

PREDICTION MARKETS: REALITY AND THEORY

Daniella Acker

Accounting and Finance
Discussion Paper 16 / 5

18 October 2016



School of Economics, Finance and Management
University of Bristol
8 Woodland Road
Bristol BS8 1TN
United Kingdom

PREDICTION MARKETS: REALITY AND THEORY

Daniella Acker*
University of Bristol

18 October, 2016

ABSTRACT

Data on individual trades in prediction markets relating to the 2008 and 2012 US Presidential elections reveal that traders vary enormously in their behavior. This contrasts with the standard prediction-market models, which assume relatively homogeneous participants who differ only in their beliefs and wealth. We show that risk-lovers have particularly strong distortionary effects on market outcomes even when beliefs are symmetrically distributed around the truth. Simulations of a model which allows traders to have different motives and tastes for risk indicate that including such traders produce the market outcomes we observe, such as herding, persistent contrariness, a skewed profits' distribution and favorite-long-shot bias. The attraction of such markets to risk-lovers means that caution must be exercised when using prediction-market prices for forecasting

Keywords: Prediction markets; Risk-lovers; Herding and contrariness;
Favorite-long-shot bias

JEL codes: G10, G12, G14, G17

*Contact details:
daniella.acker@bristol.ac.uk
Department of Accounting and Finance
School of Economics, Finance and Management
University of Bristol, 12 Priory Road
Bristol, BS8 1TN

PREDICTION MARKETS: REALITY AND THEORY

1. INTRODUCTION

The primary contribution of this paper is to the understanding of the workings of prediction markets, markets which allow members to trade on the outcome of real world events and which are sometimes used as a useful ‘laboratory’ to investigate wider issues, such as asset mispricing (see Brown and Yang, 2016, and Brown, 2014, for example). Prediction and betting markets have flourished in recent years, and a subset of them, political prediction markets, have become especially influential. A number of studies have presented evidence that the prices established in prediction markets exhibit some of the standard characteristics of market efficiency and accurately forecast final outcomes.¹ As a consequence political prediction markets have begun to rival opinion polls as the go-to measure of the direction of the ‘political wind’, eliciting headlines such as “Betting markets give Trump an 80% chance of securing the nomination” *Financial Times*, 29 February, 2016, and comments such as this by Summers, 2016, in an op-ed in the Washington Post: “... because markets aggregate the views of a huge number of participants, and because they are constantly assessing the future ..., they are like canaries in coal mines: very valuable in giving warning when conditions change. That is why several studies have shown that prediction markets do a better job of forecasting elections than pollsters ...”.² This apparent accuracy has led Snowberg, Wolfers and Zitzewitz (SWZ), 2013, to suggest that prices established in prediction markets might usefully be incorporated into more general economic forecasting models.

¹ See, for example, Tetlock, 2004, Borghesi, 2006, Leigh and Wolfers, 2006, Zitzewitz, 2006, Wolfers and Zitzewitz, 2007, and Berg *et al.*, 2008.

² https://www.washingtonpost.com/opinions/why-the-fed-needs-to-prepare-for-the-worst-right-now/2016/01/10/b754a962-b7a3-11e5-99f3-184bc379b12d_story.html

The initial motivation for this research was access to a highly detailed set of observations on a group of markets offered by Intrade.com on the 2008 and 2012 US Presidential elections.³

These observations cover almost the entire universe of trades relating to these elections, time stamped and labelled with the numerical (anonymised) ID of the buyer and the seller. This allows us to track individual trading behavior in great detail and reveals a degree of diversity that is hard to reconcile with standard models of prediction markets, such as that in Wolfers and Zitzewitz (WZ), 2007. Some traders appear highly biased, others neutral; some appear financially unsophisticated, others highly sophisticated; some appear timid, others reckless; and some begin trading early on, others only in the last few weeks of the contract. The aggregate market characteristics that we document are also inconsistent with the standard model: markets are polarised, with an imbalance between the number of buyers and the number of sellers; furthermore, certain traders exhibit persistent contrariness, trading consistently against the crowd. Profits are highly skewed and losses tend to be especially common amongst persistent contrarians; and finally there is evidence of favorite-long-shot bias – favorites tend to be underpriced and long-shots overpriced. The main aim of this paper is to build a model of a prediction market which incorporates some of the characteristics of traders we document and to assess whether the resulting model can reproduce the market characteristics we observe.

The main innovation of our model is to incorporate traders who have diverse attitudes to risk as well as, rather than solely as in WZ's model, diverse beliefs and different levels of wealth. Heterogeneity in the taste for risk is especially likely to be a feature of unconventional markets such as prediction markets, which are not dominated by financial institutions and

³ Snowberg, Wolfers and Zitzewitz, 2013, describe Intrade as the industry standard amongst prediction markets. A lawsuit was filed against it in late November 2012 by the US Commodity Futures Trading Commission forcing it to exclude US traders, after which Intrade suspended all trading.

which are attractive to risk-lovers.⁴ Consequently trade in such markets will likely take place between traders who have not only different beliefs but also quite different – even opposite – attitudes to risk. Yet, although there is a vast array of models which investigate market price formation by traders with heterogeneous beliefs – Harris and Raviv, 1993, being a seminal article – the effect on markets of heterogeneous tastes for risk has been relatively neglected. One contribution of our paper is to provide such a model, which, though targeted specifically at prediction markets, has relevance for other types of market which are likely to attract traders with diverse attitudes to risk and especially traders who seek risk.

