs on the market, this produces a bias of the price towards 50%, similar to the so-called favourite/longshot bias.[2]

We confirm this prediction using a data set of actual prediction markets prices from 1,787 market representing a total of more than 500,000 transactions.

The remainder of this article is organised as follows: first, in Section 1, we discuss the factors which may influence the calibration of prediction markets and the relevant theoretical and empirical research. In Section 2, we model the effect of time discounting preferences on market prices. We show that traders' decisions to invest in prediction markets are influenced by the market duration, resulting in prices that display a longshot bias. In Section 3, we describe our data set of transaction prices for prediction markets lasting more than one month. The size of our data set allows us in Section 4 to use non-parametric estimation methods to calculate local estimates of market price miscalibration. In particular, we are able to study the calibration of market prices near zero or one with more precision than previous researchers have been able to achieve. Section 5 presents an economic analysis of the data; we show that accounting for the time value of money affects the ability to exploit the systematic bias in prices. Section 6 discusses the results and concludes.

1. Calibration of Prediction Market Prices

The calibration of prediction markets' estimates has been one of the main questions investigated by previous prediction market research. Such research has identified

[2] For convenience, we will hereafter use the term 'longshot bias' to refer to the combination favourite/longshot bias.

several possible reasons why prices could be biased or miscalibrated, systematically departing from the event probability.

First, in the case of risk-neutral traders with limited budgets, a longshot bias, when the prices are biased towards 1/2 (see, for instance, Ottaviani and Sørensen, 2009, 2010*a*) can arise even if their beliefs are well calibrated. The reason for this is that the market clearing condition typically requires that the price be located not at the median of the distribution of traders' beliefs but at a percentile closer to 1/2 than the median (Ali, 1977; Manski, 2006). Extending this framework, Ottaviani and Sørensen (2010*b*) examine a fully revealing rational expectations equilibrium, in which traders make correct inferences from prices, given common knowledge of the information structure and prior beliefs. They find that, even in this situation, the price underreacts to new information, thereby leading to prices biased towards 1/2.

Second, if traders are not risk neutral but risk averse, as in Gjerstad (2005), Wolfers and Zitzewitz (2006*a*) and Ottaviani and Sørensen (2010*b*), traders' risk aversion can automatically create a longshot bias. However, the bias would be small if traders' risk aversion falls within the range typically observed in empirical studies.

Third, when prediction market prices play a role in the decision processes of individuals or organisations, there may be an incentive for some agents to manipulate the market by buying or selling contracts to move the price up or down. What role might manipulators play in shaping prediction market prices? Results from theoretical and empirical studies are mixed on this issue. Hanson and Oprea (2009) propose a model in which manipulation causes prediction market prices to be better calibrated due to the liquidity they provide. However, in the more general discussion of manipulation of prices of financial assets, the seminal work of Allen and Gale (1992) shows that manipulation can affect financial market prices even with rational traders. On the empirical side, some experimental studies find prediction markets to be resilient with respect to manipulation (Camerer, 1998; Hanson *et al.*, 2006; Rhode and Strumpf, 2006). However, Hansen *et al.* (2004) report successful attempts at manipulating prices in the Iowa Electronic Market.

Fourth, Ottaviani and Sørensen (2007) show how a prediction market can create an incentive to manipulate the outcome itself when participants have an influence on it (typically in firms). In such a situation, the price still represents the probability of the outcome but the measurement tool (i.e. the market) cannot be separated from the measured probability, as the creation of the prediction market may foster manipulation of the outcome, thereby affecting the probability of the event.

Fifth, evidence from behavioural finance suggests that traders are subject to biases when estimating subjective probabilities and making decisions about financial transactions. Thus, one might expect prices emerging from traders' average beliefs to be systematically biased as well. In particular, Kahneman and Tversky (1979) propose a model of decision making in which individuals weight probabilities when making decisions. Small probabilities are typically overweighted and large probabilities underweighted. Thaler and Ziemba (1988) suggest that this is the main reason for the presence of longshot bias in prediction markets.

Sixth, relatively recent psychological studies have shown that subjective probabilities may display *partition dependence*. Fox and Clemen (2005) define partition dependence as 'the tendency for judged probabilities to vary systematically with the way a state space is partitioned into events for which probabilities are assessed'. This phenomenon is due, according to Fox and Clemen (2005), to the fact that 'people anchor their judgements on equal probabilities for each event in the specified partition (the ignorance prior distribution) and adjust insufficiently to account for their beliefs about how the likelihood of the events differ'. Sobel and Raines (2003) propose an information model in which Bayesian bettors begin with an ignorance prior (equal probabilities) and update the prior according to private information. As a result, longshots are overestimated and favourites underestimated. Sonnemann *et al.* (2011) find an ignorance prior bias in experimental prediction markets and some suggestive evidence of it in real prediction markets.

Although these elements tend to suggest that market prices may not be good estimates of underlying probabilities, Oliven and Rietz (2004) argue and provide some evidence that prediction market prices are primarily determined by a minority of active traders, the so-called market makers. They find that these traders tend to behave more rationally than others. For this reason, an efficient price could emerge from a set of biased traders.

This article contributes to this literature by showing that the time dimension plays a critical role in the calibration of the prediction market prices, limiting the ability of prediction markets to provide calibrated probabilities for events beyond the short term. To do so, we generalise Manski (2006)'s model of a prediction market with risk-neutral traders having budget constraints to a two-period situation.

2. Model

Previous papers have implicitly assumed that a trader's investment decision is a one-period problem (Gjerstad, 2005; Manski, 2006; Wolfers and Zitzewitz, 2006*b*; Ottaviani and Sørensen, 2010*b*). However, this assumption is inadequate for studying prediction markets in which a significant amount of time can elapse between the purchase of a contract and the final resolution of the uncertainty. In these cases, the time value of money should be taken into account. We might also expect that a trader who derives some utility from gambling might prefer not to commit a significant amount of money over a long period (Diecidue *et al.*, 2004); doing so would further constrain the amount of money available for subsequent gambling activity during the corresponding period.

To model the trader's choice, we introduce a time dimension into the framework proposed by Manski (2006) which is the prediction market equivalent of the framework proposed by Ali (1977) for racetrack betting. In this framework, a risk-neutral trader with budget y has the opportunity to purchase some quantity of a single contract. The trader's problem is to choose a quantity x of contracts to maximise his or her profit, given the contract price p and his or her belief q that the contract will be successful. The trader chooses to invest xp from his or her budget y. At resolution, the contract yields a value zero or one. Given that the trader has a belief q that the contract will be successful, his or her expected gain is qx.[3] In previous models, time plays no role. It is as if

[3] In line with previous literature, we take belief as a primitive. Traders can update their prior beliefs with the market prices they can observe. We do not model explicitly how traders update their prior with market prices (Ottaviani and Sørensen, 2010*b*). We analyse the consistency between prices and (posterior) beliefs. We thank a reviewer for stressing this point.

the contract's value is revealed to the trader as soon as the contract is purchased. To capture the time dimension of prediction markets, suppose that the outcome of the contract takes place later in time. We introduce (the inverse of) a time discounting factor $\beta > 1$ in the utility function of the trader such that contracts that are resolved later in time have a lower utility, x/β.