Simulations confirm that the presence of risk-lovers can *by itself* account for many of the features we observe at the aggregate level. Risk-lovers will exhibit persistent contrariness and, since risk-lovers are prepared to pay a premium for the utility of the gamble, contrarians will make losses on average. The presence of risk-lovers will also result in a propensity for traders to crowd on the buying or the selling side of the market. Even when traders' beliefs are tightly and symmetrically distributed around the true probability risk-lovers can also be a prime source of favorite-long-shot bias, though we show that price inaccuracy can also occur in mild form even if traders are all risk averse. Conversely, risk-loving traders can ameliorate the effects of bias for or against a contract – driven by partisanship, for example – so that such bias is most likely to cause price inaccuracy if all traders in the market were risk-averse. Finally we establish that both thin trading (which is especially likely well before the contract's event occurs) and a reluctance to short (possibly as a result of lack of financial sophistication) can undermine the usefulness of prediction markets as a forecasting tool.

⁴ In a paper based on a dataset which overlaps with our own, Rothschild and Sethi, 2015, argue that prediction market traders are “predominantly young and male”, a group that one might expect to contain an especially high proportion of risk-lovers.

The paper is in four main sections and a summary. In the next section we describe our data and document the wide variety of trading behavior and market outcomes. We then set out our model and explain how we use it to incorporate more diverse motives and thereby create the virtual market on which we base our simulations. In the final two sections we present the results of our simulations, and show the extent to which they capture the outcomes we observe in our data series and what they imply for the market's accuracy.

2. DATA DESCRIPTION AND TRADER CHARACTERISTICS

2.1 Data

Prediction markets provide platforms on which traders can buy and sell shares relating to the outcomes of particular events, the events being defined as Yes/No propositions. Holders of long positions at maturity receive a fixed pay-out if the event occurs and nothing if it does not; while holders of short positions receive the pay-out if the event does *not* occur and nothing if it does. The pay-out in the Intrade markets was \$10.

Intrade shares could be traded before the maturity date and counter-parties to the trades were the traders themselves. Any trader could be a 'market-maker', posting a bid to buy or an offer to sell. There were no restrictions on shorting. However, shorting was not the exact opposite of going long, because Intrade operated a 100% margining system. To protect against default risk each trader was required to maintain a (non-interest bearing) margin account equal to the trader's maximum eventual loss. So, a trader who was long Q shares at price $\$P$ would hold a margin account worth $\$QP$, while a trader who was short would hold a margin account worth $\$Q(10-P)$. Consequently, although the net cashflows associated with taking a long position

were the exact opposite of those associated with taking a short position, both buyers and sellers had to expend cash when opening a position, and the net expected return-on-investment was different for buyers and sellers.

Before 2011 market-makers paid no fees, while price-takers paid \$0.05 per share traded, reducing to \$0.03 for prices exceeding \$9.49 or less than \$0.51. On settlement the fee was \$0.10 per share for shares that were in-the-money, with no fee for out-of-the-money shares. From 2011 the fee structure changed to a flat desk fee of \$4.99 per month per customer.

Our data consist of trades in 136 contracts relating to the US Presidential elections in 2008 and 2012. The contracts are listed in table 1 and include, for each election, contracts for the overall winner of the Presidency (both by name and by party) and the winner of the 21 states in which there was most trading. We therefore have all trades in the main contracts on these two US elections. These contracts generated 491,475 time-stamped trades, labelled with the numerical IDs of the buyer and the seller. There were 5,337 traders in the 2008 election and 6,334 in 2012, with 1,083 traders operating in both elections.

Each event of the 2008 and 2012 elections – that is, the Presidential vote in each State and the overall Presidential result – was represented by a set of contracts that covered all possible outcomes. A typical set consisted of three contracts; for example ‘Republican candidate to win Florida’, ‘Democratic candidate to win Florida’, and ‘Any other candidate to win Florida’, with very little trading in the ‘Any other candidate’ contracts. We refer to the Republican-to-win/Democrat-to-win contracts as ‘mirror’ contracts.

Since the likelihood of anything other than a Republican or Democratic candidate winning was approximately zero, one of the two mirror contracts within each set was effectively redundant: selling a 'Republican-to-win' contract, for example, was equivalent to buying a 'Democrat-to-win' contract. For some tests reported below we construct 'artificial' contracts by aggregating trades in sets of mirror contracts, reversing the sign of the trades in one of the contracts, and deducting its price from \$10. We define each aggregate contract so that trades in pro-Republican contracts (those which paid off if the Republican candidate won) are reversed, in effect creating a market consisting only of pro-Democrat contracts. The only contracts which we do not aggregate in this way are the 'named' Presidential contracts, such as Barack Obama and John McCain to win the 2008 elections. This is because, until all other candidates dropped out of the race, the two main contenders were not the only ones on whom there were contracts to win.

2.2 Diverse characteristics of traders and market outcomes

Characteristics of traders

An initial, broad, but telling indication of differences in traders' characteristics is provided by the very existence of mirror contracts. With no shorting restrictions there is no obvious reason for both sides of the mirror to exist. In correspondence with the author Intrade explained that many traders apparently did not grasp this: until it operated markets in both Republican and Democratic contracts Intrade was plagued with questions from traders and the media asking why the 'missing' contracts were not available. As a further possible indicator that many traders lacked the financial sophistication to understand the concept of shorting nearly 70% of traders in each election never held short positions. At the opposite extreme we identified approximately 25 traders in each election who appeared to be operating highly sophisticated

algorithmic trading programs, trading at high frequency and identifying arbitrage profits to be made when the prices of sets of contracts did not sum to \$10.⁵

Figure 1 shows the behavior of a trader who traded over 2.5 million shares over the two elections, chalking up losses of nearly \$10 million. The figure shows only the 2012 activity but it was similar in 2008. The trader went heavily long on pro-Republican contracts and also heavily short on pro-Democrat contracts. Whilst this may have been due to extreme confidence in Republican victories, it also suggests another possible motive – an extreme desire to back the Republican Party, maybe even to influence the market. Other traders provide similar, though less enthusiastic, examples of apparent partisanship: depending on the election, some 30% to 40% of traders traded only in pro-Democrat contracts, while 20% to 30% traded only in pro-Republican ones; of these, nearly 90% never held short positions in these contracts.

At the opposite extreme to the trader depicted in figure 1 our data also reveal the existence of traders such as one of the earliest participants in the 2012 market, who limited his or her activity to the purchase of one share. The position was held until maturity 15 months later and generated a profit of \$4.94.⁶ These examples illustrate the great differences in the sums which traders are willing to expose to the risks that are the inevitable consequence of participating in this type of market.⁷ Table 2 shows that there are other major differences in traders' behavior. They trade vastly different numbers of shares: the median number of shares traded per trader was 60 in 2008 and 142 in 2012; the minimum and maximum were 1 and over half a million in 2008, and 1 and over 2 million in 2012. They differ greatly in the

⁵ These opportunities were short-lived and the pricing errors which caused them were small.

⁶ The share was bought from a trader who dealt in more than half a million 2012 Presidential Election shares on Intrade over the course of 22 months, and made a loss of over \$111,000.

⁷ It is presumably considerations such as these that led Nate Silver to describe the figure 1 trader as an 'idiot degenerate gambler'. <http://fivethirtyeight.com/features/intrade-betting-is-suspicious/>.

total trades they carry out: in 2008 the median, minimum and maximum were 6, 1 and 47,549; in 2012 the equivalent figures were 10, 1 and 104,170. They differ in the number of contracts traded: the median number of different contracts traded in 2008 was 1, with a minimum and a maximum of 1 and 58, respectively; the equivalent figures in 2012 were 1, 1 and 60. And they differ in the number of days they hold shares: the mean holding period per share was 16 days for the median trader in 2008 with a minimum and a maximum of, respectively, 0 and over 1,000 days; in 2012 the equivalent figures were 7, 0 and almost 1,500.⁸

All these differences are consistent with traders differing widely in the wealth they are willing to commit to these markets, in their motives and in their attitudes to risk.

One final indicator of heterogeneity amongst traders is provided by figure 2 which shows that the Intrade markets were thin early on in each Presidential campaign, most traders not participating until the last three months or so of each campaign.

Overall, our data suggest that it is implausible that traders in prediction markets differ only in their beliefs about the true probability of the event occurring: differences in financial sophistication, partisanship, attitudes to risk and wealth are additional characteristics which appear likely to affect prediction markets and which should be incorporated into a model of these markets.

Market outcomes

Such a model ought also to account for the market features that our data reveal. One is the skewness of the profits' distribution. Although the figure 1 trader lost by far the most in the two elections, table 2 shows that a minority of traders made relatively large losses,

⁸ The increase in the number of shares traded and decrease in the mean holding periods between 2008 and 2012 are to be expected given the change in the fee structure.

counterbalanced by a majority of traders who made relatively small profits. The mean net profit per trader in the 2012 election market is zero, as would be expected.⁹ The mean net loss per trader – that is, the average transactions costs – in the 2008 election market equals \$45. However, for both elections, the median profit is positive (54% of traders made profits in 2008 and 66% in 2012).