Thus, the trader solves

$$\begin{aligned} \max_x V(x) &= y - xp + q\frac{x}{\beta} \\ \text{s.t. } xp &\leq y. \end{aligned} \tag{1}$$

For convenience, we model the seller as a trader who *buys the reverse contract.* Imagine a market maker who simultaneously sells the original contract to those who want to buy and sells the reverse contract to those who want to sell the original contract. For the market maker, the equilibrium price not only equates supply and demand but it also balances the books, putting the right amount of money on each side of the bet so that the market maker can transfer all of the money to the side that wins the bet. If the market maker takes a commission from each bettor, then the market maker is called a bookie. Thus, a seller solves the problem

$$\begin{aligned} \max_x V(x) &= y - x(1-p) + (1-q)\frac{x}{\beta} \\ \text{s.t. } x &\leq y. \end{aligned} \tag{2}$$

In this framework, there are only two relevant periods for the trader: the period where the decision to buy or sell the contract is made and the period where the final outcome of the contract is revealed. This is a simplistic representation of reality as in practice traders can continuously change their market positions. It is, however, the simplest extension of existing models that assume that traders only take into account their beliefs of the final probability of the event.

The introduction of the time dimension has two effects. First, it reduces the potential number of traders willing to participate in the market:

PROPOSITION 1. *If traders solve (2), any value of $\beta > 1$ results in a region $1-(1-p_1)\beta < q < p_1\beta$, such that any trader whose belief q falls within this region neither buys nor sells the contract.*

Proof. The Kuhn–Tucker conditions (KTC) for a trader who buys the contract imply that he or she invests all of his or her budget y to buy if $q > p\beta$ and invests nothing if $q < p\beta$.

The KTC for a seller imply that he or she will buy the reverse contract if $1-q > (1-p)\beta$, that is if $q < 1-(1-p)\beta$. When $\beta > 1, 1-(1-p)\beta < p\beta$. Thus, the no-trade region is $1-(1-p)\beta < q < p\beta$.

From Proposition 1, it is evident that, as β increases, the no-trade region expands. Figure 1 shows these boundaries for two different levels of discounting. For each price, traders whose beliefs fall within the boundaries will neither buy nor sell contracts on the market.

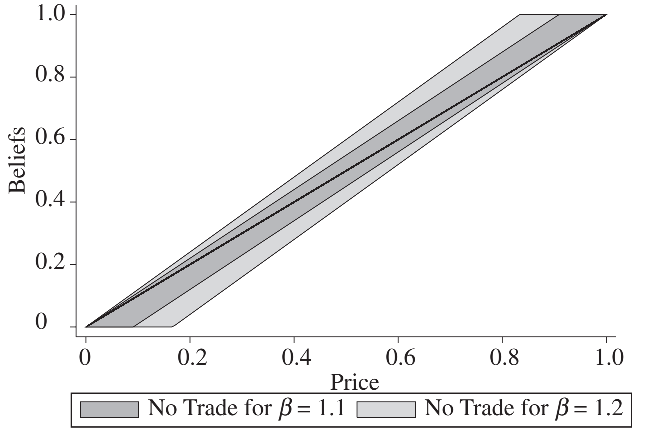

Fig. 1. *Regions of No Trade with Discounting*

How does this affect the market price? In the most general case, it is possible to define boundaries for the market price as a function of β:

Proposition 2. *If traders solve* (1) *and* (2), *and if F is the distribution of beliefs q of the traders, the market price is such that:*

$$F[1-(1-p)\beta] < 1-p < F(p\beta).$$

Proof. The market clearing condition is:

$$\int_{p\beta}^{1} \frac{y}{p} f(q)\mathrm{d}q = \int_{0}^{1-(1-p)\beta} \frac{y}{1-p} f(q)\mathrm{d}q. \tag{3}$$

This reduces to $(1-p)[1-F(p\beta)] = pF[1-(1-p)\beta]$.

The Proposition can be proven by contradiction. Suppose $1-p < F[1-(1-p)\beta]$, then $F[1-(1-p)\beta][1-F(p\beta)] > pF[1-(1-p)\beta]$ and, as a consequence $F(p\beta) < 1-p < F[1-(1-p)\beta]$.

However, when β is strictly greater than 1, $1-(1-p)\beta < p < p\beta$, which implies that $F[1-(1-p)\beta] \le F(p) \le F(p\beta)$. Hence, $F[1-(1-p)\beta] \le F(p\beta)$, which contradicts the previous result.

When $\beta \mapsto 1$, the model is identical to the one proposed by Manski (2006). The following results can be established.

Corollary 1. *Let* $\beta = 1$ *in* (2) *and let F be the cumulative distribution function of beliefs of the traders. The equilibrium price p satisfies the condition:*

$$p = 1 - F(p).$$

This Proposition is established by Ali (1977) (4) and Manski (2006) (1).

Proposition 3. *Let* $\beta = 1$ *in* (2). *The equilibrium price p differs from the median m of the distribution of beliefs F:*

(*i*) *If* $p \geq 1/2$, *then* $p \leq m$;
(*ii*) *If* $p \leq 1/2$, *then* $p \geq m$.

This Proposition corresponds to Theorem 2 of Ali (1977). It shows that in a one-period model, the equilibrium price in a prediction market may not represent the underlying median belief of the traders. Thus, a longshot bias may exist with risk-neutral traders.

It is possible to go further and to predict how the market price is likely to be influenced by an increase in β. In general, the no-trade region is not symmetric around p. Specifically, the no-trade region is larger above the price when the price is high and larger below the price when the price is low. This is likely to create an imbalance between supply and demand. For high prices, for example, potential buyers are more likely to abstain than potential sellers, because the no-trade area is larger above the price than below. The asymmetry in the no-trade area could therefore result in an increase in the longshot bias. If the demand decrease is greater than the supply decrease, prices would move downward. The reverse would happen for prices below 50%.

The following gives a sufficient condition regarding the distribution of beliefs for the no-trade region to be asymmetric as described above:

Condition 1. *Let f be the density function of the distribution of beliefs on the possible prices and p_1 be the actual equilibrium price on the market when $\beta = 1$:*

(*i*) *If* $p_1 > 1/2$, $f'(p_1) > 0$.
(*ii*) *If* $p_1 < 1/2$, $f'(p_1) < 0$.

Any unimodal distribution symmetric around m is consistent with Condition 1 (this class includes for instance many kernel functions). Proposition 3 implies that for a unimodal and symmetric distribution of beliefs, the price p will be between the median m of the distribution and 1/2. It is trivial to see that Condition 1 is respected in such situations.

Proposition 4. *If traders solve* 1 *and* 2, *and if Condition* 1 *holds, an increase in β in the neighbourhood of* 1 *increases the price if $m < p < 1/2$ and decreases the price if $1/2 < p < m$.*

Proof. From the market equilibrium condition, write the difference between demand and supply as follows:

$$\Phi = \int_{p\beta}^{1} \frac{y}{p} f(q)\mathrm{d}q - \int_{0}^{1-(1-p)\beta} \frac{y}{1-p} f(q)\mathrm{d}q = 0. \tag{4}$$

Using the implicit function theorem,

$$\frac{\mathrm{d}p}{\mathrm{d}\beta} = -\frac{\mathrm{d}\Phi/\mathrm{d}\beta}{\mathrm{d}\Phi/\mathrm{d}p}.$$

The effect of an increase in price on the market balance is unambiguous: $\mathrm{d}\Phi/\mathrm{d}p \leq 0$. The derivative $\mathrm{d}p/\mathrm{d}\beta$ has the sign of

$$\frac{\mathrm{d}\Phi}{\mathrm{d}\beta} = f[1-(1-p)\beta] - f(p\beta). \tag{5}$$

Condition 1 implies that for $\beta = 1 + \varepsilon$ with ε small and positive, $f(p\beta) > f[1-(1-p]\beta)$ if $p > 1/2$ and symmetrically $f(p\beta) < f[1-(1-p)\beta]$ if $p < 1/2$.