Linked to this is the tendency of the markets to polarisation. Traders often herd ie crowd to one side of the market or the other, rather than both sides of the market being equally represented. Table 3 demonstrates this characteristic, summarising by contract the daily imbalances between buyers and sellers, for days on which at least 10 traders traded in the contract. It is clear that, in both elections, herding and its inevitable opposite, contrariness, were frequent. There was a particular tendency for the crowd to trade on the buying side, which is likely to be partly due to those traders who appeared not to understand the concept of shorting.¹⁰

We also find that the tendency to trade with or against the crowd appears to be a relatively *stable* characteristic of a trader. Appendix 1 describes how we calculate a contrariness rank for each trader. For those traders who traded in both elections we regressed the 2012 contrariness rank on the 2008 rank, resulting in a coefficient of 0.12, with a *t*-statistic of 3.68. Including a trader's 2008 profit/loss per share (total profit or loss, divided by the number of shares traded) as an independent variable in the regression had no major effect on this result, and the profit/loss itself was not a significant explainer.

⁹ The mean net profit was, in fact, very slightly negative, due to the transactions fees charged before 2011. We do not deduct the monthly desk fees as we do not have sufficient information to enable us to apportion fees between trades – a trader may be carrying out other trades on which we do not have data.

¹⁰ However, we repeated the analysis with the artificial aggregated mirror contracts described earlier and, while there was some reduction in the tendency to herd on the buying side, there was little substantial change to the results. We also repeated the analysis excluding all the figure 1 trader's trades, with no appreciable difference.

Our data also suggested that contrariness is associated with loss-making. We found that the median trader in the most contrarian quintile made a loss of \$3 in the 2008 election and \$27 in the 2012 election; whereas the equivalent figures for the median trader in the least contrarian quintile were a profit of \$156 and \$2,092 respectively.¹¹

SWZ, 2013, argue that political prediction markets exhibit favorite-long-shot bias. As evidence they use data from Intrade's 2004 election markets and show that all the contracts that traded above a price of 50 on the night before the election won, and all those that traded below a price of 50 lost "implying that those markets which predicted a high probability of one candidate winning under-predicted the actual probability of victory, while those that predicted a low probability of victory over-predicted victory." (SWZ, 2013, p.667). Our data show a similar pattern.¹² Pooling the 2008 and 2012 contracts, and combining the mirror contracts into single pro-Democratic contracts, we found that in the month before the relevant election the two contracts in the \$9 - \$10 price range both paid off; so too did the five contracts in the \$8 - \$8.99 range, the seven contracts in the \$7 - \$7.99 range, nine of the ten contracts in the \$6 - \$6.99 range, and both the contracts in the \$5 - \$5.99 range. On the other hand, the two contracts in the \$2 - \$2.99 range failed, as did the four contracts in the \$1 - \$1.99 range, and the five contracts in the \$0 - \$0.99 range. Prices in other periods showed a similar pattern, though by the final day almost all the prices were either above \$9, and the contracts paid off, or were below \$1 and they did not. So by that date favorite-long-shot bias was minimal.

¹¹ Differences between mean profits and losses in the two quintiles are even greater but are skewed by the figure 1 trader, who is the most contrarian of all the traders, so we do not report them.

¹² Of course in both SWZ's and our cases the results of the different election contests in each particular year cannot be regarded as independent observations.

Overall, the examination of our data set reveals what appears to be highly diverse trading behavior and some notable outcomes of that trading behavior. In the next two sections we build a model of a prediction market which allows us to incorporate traders with diverse motives and attitudes, and through a set of simulations to examine whether such a market can generate the outcomes observed in our data and what the interplay of these diverse motives and attitudes implies for the forecasting accuracy of prediction markets.

3. A MODEL OF A PREDICTION MARKET

3.1 Model

We assume that all traders have the same general utility function but differ in their wealth and attitude to risk. Formally we assume,

$$U(W_k) = \frac{W_k^{(1-\gamma_k)}}{1-\gamma_k}, \quad (1)$$

where W_k is trader k 's wealth and γ_k is an indicator of the trader's attitude to risk. The trader is risk-averting, risk-loving or risk-neutral depending upon whether γ_k is greater than, less than, or equal to zero.

A trader's wealth at the end of the contract will be $W_{0,k} + x(1 - p)$ if the contract pays off, and $W_{0,k} - xp$ if it does not, where $W_{0,k}$ is trader k 's initial wealth, x is the number of contracts demanded (which will be negative if the trader is a seller) and p is the contract's price. Maximisation of the expected value of $U(W_k)$ generates for a risk-averse trader the following expression for the optimum number of contracts demanded, x_k^* .

$$x_k^* = W_{0,k} \frac{\theta_k^{1/\gamma_k} - 1}{1 + p(\theta_k^{1/\gamma_k} - 1)} \text{ for } \gamma_k > 0, \quad (2)$$

where $\theta_k \equiv \frac{q_k(1-p)}{p(1-q_k)}$; and q_k is the trader's belief of the probability of the event occurring.¹³

Since the sign of x_k^* is determined by whether θ_k exceeds, or is less than, 1, which in turn is determined by whether q_k exceeds or is less than p , their decisions about whether to buy, sell or do nothing are entirely determined by traders' beliefs *not* by their degree of risk aversion; the latter influence only the quantities to be bought or sold.¹⁴

Where traders are not risk averse, i.e. $\gamma_k \leq 0$, equation (2) does not represent their demand function because the second-order conditions for a maximum are not satisfied. With no budget constraint an investor who is risk-neutral or a risk-lover will choose the position which generates the most expected utility and deal in an infinite number of contracts. A natural budget constraint to impose is the opening wealth endowment, $W_{0,k}$, so that maximisation of expected utility occurs at a corner solution: such traders use all their wealth either to buy the contract or to sell it. If they are buyers then they will buy $W_{0,k}/p$ contracts; if they are sellers they will sell $W_{0,k}/(1-p)$.¹⁵ Formally, their demand for the contract will be determined as follows:

¹³ WZ's, 2007, model is a special case of this utility function where $\gamma_k \rightarrow 1$ in which case the utility function becomes logarithmic in wealth, and demand functions are linear in θ_k .

¹⁴ This will be the case for risk-averters with any utility function. A risk-avorter will never, for example, buy a contract for which $q_k < p$ because this would mean that expected return contributes negatively to expected utility, in addition to the negative contribution from exposure to risk.

¹⁵ This constraint is implicitly imposed on risk-averters: in equation (2) the fraction of opening wealth that is used to set up the optimum position, $W_{0,k}/x_k^*p$ for a buyer or $W_{0,k}/x_k^*(1-p)$ for a seller, cannot exceed 1.

$$x_k^* = \begin{cases} \frac{W_{0,k}}{p} & \text{if } \theta_k \left(\frac{p}{1-p} \right)^{\gamma_k} > 1 \\ -\frac{W_{0,k}}{1-p} & \text{if } \theta_k \left(\frac{p}{1-p} \right)^{\gamma_k} < 1 \\ 0 & \text{otherwise} \end{cases} \quad \begin{matrix} (3a) \\ (3b) \\ (3c) \end{matrix}$$

Thus, in contrast to the risk-averse traders, the value of γ_k , as well as the value of q_k relative to p , plays a role in determining whether the risk-loving traders buy the contract or sell it; but the *amount* these traders buy or sell is determined by $W_{0,k}$ and p alone.

Equations (2), (3a), (3b) and (3c) specify the complete set of traders' demands. The equilibrium price is the price at which $\sum_{k=1}^{k=N} x_k^* = 0$, where N is the number of traders. Because we assume traders differ in their attitudes to risk closed-form solutions for the equilibrium price and traders' resultant demands and profits etc. are not feasible. Instead we derive these variables and the characteristics of the resulting market using simulations, each set of which assumes a particular distribution of γ_k s, $W_{0,k}$ s and q_k s and imposes a true probability of the event occurring.

3.2 Simulation details

Many of the market outcomes revealed in section 2.2 can be trivially modelled by appropriate assumptions about traders' beliefs. But in the absence of any compelling evidence to the contrary we follow WZ, 2007, in treating trader's beliefs as symmetrically and tightly distributed. Furthermore we shall assume that the mean of traders' beliefs equals the true probability of the event. This allows us to isolate the effects of the factors on which we are focusing such as attitudes to risk and bias. Specifically, in all our simulations we assume that

traders' beliefs are uniformly distributed with a range of ± 0.09 around whatever is the assumed true probability of the event. These true probabilities are allowed to vary between 0.1 and 0.9 in steps of 0.1. We also make the common assumption that wealth and beliefs are uncorrelated.

The baseline set of simulations assumes a market populated by 1,000 traders with γ_k s within the range ± 1 around a mean, $\bar{\gamma}$, which is constant for any particular set of simulations but in different sets takes on values from -2 to 2 in steps of 0.5.¹⁶ So, for example, in the first set the γ_k s range from -3 to -1 and all traders are risk-lovers; in the second set the γ_k s range from -2.5 to -0.5; and so on. We are therefore simulating the effects both of diverse attitudes to risk and of different proportions of risk-lovers in the market. In these baseline simulations traders' wealth is uniformly distributed in the range 0 to 1,000, and wealth is not correlated with the γ_k s. Furthermore traders are assumed to have no inherent bias towards buying or selling the contract. We summarise the baseline results in section 4, discussing the results in the context of other, related literature.

In subsequent sets of simulations, results of which are summarised in section 5, we examine the effect on market prices of the following variations:

- We allow traders to have an innate preference either for buying or for selling a contract. One possible reason for such a preference is political partisanship: traders wish to show their support for a political party by buying contracts that pay off if it wins, or they might wish to affect the price of a contract in the hope that this will build momentum behind the party. They might also wish to diversify their disappointment if

¹⁶ We also carried out simulations for different ranges of γ_k around $\bar{\gamma}$ but these simulations produced very similar results.

their favored party loses. To allow for such a preference we add a constant to each trader's demand equal to a fraction of the trader's wealth. The fractions themselves are assumed to be uniformly distributed around the range 0 to 1, in the case of strong pro-contract bias, and 0 to 0.1 in the case of weak pro-contract bias; similarly strong and weak anti-contract biases are represented by fractions with negative signs.