An increase in the discount factor for values close to 1 will therefore lead to an increase in the longshot bias.

Proposition 4 may be intuitively explained by the effect of the disappearance of trade. With an increase in β, the no-trade region expands around the price p. Under Condition 1, when $p > 1/2$, the decrease in demand is larger than the decrease in supply, leading to a drop in price. Conversely, when $p < 1/2$, the decrease in supply is larger than the decrease in demand, leading to an increase in price.

It is important to stress that the technical Condition 1 is consistent with a wide range of belief distributions. Condition 1 holds for unimodal distributions on [0,1] that have positive skewness when the median m of the distribution is above 1/2, and a negative skewness when m is below 1/2. This case includes, for instance, all beta distributions $F_B(a,b)$ with $a > 1$, $b > 1$. Moreover, as saw above, Proposition 3 ensures that any unimodal distribution symmetric around m is consistent with Condition 1.

To check that this result is robust to risk-averse utility functions and different levels of discounting, we performed a series of market-price simulations as a function of the mean beliefs of the traders and the discounting parameter β. We created fictitious markets with 10,000 traders having beliefs distributed as $F(q)$ and a specified utility function and discount factor β. We considered risk neutral and log utility functions. Beliefs were distributed according to a Beta distribution $B(\mu,1-\mu)$ where $\mu = \mathrm{E}(q)$ represents the traders' mean beliefs. The discount parameters chosen ($\beta = 1.1$ and 1.2) are in the range of the values found in the empirical literature (Andersen *et al.*, 2008). A numerical algorithm calculated the market clearing price. All simulations show the same effect of discounting on prices (Figure 2). Discounting systematically distorts the equilibrium price, with the anticipated stronger effect for prices near zero or one.

3. Data

Empirical analysis of long-term prediction markets (expiration beyond one month) is complicated by the fact that such markets are much less prevalent than short-term markets. For example, markets for US presidential elections may open years before election day but only six US Presidential elections have taken place since the creation of the Iowa Electronic Markets in 1988. For this reason, empirical studies of long-term prediction markets have so far relied primarily on relatively small samples. Several such studies have found that the prices in such prediction markets are in general not far from the empirical probability of the underlying event (Servan-Schreiber *et al.*, 2004; Wolfers and Zitzewitz, 2006*a*). However, the small sample sizes in these studies result in low levels of precision. In particular, although the theory may predict some specific miscalibration for some range of prices, including prices near zero or one, the theory cannot be adequately tested without a precise local estimation procedure that in turn requires a large data set.

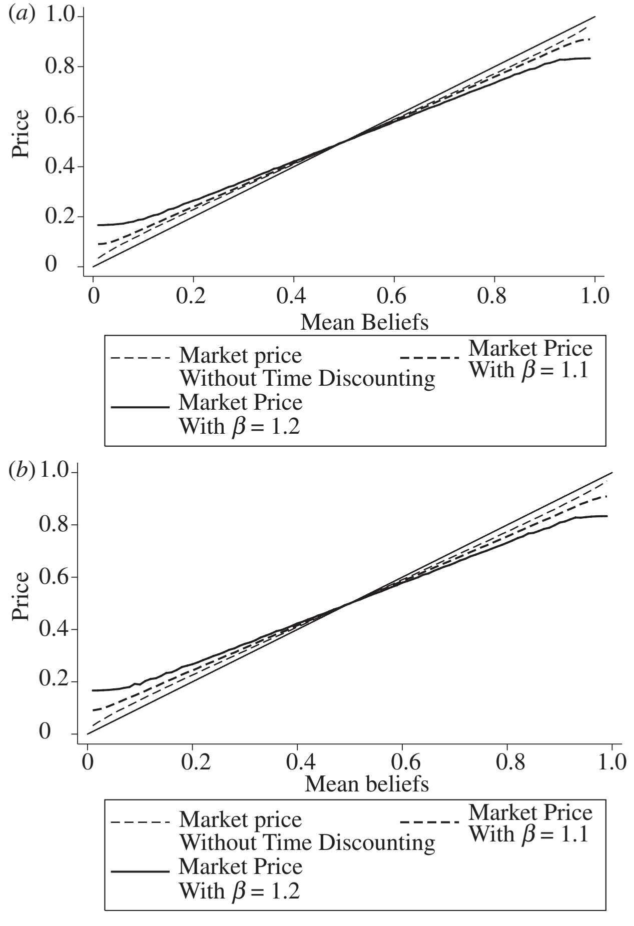

Fig. 2. *Simulation of Market Prices with Traders who are Risk Neutral (Top) and Risk Averse with Log Utility Functions (Bottom)*

Our data set includes transactions from 1,883 Intrade markets on future events. These markets are often linked to other markets in what we call a competition, in which the outcomes of the markets are interdependent. For instance, if two candidates are vying for a Senate seat, then the two markets (one for each candidate) constitute a competition; they are bound to each other by the fact that only one candidate can win. Our data also contain few markets for state-level results of a given election (e.g. 'GW Bush wins in Florida in 2004' and 'GW Bush wins in Virginia in 2004'). We consider them as not independent from each other and we cluster them in the same competition.

In the remainder of this article, we will use the following terminology: A *market* is an online tool where contracts on a future event are exchanged (e.g. 'Candidate A will win the 2004 US presidential election'). A *competition* is a set of interdependent markets (e.g. '2004 presidential election'). A *contract* is an asset exchanged in a market. The contract is linked to the realisation of a particular outcome of a specified future event. If the indicated outcome occurs, the contract's value is 1, and zero otherwise. A *transaction* occurs when two traders exchange a specified number of contracts at a given price between 0 and 1.

Our raw data set contains data for 1,883 markets and 512,828 transactions. From this raw sample, we eliminated markets with very low liquidity (less than five observations, or a total volume of less than 10 contracts). We gathered markets in competitions of interdependent markets, leaving us with a final data set of 597 competitions, 1,787 markets and 512,612 transactions. Low liquidity may create several problems, but for our purposes, low liquidity implies few traders and volatile prices. In either case, price data from such markets may not represent traders' median beliefs with adequate precision Christiansen (2007).

Table 1 presents some descriptive statistics for our data set, including subsets of sports, political and other[4] markets. Figure 3 shows a histogram of market prices for the entire data set. Many markets are from competitions with more than two possible outcomes, resulting in a distribution of prices that is skewed towards zero. Because only one outcome is possible, more markets end with a final price of zero than a final price of 1.

4. Statistical Analysis

4.1. *Analysis Principles*

In this Section, we focus on the empirical estimation of the calibration of the market prices[5] Let A be a state of nature taken from the set $\mathcal{A}$ of all possible states of nature. Let y be the final value of the contract on a given market, $y \in \{0,1\}$. A well-calibrated market price is defined by

$$p\,|_A = \mathrm{E}(y\,|\,A). \tag{6}$$

When market prices are well calibrated according to (6), they are accurate estimates of the probability of the contract's success. This relation cannot be tested directly, though, because $\mathrm{E}(y \mid A)$ is unknown. However, every market has a resolution of the underlying event, and hence prices can be compared with the corresponding contracts' final outcome. Let $\mathcal{A}_p$ be defined as $\{A \in \mathcal{A}_p \Leftrightarrow p\,|_A = p\}$, for each $A \in \mathcal{A}_p$, we have:

$$p = \mathrm{E}(y\,|\,p). \tag{7}$$

If the price is a calibrated estimate of the probability, contracts exchanged at a given price should, on average, yield a result equal to that price. Said differently, when buying a contract at a price p, the probability of success of the contract should be p. If the price departs systematically from the probability of the event, observed prices will not satisfy (7) uniformly. If there is a longshot bias, prices will be too high for low probabilities and too low for high probabilities:

[4] The 'other market' category includes a variety of different markets, such as markets on financial or economic events or other popular events such as the Academy Awards or reality TV shows.

[5] We stress that we do not address the calibration of the traders as such. We do not study the distribution of traders beliefs but only the link between prices and observed probabilities.

Table 1
Descriptive Statistics

	All markets			Sport markets			Political markets			Other markets		
	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max	Mean	Min	Max
Vol. per transaction	24.17	1	1,687	29.51	1	1,687	14.29	1	564	17.75	1	120
Vol. per contract	115,176.37	16	1,550,189	188,137.55	16	1,236,944	15,635.79	19	1,550,189	2,184.84	17	32,704
No. obs per contract	5,014.12	5	67,547	7,942.98	5	67,547	1,486.18	5	62,956	142.71	5	2,609
Price	22.18	0.10	99.90	16.20	0.10	99.90	39.23	0.10	98.80	25.08	0.10	95.00
Duration	205.00	27	649	236.38	36	431	256.55	33	649	42.00	27	107
No. of transactions	512,612			336,575			140,041			35,996		
No. of competitions	597			241			142			214		
No. of contracts	1,787			1,064			302			421		

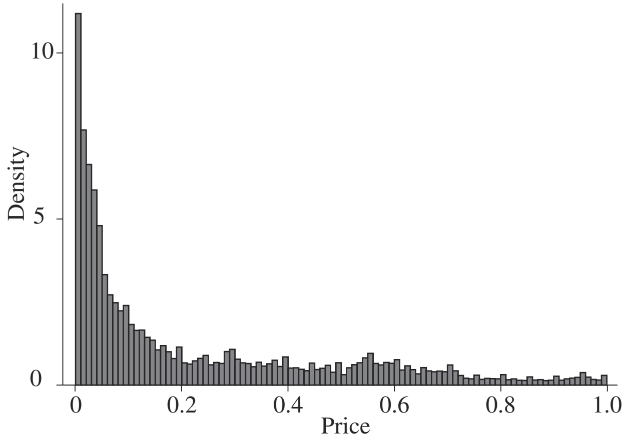

Fig. 3. *Distribution of Prices*

$$\begin{aligned} p &> \mathrm{E}(y) \text{ for } p \in [0, p_0], \\ p &< \mathrm{E}(y) \text{ for } p \in [p_1, 1], \end{aligned} \tag{8}$$

for some $0 < p_0 \leq p_1 < 1$. As above, (8) cannot be tested directly. However, it is possible to estimate $\mathrm{E}(y \mid p)$, and (8) implies:

$$\begin{aligned} p &> \mathrm{E}(y \mid p) \text{ for } p \in [0, p_0], \\ p &< \mathrm{E}(y \mid p) \text{ for } p \in [p_1, 1]. \end{aligned} \tag{9}$$

To assess the calibration of market prices, we estimate the empirical probability associated with each price, $\mathrm{E}(y \mid p)$, using both parametric and non-parametric techniques.

The principle of this analysis, relying on the comparison of $\mathrm{E}(y \mid p)$ to p, systematises previous approaches in the study of the calibration of odds on betting markets and prices on prediction markets (Coleman, 2004). Initial studies of the longshot bias in horserace betting used the ranks of starters, ranking starters by their odds (derived from bets placed before the race) and numbering them from 1 to N (Griffith, 1949; Ali, 1977; Bird *et al.*, 1987; Terrell, 1998). For each rank, the observed relative frequency of winning is compared with the probability of winning as derived from the odds. A discrepancy between the relative frequency of winning and the odds-based probability [violation of (7)] indicates miscalibration, and systematic miscalibration suggests a bias. A difficulty with this approach is that starters with identical ranks may have quite different odds. The binning of starters by rank can result in substantial variance in odds-based probabilities within each bin, which in turn reduces the precision of the estimation.

Other studies use odds/market price directly to bin the starters (Snyder, 1978; Williams and Paton, 1997; Jullien and Salanie, 2000) and compare the implied probabilities p to the relative-frequency-based estimates of $\mathrm{E}(y \mid p)$. A problem with this solution is that each race includes several starters but only one winning outcome. Using the odds of all the horses in the race can lead to a spurious analysis, for example, when two horses from a given race are placed in the same bin (Busche and Hall, 1988). The outcomes (winning or losing) for each of these horses are by definition not independent events.

The binning approach has also been used to analyse prediction markets by placing contracts with similar prices into the same bin. Such studies face the same problem of non-independence of contracts within one market.

In this article, we use a non-parametric local linear regression to estimate $\mathrm{E}(y \mid p)$. Our approach provides precise estimates, even for very high and very low probabilities. An innovation of our empirical strategy is the use of a clustered bootstrap to account for the non-independence of contracts within markets whose outcomes are dependent (e.g. all markets linked to one election). Using groups of non-independent markets as clusters in a bootstrap resampling, we are able to estimate a confidence interval for the entire calibration curve.

We also use a parametric estimation procedure to study factors that may influence price calibration. We use a clustered bootstrap here as well, allowing us to calculate accurate confidence intervals for the parameters.

The computation of a confidence interval by clustered bootstrap allows us to keep all the information in the data set while controlling for the non-independence of prices within clusters. It is a significant improvement relative to the previous literature.

4.2. *Non-parametric Estimation*

4.2.1. *Methodology*

With a large number of observations, it is possible to calculate a local estimate of the conditional expectation $\mathrm{E}(y \mid p)$, and thereby determine precisely how much the data deviate from (7). We use a local regression estimator, which has the advantage of being more flexible than a kernel regression estimator. It is also more accurate than kernel regression for transaction prices at the high and low extremes, a property that is important for our analysis. (For a description of the boundary problem with non-parametric estimators, see Härdle (1992)).

The local regression estimator estimates a local regression line for each price p, by estimating a constant parameter $\beta_0(p)$ and a slope parameter $\beta_1(p)$. To do so, the following program is solved:

$$\min_{\beta_0,\beta_1} \sum_{i=1}^{n} K_h(p-p_i)[y_i-\beta_0-\beta_1(p_i-p)]^2,$$

where p_i represents the ith of n observations used in the estimation; h is the width of an estimation window around p and K_h is a kernel (i.e. a weighting function) defined by[6]:

$$K_h(p-p_i)=\frac{3}{4}\left[1-\left(\frac{p-p_i}{h}\right)^2\right]\mathbb{1}_{\{|p-p_i|\leq h\}}.$$

[6] In all of our estimations, we use an Epanechnikov kernel. The estimations are, however, not sensitive to the particular kernel chosen.