- We assume traders' wealth is log normally distributed so as to generate a mean of approximately 500 and a Gini coefficient of approximately 0.8.¹⁷
- We mimic a thin market by severely restricting the number of traders.
- We capture the effects of the possibly self-imposed restriction on shorting noted in section 2.2 by running simulations where 67% of traders – the percentage of all traders in our Intrade markets who never shorted – are assumed to choose not to trade rather than take the short side of the contract. This tendency was randomly assigned and was assumed to be uncorrelated with beliefs, degree of risk aversion, or wealth.

4. SIMULATION RESULTS: BASELINE SIMULATIONS

In this section and the next we report on the model's ability to account for the main aggregate features of the prediction markets revealed in section 2.2 – herding, persistent contrarianism, the association of contrarianism with loss-making and the skewness in profits – and the likely usefulness of prediction markets as a forecasting tool.

¹⁷ We also simulated a market where traders' wealth and their γ_k s were either positively or negatively correlated (see Paravasini *et al.*, 2016, for example), but this made little difference to the results so we do not report them.

4.1 Herding and consistent contrariness

There is an extensive literature on herding in conventional stock markets, Lakonishok *et al.*, 1992, being one of the earliest examples. In a recent study, Wei *et al.*, 2015, investigate mutual funds which consistently act in a contrarian manner, regularly trading against the herd or crowd. They summarise various reasons for such behavior, such as possession of private information, superior information-processing abilities, (over)confidence, and tournament-induced risk-taking and find evidence that contrarian mutual funds tend to be more profitable, concluding that their contrarian positions are driven by superior information. Of course, for the information-related arguments to account for *persistent* contrariness it is not enough for contrarians persistently to possess superior information; their information must also *persistently* result in disagreement with the crowd's assessment of the asset's value so that the contrarian buys when most others are selling, or sells when most others are buying.

In our baseline simulations most of the factors which might produce contrariness are ruled out: no trader or group is assumed to possess consistently superior information, and certainly no traders are assumed to possess such information which also happens to lead them persistently to trade against the majority; no group processes information more efficiently than any other; no group is more confident than another, and there is no tournament-motivated behavior. Indeed, most of these factors are unlikely to be present in our sampled Intrade markets. The only way in which the baseline model differs from WZ's, 2007, model is that it allows for the presence of traders with diverse attitudes to risk, and specifically it allows for the presence of risk-loving traders. Yet as tables 4 and 5 show this in itself is sufficient to ensure market polarisation and persistent contrariness.

Table 4 illustrates the tendency of the simulated market to polarise, that is, for the majority of traders to crowd on one side of the market and a minority to trade in a contrarian fashion against the crowd. It shows market outcomes derived from, in each case, 100 simulations of the baseline virtual market.

Each cell relates to a particular assumed average value of γ_k , $\bar{\gamma}$, and a particular assumed true probability, π . As one moves from the extreme left column of the table to the extreme right the percentage of risk-lovers in the market decreases from 100% to zero; when $\bar{\gamma} = 0$ the percentage is approximately 50%. As one moves from the top of each column to the bottom the true probability of the contract paying off increases from 0.1 to 0.9. Each cell shows the average ratio (over the 100 simulations for that particular combination of $\bar{\gamma}$ and π) of sellers to buyers. Departures from 1 indicate polarisation; a figure above 1 denotes an excess of sellers; below 1, an excess of buyers.

The columns of table 4 to the right of $\bar{\gamma} = 0$ relate to markets that are dominated by risk-averse traders: the further to the right the greater this dominance. As noted earlier, for such traders the key influence on their decisions whether to buy or to sell is their beliefs. These beliefs are by assumption symmetrically distributed. Consequently, when the market is entirely composed of such traders roughly half of them will choose to be sellers and half to be buyers. The polarisation measure therefore becomes close to 1 in these right-hand columns.

As Manski, 2006, has shown, when traders are risk neutral – which corresponds most closely in our case to $\bar{\gamma} = 0$ – the ratio of buyers to sellers must be $p^*/(1-p^*)$, where p^* is the equilibrium price. If p^* equals 0.33, for example, buyers will want to buy twice as many contracts as sellers will want to sell, $W_{0,k}/0.33$ against $W_{0,k}/(1-0.33)$; consequently for

equilibrium there needs to be twice as many sellers as buyers. The polarisation measure will in this case be 2. On the other hand, at a p^* of 0.67 there needs to be twice as many buyers as sellers, giving a polarisation measure of 0.5. In general, absolute polarisation increases as the probability of the event – and hence, as we show shortly, its price – becomes more extreme.