The smaller the window size h, the more flexible the estimation. As the size of the window decreases, though, the variance of the estimated curve increases. Thus, the choice of window width requires a trade-off between the flexibility of the estimation and the smoothness of the estimated curve. We use a window width of 0.10. This allows for a flexible estimation but with a reasonably small variance. The estimator of $\mathrm{E}(y \mid p)$ is then given by

$$\mathrm{E}\widehat{(y \mid p)} = \widehat{\beta}_0.$$

Formally, this estimator is equivalent to a local weighted average of the dependent variable $y \in \{0,1\}$ (Fan, 1993). To assess the precision of this non-parametric estimation, we calculate a confidence interval for each of 99 prices (0.01–0.99). The calculation follows the technique proposed by Härdle (1992), using a clustered bootstrap approach. Because of the non-independence of the prices within each competition, we use competitions as our clusters. For example, if there are five separate markets, one for each possible winner of an Academy Award, then those five markets would constitute one cluster. For political markets, we considered the markets linked to the re-election of George W. Bush in 2004 as one competition. In addition, all markets associated with the 2004 Democratic primaries were considered a separate single competition.

There are huge variations in the number of transactions across competitions. Some, like the US presidency, have tens of thousands of transactions, and others have only a few dozen. Because of this disparity, for each point from 0.01–0.99 we randomly select 10 transactions (prices) per competition, for a total of $n = 6{,}020$ transactions. Doing so avoids giving too large a weight in the estimation to large competitions.[7]

4.2.2. *Results: overall calibration*

Figure 4 shows the results of the non-parametric estimation of the calibration of prediction market prices. The heavy dashed line is computed as the average of the different estimates, while the grey-shaded area is the confidence interval computed by the clustered bootstrap procedure with 400 replications. The S-shaped curve indicates a clear longshot bias. For example, a price of 0.20 is on average associated with a relative frequency of 15.3%; conversely, a price of 0.80 is on average associated with a relative frequency of 87.4%. These results indicate miscalibration at a considerably higher level than reported previously for short-term markets.

Noticeably, this S-shaped curve does not present a bias at the very extreme prices towards 0 and 1 similar to our simulations displayed in Figure 2. In our results, prices seem, at the limit, well calibrated in 0 and 1, while it is not the case in our simulations. Our model does not explain this particular feature of the bias in prices' calibration. As a matter of fact, our model is unable to explain prices close to extremes as it predicts that the liquidity disappears on the markets in such situations. For instance, for prices close to 1, no traders are willing to buy long-term contracts when their beliefs are between p and $p\beta > p$. For a given level of time discounting β, there

[7] To estimate the confidence interval by bootstrap, in each iteration, we draw new samples of prices for each contract. As can be seen from the narrow intervals in Figure 4, randomly choosing prices results in little variation in the point estimate.

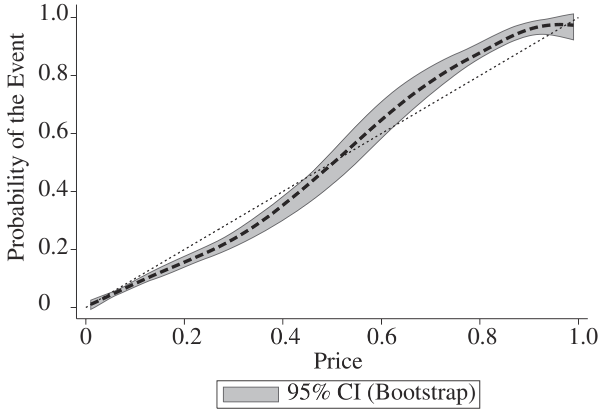

Fig. 4. *Non-parametric Estimation of the Calibration Curve*

is a maximum price which can be observed on the market in our model: $p^* = 1/\beta$. Above this price, there are no more traders willing to buy the contract. Interestingly, this is not what is observed in our estimation where a S-shape curve is observed. One possibility is that, unlike in our simple model, there is an heterogeneity in discount factors in the population of traders. In such a situation, traders with the lowest discounting factors will be the only ones trading in the extreme ranges of prices, hence leading to a smaller level of bias in these areas. This explanation is however not entirely satisfying as no rational trader should have a time discounting factor lower than the interest rate in the economy. Another solution is that some traders have a non-economic motive such as a preference to gamble, possibly a preference to gamble on specific events. Some traders may also have a preference on the level of the price itself if they want the market price to give a signal to third parties (other traders, general public, media). Whilst these traders with non-economic motives may not be very numerous, they could be the only ones left on the market in the extreme ranges of prices for long-term prediction markets. Trades above the economically rational threshold p^* from our model could then typically produce a calibration curves in S, as the possible distance between the upper (lower) bound of the calibration curve and the 45° line is necessarily decreasing as we get closer from prices close to one (zero). This could give the impression that the prices are better calibrated at the extreme even if traders in these ranges of prices tend to trade for non-purely economic motives.

4.2.3. *Results: variations in miscalibration across markets*

To go further, we can consider variables potentially related to miscalibration. Looking across three different categories of competitions (sports, political, other) and ignoring the time dimension, we find a meaningful difference only when comparing political to non-political competitions. Figure 5 shows the non-parametric estimations for these two categories. Market prices for political events show a very strong longshot bias, while the other markets are only slightly miscalibrated. The extent of the bias in political markets is noticeably greater than the overall results in Figure 4. Here, for a price of

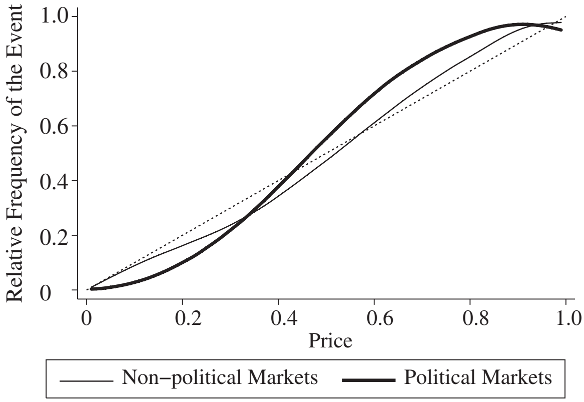

Fig. 5. *Non-parametric Results for Political and Non-political Markets*

0.20, the relative frequency of the event occurring is 10.9% and, for a price of 0.80, the relative frequency is 92.8%.

To our knowledge, no reason is given in the literature to explain why political prediction markets would show a stronger longshot bias. A possible explanation could be the existence of more manipulative traders on these markets. This phenomenon and possible explanations merit further study.

4.2.4. *Evolution of price calibration over time*

Following the results of Section 1, one should expect that price calibration worsens as time to contract expiration increases. To test this non-parametrically, we divided the sample between prices with more or less than 100 days left to expiration. Figure 6 shows the overall result, and Figure 7 shows the result when looking specifically at the categories of political and non-political markets (upper panels) and when splitting the non-political markets between sport ones and non-sport ones (lower panels). The subdivision of the sample in smaller subsample increase the variance of the estimates, in particular for prices observed more than 100 days before maturity. However, overall the results are consistent with our model's predictions; the longshot bias is stronger for the longer time horizon, and the effect appears to hold for both political and non-political markets.