Columns to the left of $\bar{\gamma} = 0$ in table 4 represent a market that is becoming increasingly dominated by risk-loving traders. Equations (3a) and (3b) demonstrate that as their degree of risk-loving increases these traders become increasingly willing, given their beliefs, to buy low-probability contracts and sell high-probability contracts. Consequently, the columns further to the left of $\bar{\gamma} = 0$ in table 4 involve an increasing number of risk-loving buyers of low-probability contracts, reducing the polarisation measure in the upper rows. Similarly, the increasing number of risk-loving sellers of high-probability contracts increases the polarisation measure in the lower rows.

Table 5 shows that risk-lovers will tend to exhibit *persistent* contrariness. It is worth emphasising that each of these simulations randomly allocates each trader a different view of the true probability, but for each trader the rate of risk aversion remains fixed across all simulations. Any persistence is therefore due to the trader's taste for risk and not connected with the trader's beliefs. The table shows the average percentage of times over all the simulations that traders, sorted into quintiles by their γ_k s, acted as contrarians. The results presented show, therefore, that risk-lovers are the ones who will tend consistently to take contrarian positions, and the more risk-loving they are the more frequently they will trade against the crowd. If there are no risk-lovers in the market then, as the simulation results for $\bar{\gamma} = 2$ in table 5 suggest, all traders will typically be contrarians roughly 50% of the time and trade with the crowd the other 50%.

In summary, the results in tables 4 and 5 show that market polarisation, and the persistent contrariness of particular traders can be explained by a single factor – the presence of risk-loving traders.

Table 6 shows the relationship between risk aversion and profits implied by our baseline model. It shows that our simulated market suggests that in the presence of a significant proportion of risk-loving traders profits will be highly skewed. Furthermore it implies that contrarians are especially likely to make losses because contrarians are the risk-lovers, who are prepared to pay a premium for the utility of the gamble.

The existence of a group of traders – risk-lovers – which persistently makes losses provides a reason for rational risk-averse traders to participate in prediction markets, since they can exploit this tendency. As the right-most block of table 6 confirms, without these traders the expected profit for all traders would be zero and therefore it would be irrational to participate in the market. But the fact that risk-lovers tend to be persistent loss-makers suggests that over time they will comprise a declining proportion of the market, in which case we would expect to see, as prediction markets develop, the characteristics of market outcomes that are due to risk-loving traders gradually disappearing. And consequently so too would the rationale for risk-averse traders to participate. The sustained vibrancy of prediction markets may therefore rely upon the pool of risk-loving traders being constantly replenished.

4.2 Price accuracy and favorite-long-shot bias

Favorite-long-shot bias is apparently a pervasive phenomenon. It has been noted in horse-racing and football betting markets (see Snowberg and Wolfers, 2008, for a review), in fixed-odds US college basketball and football markets (see Berkowitz *et al.*, 2016), and, in a similar context to ours, 50 US Presidential Electoral College races and 24 Senatorial races in 2008 (Rothschild, 2009). Ottaviani and Sørensen, 2008, discuss a number of reasons proposed for its existence.¹⁸ Some of these, such as market power by informed bettors or betting by insiders, will not normally apply to Presidential election markets where relevant information is generally available to all potential traders. Mis-estimation of probabilities is one possible explanation. This invokes two ideas that underlie the Kahneman and Tversky, 1979, Prospect Theory: that people have a strong preference for certain, over extremely likely, outcomes and an inability to distinguish between small and tiny probabilities. In a study of ‘exotic’ US racetrack bets – such as which horse will come first and which will come second – Snowberg and Wolfers, 2010, find that the mis-estimation explanation outperforms the alternative, that favorite-long-shot bias is due to the presence of risk-loving bettors.

The baseline model rules out mis-estimation of the form that generates favorite-long-shot bias but allows for the inclusion of risk-loving traders. Table 7 shows the mean equilibrium prices derived from the 100 baseline simulations, with each cell giving the mean price for a particular assumed average value of γ_k , and a particular assumed true probability.¹⁹ Of course, subject to slight variation across simulations, the prices are symmetrical around a true probability of 0.5. For example, the mean price when the probability is 0.1 equals 1 minus the price when it is 0.9. Figure 3 presents a subset of these results, together with the 45° line.

¹⁸ The article relates to pari-mutuel betting markets but most arguments can also be applied to prediction markets of the type we analyse here.

¹⁹ Results for the median price were very similar to those presented here; and the minimum and maximum prices over the 100 simulations in each case were within 1 percentage point of one another.

When all or nearly all traders are risk averse ($\bar{\gamma} > 0.5$) these results are consistent with the simplest version of the WZ, 2007, framework: the equilibrium price closely approximates the mean of traders' beliefs, which is the true probability given our assumptions. However, as the market becomes increasingly dominated by risk-loving traders, the equilibrium price diverges much more noticeably from the true probability; and the more risk-loving the traders are assumed to be, the greater is the inaccuracy. Furthermore, as figure 3 highlights, this inaccuracy takes a specific form: as risk-lovers come to dominate the market, equilibrium prices tend to overstate the probability of the occurrence of an event with a low true probability and to understate the probability of an event which has a high true probability, i.e. equilibrium prices exhibit favorite-long-shot bias.