4.3. *Parametric Estimation*

4.3.1. *General methodology*

To model a possible non-linear relationship between observed prices and empirical frequencies, we use the Lattimore *et al.* (1992) function which is adapted to model the S curve that one could expect from a longshot bias. This function allows for the calibration curve to be systematically below the 45° line before a crossing point and above the 45° line after this crossing point. The function is characterised by two parameters: γ which controls the curvature and δ which controls the crossing point:

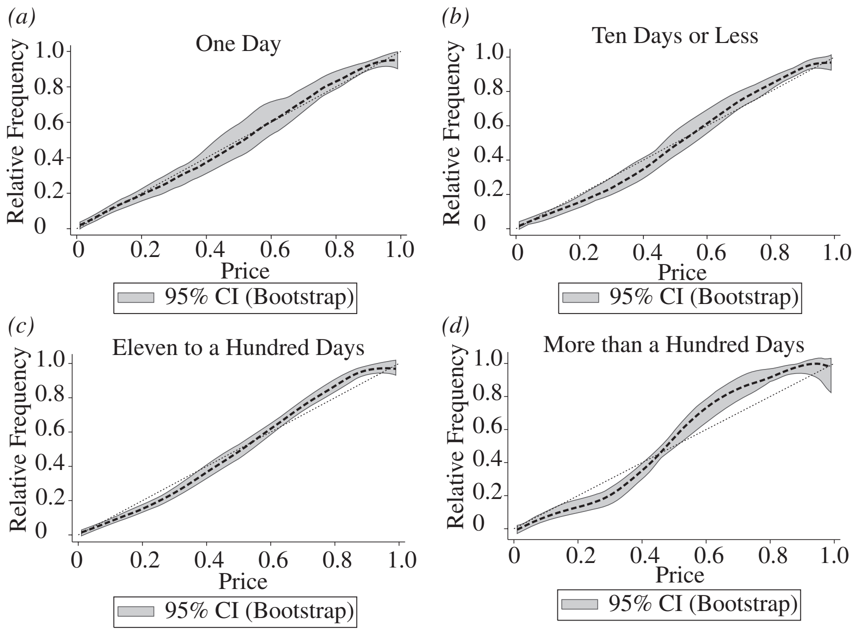

Fig. 6. *Calibration of Prices over Different Times to Expiration*

$$p = \frac{\delta\pi^{\gamma}}{\delta\pi^{\gamma} + (1-\pi)^{\gamma}}. \tag{10}$$

In this expression, $\gamma = 1$ implies no curvature (the function corresponds to the 45° line), and $\delta = 1$ implies a crossing point of 0.5. The function is invertible:

$$\pi = \frac{1}{[\delta(1-p)/p]^{1/\gamma} + 1}. \tag{11}$$

As a result, one can estimate parameters γ and δ using maximum likelihood. Given the non-independence of the prices (observations), though, the product of likelihoods over all observations is a pseudo-likelihood. Let Y be a binary variable indicating the success of contract i, $Y_i \in \{0,1\}$. The pseudo-likelihood PL is

$$PL = \prod_i \left\{\frac{1}{[\delta(1-p)/p]^{1/\gamma}+1}\right\}^{Y_i} \left\{1 - \frac{1}{[\delta(1-p)/p]^{1/\gamma}+1}\right\}^{1-Y_i}. \tag{12}$$

Following the same principles as above, these estimations are performed with a small number of observations ($n = 10$) selected randomly from each competition. As before, this procedure ensures that the results are not driven by a few competitions that have a large number of transactions. Also as above, to compute confidence intervals, we use a clustered bootstrap with the competitions as clusters. For each bootstrap sample, we draw a new set of observations within each competition to ensure that our results do not depend on a specific selection of transactions within each competition.

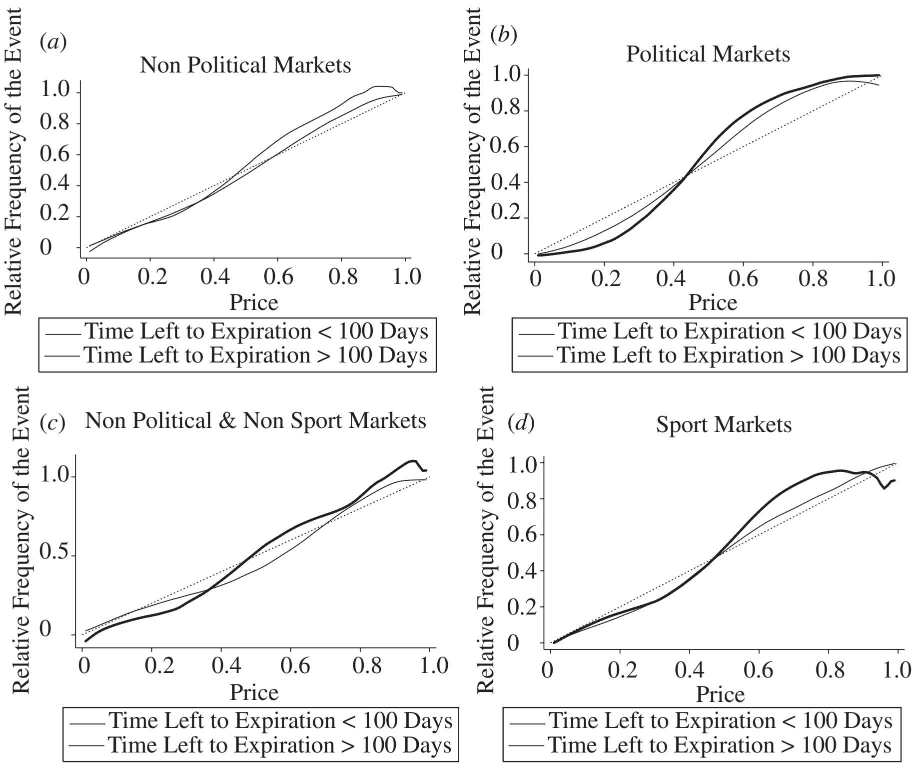

Fig. 7. *Calibration of Prices over Different Times to Expiration, by Market Type*

4.3.2. *Empirical results*

Table 2 presents the results for the maximisation of the pseudo-likelihood (12). The first column shows the results when γ is assumed to be a constant. The results indicate that the estimate of δ is not significantly different from 1, in turn indicating that the crossing point is not significantly different from 0.5. However, γ is significantly different from 1, indicating a significant curvature of the calibration curve, confirming the non-parametric results.

In columns 2, 3 and 4 of Table 2, the hypothesis that γ is a constant is relaxed to study whether the curvature of the calibration function varies across different categories of contracts. To accomplish this, we replace the constant γ with a linear function of potential explanatory variables. We first include the time left to expiration. This variable is negative and significant, indicating that γ decreases as time to contract expiration increases; the lower γ implies more curvature in (10). This confirms our model's prediction regarding traders' discounting of future winnings and the corresponding result in the non-parametric analysis above. In column 3, we add variables to indicate the event categories sports (241 competitions, 1,064 markets), political (142 competitions, 302 markets) and other markets (214 competitions, 421 markets). Table 2 shows that when the corresponding variables are included, the coefficient for political markets is the only one that is significant, confirming the non-parametric result that

Table 2
Estimation of the PL Model

Coefficients	Models (1)	(2)	(3)	(4)
μ				
Constant†	0.94**	0.94**	0.95**	0.96**
γ				
Constant†	0.80**	0.88**	0.98**	0.94**
Time left to expiration		−0.001*	−0.001*	−0.001*
Sport			−0.09	−0.10
Politics			−0.33*	−0.35*
Volume				> 0.01
Volume total				> 0.01
N competitions	597	597	597	597
N prices	512,612	512,612	512,612	512,612

Notes. Significant at 10%, *5%, **1%, ***0.1%.
†Test of significance relative to one is indicated in subscript.
CI calculated by clustered bootstraps with the independent competitions as clusters.

political markets display a stronger longshot bias than non-political markets. The coefficient for the time left to expiration remains negative and significant.

In column 4, we include two variables related to market volume, volume per transaction (number of contracts, recorded for each transaction price) and the total volume of the market (total number of contracts traded, recorded for each market). Greater volume per transaction suggests traders more willing to trade on their information, and larger market volume suggests more traders. Taken together, these would indicate more liquidity in the market and more information, suggesting that the market would be likely to incorporate more information into prices and do so efficiently, leading to better calibration.

Our results show that neither volume variable has a significant effect on calibration of the market prices. In our data set, the 10th and 25th percentiles for market volume are 269 and 858 respectively. This may provide some guidance for corporate prediction markets, which are often characterised by relatively low volumes. According to our results, low volume appears not to introduce additional systematic bias in market prices, but low volume may still require a substantial amount of trading for a corporate market.

5. Economic Analysis

The previous analysis has shown that prediction market prices can be significantly miscalibrated. Does this imply that one can design investment strategies to exploit the bias to achieve positive economic returns? To test for this, we divided our sample into two halves. The first subset includes markets that began between the 22 August 2002 and the 14 January 2005 (median starting date in our data set). The second subset consists of markets that began between the 15 January 2005 and the 6 February 2007.

A non-parametric analysis of the first sample revealed a significant S-shaped calibration curve, consistent with the results for the entire data set. This analysis suggests that one should buy when the price is between 0.08 and 0.10 or above 0.61 and sell when the price is between 0 and 0.08 or between 0.10 and 0.61. Call this strategy A. Given the S-shape, it is not possible to reject the hypothesis that the actual calibration curve is a standard longshot bias, and so we also implement a simple strategy B consisting of buying over 0.50 and selling below 0.50. Arguably these strategies are simplistic, an optimal strategy would require one to invest more in assets where expected profits are higher (like the Kelly criterion strategy). Therefore, these simple strategies give us a conservative estimate of the possible profits.

When applying both strategies to the second half of the data set, strategy A (B) yields an average rate of return of 5.25% (9.55%). The fact that strategy B performs better than A suggests that the key issue is the presence of a longshot bias. Moreover, the calibration curve estimated using the first sample deviates only slightly from a simple longshot bias and is not significantly different. Thus, the strategies seem to produce positive returns due to a persistent longshot bias.

The analysis above assumes a discount factor of one. As the discount rate increases, though, the rate of return decreases very quickly. A discount rate of 10% yields a gain of 2.47% on average with strategy A. A discount rate of 15% completely cancels out the gains for strategy A, and a discount rate of 25% completely cancels out the gains for strategy B. Thus, if traders have discount rates of this magnitude, there is no free lunch; the miscalibration would be insufficient for traders to achieve excess returns.

Of course, we cannot say whether traders have high discount rates and hence that no systematic opportunity of gain exists, or whether the discount rate only cancels a part of the possible gains, leaving some margin for profitable strategies. A discount rate of 15–25% is somewhat high, even more given that the strategies we tested are simplistic and underestimate the maximum level of profits which can be reached due to the discrepancy between market prices are probabilities. Empirical studies of individual discount rates have shown that they can be quite high (Andersen *et al.*, 2008). In the case of prediction markets, even if the trader cares less about an economic return *per se*, a high implicit discount rate may stem from the trader's desire to use his or her money to play repeatedly, which would indicate a high option value for the money.

6. Discussion and Conclusion

Our study is the first to show that prediction markets are mechanically affected by time discounting. First, the domain of beliefs compatible with an incentive to trade shrinks as the time to expiration increases. For this reason, long-term prediction markets are typically characterised by a low level of trading volume. At the limit, long-term prediction markets may be too unattractive to generate any trade at all. Second, the reduction of trade for long-term prediction markets is asymmetric with traders willing to buy high price contracts more likely to leave the market as the time to expiration of the market increases. As a consequence, prices of long-term prediction markets are systematically biased towards 0.50.

Using a large dataset of prediction market prices, we have been able to estimate the calibration of market prices more precisely than has been possible in previous studies. As predicted, prices on long-term prediction markets exhibit systematic biases when considered estimates of underlying probabilities. Our study casts doubts on the use of prediction markets to forecast long-term events. Prediction markets for events occurring a year or more in the future will either fail to generate a price or the price will be systematically biased.

In regards to these results, several questions can be raised relative to the possible limitations of our model and relative to the meaning of the empirical finding.

Our model shows that, as a consequence of time discounting, calibration of prices improves as the time to expiration decreases. One could also consider other explanations. For instance, calibration could be a function of the amount of information available; more information regarding the likelihood of future events might lead to better calibration. Our results show that the improvement in calibration over time can be explained without appealing to information effects. However, understanding the possible relationship between information and calibration in prediction market prices is an interesting question for future research.

Importantly, our model is unable to explain what happens for extreme prices, close to 0 or close to 1. Our model and simulations predict an absence of extreme prices, even when trader's beliefs are extreme. When assuming rational traders only motivated by economic profit, time discounting should prevent trades to occur in these ranges of prices as sellers, in the former case, and buyers, in the latter case, withdraw from the market. For example, considering a situation where the interest rate on saving accounts is 5%, there should be no trade for contracts whose prices are above 0.95 or below 0.05 where there is one year or more until expiration. However, we observe transactions at these prices in markets that expire in a year or more, and our empirical analysis has found an S-shaped calibration curve, over the whole range [0,1]. The economic unattractiveness of long-term markets leads to questions about the possible adverse selection of traders' motivations. These traders may have a utility of gambling *per se*, or the intent to affect the prices. This, in turn, raises questions about the effect of these motivations on the calibration of market prices.

Our result of a time-specific bias on prediction market and the S-shaped pattern observed, typical of a longshot bias, could look similar to the result of Rubinstein (1976) on the pricing bias on short maturity out-of-money calls. This bias has been interpreted as a form of favourite/longshot bias by later papers (Berg and Rietz, 2002; Cain *et al.*, 2002). However, the two results are different in nature. Rubinstein found a bias, for out-of-money calls close to expirations. On the contrary, we find a bias for contracts which are far away from expiration. In the case of out-of-money calls, these are options which are unlikely to be profitable but have the potential to provide a large profit with a very small probability. These longshot characteristics could well increase buyers' interest in such contracts, creating an increasing difference between the theoretical price of the call and its observed price. In the literature on longshot bias, at least two different behavioural explanations have been put forward to explain the existence of such a bias: a preference for skewness (Golec and Tamarkin, 1998) and a propensity to overestimate small probabilities (Thaler and Ziemba, 1988). In the case of our long-term prediction market prices, we show that such a bias should naturally emerge due to the standard preferences of traders (time preferences), which makes these contracts less appealing to sellers, without the need to use non-standard preferences.

Finally, our results can be used to look for ways to overcome the existing limitations of long-term prediction markets. An obvious solution would be for long-term prediction markets to accommodate discounting by paying interest on committed balances from traders. In fact, for balances over US $20,000, Intrade does pay interest at the current 'bank rate for interest-accruing checking accounts'. In spite of this, we still observe a more pronounced bias on long-term prediction markets. A possible reason is that traders' subjective discount rate may be significantly above the bank interest rate. This is in particular the case if traders have a preference for gambling (Diecidue *et al.*, 2004). Investing in a one-year contract amounts to losing the opportunity to play repeatedly during the same period. Another possible reason is that traders with large accounts may be more likely to be higher performing traders. High-performing traders may have a positive expected return on their trades, and, if so they would face a real (not a subjective) opportunity cost on long-term contracts. If they have an expected return of just 2% per trade, in just five transactions they would average a bit more than 10% return. This possible expected return is lost if their capital is frozen because of investing in a long-term contract. As a result, the propensity of high-performing traders to invest in long-term prediction markets might be limited. In summary, these arguments suggest that either restricting the payment to high balances is not enough, or that the interest rate paid on balances needs to be considerably higher than the bank interest rate to eliminate bias in long-term prediction markets.

Queensland University of Technology
Duke University

Submitted: 31 January 2011
Accepted: 21 February 2012

References

Ali, M.M. (1977). 'Probability and utility estimates for racetrack bettors', *Journal of Political Economy*, vol. 85(4), pp. 803–15.

Allen, F. and Gale, D. (1992). 'Stock-price manipulation', *Review of Financial Studies*, vol. 5(3), pp. 503–29.

Andersen, S., Harrison, G.W., Lau, M.I. and Rutstrom, E.E. (2008). 'Eliciting risk and time preferences', *Econometrica*, vol. 76(3), pp. 583–618.

Berg, J.E. and Rietz, T.A. (2002). 'Longshots, overconfidence and efficiency on the Iowa Electronic Market', Working Paper, Tippie College of Business, University of Iowa.

Bird, R., McCrae, M. and Beggs, J. (1987). 'Are gamblers really risk takers', *Australian Economic Papers*, vol. 26(49), pp. 237–53.

Busche, K. and Hall, C.D. (1988). 'An exception to the risk preference anomaly', *Journal of Business*, vol. 61(3), pp. 337–46.

Cain, M., Law, D. and Peel, D. (2002). 'Is one price enough to value a state-contingent asset correctly? Evidence fromm a gambling market', *Applied Financial Economics*, vol. 12(1), pp. 33–8.

Camerer, C. (1998). 'Can asset markets be manipulated? A field experiment with racetrack betting', *Journal of Political Economy*, vol. 106(3), pp. 457–82.

Christiansen, J.D. (2007). 'Prediction markets: practical experiments in small markets and behaviours observed', *Journal of Prediction Markets*, vol. 1(1), pp. 17–41.

Coleman, L. (2004). 'New light on the longshot bias', *Applied Economics*, vol. 36(4), pp. 315–26.

Diecidue, E., Schmidt, U. and Wakker, P.P. (2004). 'The utility of gambling reconsidered', *Journal of Risk and Uncertainty*, vol. 29(3), pp. 241–59.

Fan, J. (1993). 'Local linear regression smoothers and their minimax efficiencies', *Annals of Statistics*, vol. 21(1), pp. 196–216.

Fox, C. and Clemen, R. (2005). 'Subjective probability assessment in decision analysis: partition dependence and bias toward the ignorance prior', *Management Science*, vol. 51(9), pp. 1417–32.

Gjerstad, S. (2005). 'Risk aversion, beliefs, and prediction market equilibrium', Dept of Economics, University of Arizona.

Golec, J. and Tamarkin, M. (1998). 'Bettors love skewness, not risk, at the horse track', *Journal of Political Economy*, vol. 106(1), pp. 205–25.
Griffith, R.M. (1949). 'Odds adjustments by American horse-race bettors', *American Journal of Psychology*, vol. 62(2), pp. 290–4.
Hansen, J., Schmidt, C. and Strobel, M. (2004). 'Manipulation in political stock markets–preconditions and evidence', *Applied Economics Letters*, vol. 11(7), pp. 459–63.
Hanson, R. and Oprea, R. (2009). 'A manipulator can aid prediction market accuracy', *Economica*, vol. 76(302), pp. 304–14.
Hanson, R., Oprea, R. and Porter, D. (2006). 'Information aggregation and manipulation in an experimental market', *Journal of Economic Behavior and Organization*, vol. 60(4), pp. 449–59.
Härdle, W. (1992). *Applied Nonparametric Regression*, Cambridge: Cambridge University Press.
Jullien, B. and Salanie, B. (2000). 'Estimating preferences under risk: the case of racetrack bettors', *Journal of Political Economy*, vol. 108(3), p. 503.
Kahneman, D. and Tversky, A. (1979). 'Prospect theory: an analysis of decision under risk', *Econometrica*, vol. 47, pp. 263–91.
Lattimore, P.M., Baker, J.R. and Witte, A.D. (1992). 'The influence of probability on risky choice', *Journal of Economic Behavior and Organization*, vol. 17(S 377), p. 400.
Manski, C.F. (2006). 'Interpreting the predictions of prediction markets', *Economics Letters*, vol. 91(3), pp. 425–9.
Oliven, K. and Rietz, T. (2004). 'Suckers are born but markets are made: individual rationality, arbitrage, and market efficiency on an electronic futures market', *Management Science*, vol. 50(3), pp. 336–51.
Ottaviani, M. and Sørensen, P.N. (2007). 'Outcome manipulation in corporate prediction markets', *Journal of the European Economic Association*, vol. 5(2-3), pp. 554–63.
Ottaviani, M. and Sørensen, P.N. (2009). 'Surprised by the parimutuel odds?' *American Economic Review*, vol. 99(5) pp. 2129–34.
Ottaviani, M. and Sørensen, P.N. (2010*a*). 'Noise, information, and the favorite-longshot bias in parimutuel predictions', *American Economic Journal: Microeconomics*, vol. 2(1), pp. 58–85.
Ottaviani, M. and Sørensen, P.N. (2010*b*). 'Aggregation of information and beliefs in prediction markets', mimeo, California Institute of Technology.
Rhode, P.W. and Strumpf, K.S. (2006). *Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data*, Chapel Hill, NC: University of North Carolina.
Rubinstein, M. (1976). 'The strong case for the generalized logarithmic utility model as the premier model of financial markets', *Journal of Finance*, vol. 31(2), pp. 551–71.
Servan-Schreiber, E., Wolfers, J., Pennock, D.M. and Galebach B. (2004). 'Prediction markets: does money matter?' *Electronic Markets*, vol. 14(3), pp. 243–51.
Snyder, W.W. (1978). 'Horse racing: testing the efficient markets model', *Journal of Finance*, vol. 33(4), pp. 1109–18.
Sobel, R.S. and Raines, T.S. (2003). 'An examination of the empirical derivatives of the favourite-longshot bias in racetrack betting', *Applied Economics*, vol. 35(4), pp. 371–85.
Sonnemann, U., Camerer, C., Fox, C.R., and Langer, T. (2011). 'Psychological biases affect economic market prices', Working Paper, University of Muenster.
Terrell, D. (1998). 'Biases in assessments of probabilities: new evidence from greyhound races', *Journal of Risk and Uncertainty*, vol. 17(2), pp. 151–67.
Thaler, R. and Ziemba, W. (1988). 'Parimutuel betting markets: racetracks and lotteries', *Journal of Economic Perspectives*, vol. 2, pp. 161–74.
Williams, L.V. and Paton, D. (1997). 'Why is there a favourite-longshot bias in British racetrack betting markets', ECONOMIC JOURNAL, vol. 107(440), pp. 150–8.
Wolfers, J. and Zitzewitz, E. (2006*a*). 'Five open questions about prediction markets', Federal Reserve Bank of San Francisco Working Paper Series.
Wolfers, J. and Zitzewitz, E. (2006*b*). 'Interpreting prediction market prices as probabilities', Working Paper No.11, Federal Reserve Bank of San Francisco.