In our model, therefore, risk-loving causes favorite-long-shot bias and is the primary source of it. Intuitively, in the case of low-probability contracts, even if the price equals their assessment of the event's probability risk-lovers may still want to buy the contracts because of the outside chance of a positive gain if the contracts pay off. Their demands will drive up the price and encourage risk-aversers (or those who love risk less) to sell the contracts. An equivalent argument applies to risk-lovers' willingness to sell high-probability contracts. Thus low-probability contracts will tend to be over-priced and high-probability contracts under-priced.

In fact favorite-long-shot bias can still occur even where there are no risk-lovers – though the effects are second-order – because of the non-linearity implied by equation (2): when traders' γ_k s lie between 0 and 1 then for low-probability contracts there will be moderate excess demand for the contract even if the price equals the mean of beliefs; for high-probability contracts there will be moderate excess supply. By contrast, where traders' γ_k s all exceed 1

the non-linearity works in the opposite direction, and, as the last column of table 7 indicates, favorites can be slightly over-priced and long-shots under-priced.²⁰ Thus mild favorite-long-shot bias and its opposite are both consistent with a market comprised wholly of risk-averting traders.

5 SIMULATION RESULTS: VARIATIONS ON THE BASELINE

The conclusions of section 4.1, that polarisation is induced by the presence of risk-loving traders, and is particularly acute when there are both risk averters and risk lovers in the market, are essentially unaltered by the variations on the baseline conditions. In this section we therefore focus on the price effect of these variations.

5.1 Pro- and anti-contract bias

We report only the simulated effects of strong pro- and anti-contract biases since the weak biases made little difference to our earlier results. Table 8 shows the effects of allowing traders to have an average strong bias for (or against) buying the contract. In Panel A we restrict the results to cases where there is a sufficiently high $\bar{\gamma}$ to ensure that the market consists solely of risk-averse traders; in Panel B all traders are risk-lovers. In each Panel the central block is the same as the equivalent columns shown in table 7, the baseline results. The results in Panel A, where all traders are risk averse, are shown pictorially in figure 4. We do not show the Panel B results pictorially, as they are almost the same as the table 7 baseline results, with risk-loving traders and no bias.

²⁰ A formal proof of these results is presented in appendix 2.

The results shown in Panel A of table 8 and figure 4 indicate, unsurprisingly, that equilibrium prices tend to be lower when there is anti-contract bias and higher when the bias is pro-contract. Less obviously, Panel B shows that the presence of risk-lovers in these markets markedly dampens the effect of bias, although the underlying favorite-long-shot bias remains. The reason for this difference is that risk-loving markets are likely to be far more responsive to a change in price. Faced with an increased bias against the contract there will be a fall in the contract's price and this, as conditions (3a) and (3b) indicate, will 'flip' some risk-loving traders from using all their wealth to sell contracts to using all their wealth to buy them. Consequently, only a minor price fall is required to re-establish equilibrium. A similar argument applies to the effects of an increased bias in favor of the contract: a small rise in price is all that is required to re-establish equilibrium. For risk-averting markets the response of demand to price changes is much more tempered and hence a much larger price rise or fall is required to establish an equilibrium in the event of increased bias for or against the contract. For such a market the bias against or in favor of buying the contract can render the equilibrium price inaccurate by more than 20 percentage points in our model.

One other feature of table 8 Panel B is that, in a market with risk-lovers, inherent pro- or anti-contract biases can, for certain contracts, cause prices to be *more* rather than less accurate estimates of true probabilities. Specifically, anti-contract bias counters the inaccuracy caused by the willingness of risk-lovers to buy low probability contracts; whilst pro-contract bias counters the inaccuracy caused by the willingness of risk-lovers to sell high-probability contracts. In these cases the pro- or anti-contract bias counters the favorite-long-shot bias. In the other possible combinations – pro-contract bias and high probability contracts, and anti-contract bias and low probability contracts – the two biases reinforce one another and cause prices to be even less accurate estimates of the true underlying probabilities.

5.2 Log-normally distributed wealth, and thin markets

All our results so far are based on wealth being uniformly distributed. We also simulated a market in which our other previous assumptions were maintained but wealth was assumed to be log-normally distributed so as to generate a mean of around 500 and a Gini coefficient of around 0.8. The results obtained were generally too similar to those already reported to be worth recording separately. The one set of simulations for which the assumption of a log-normal distribution made a difference was the set in which we combined this assumption with a thin market.

For the thin market simulations we reduced the number of traders from 1,000 to 10. This is a drastic reduction but in the early days of a contract such a small number of traders was not, in our data set, untypical. The reduction in the number of traders had little effect on mean and median equilibrium prices but did affect the range of prices achieved in the simulations. Figure 5 illustrates the minimum and maximum prices across simulations when wealth is log-normally distributed and there are only 10 traders. These differ according to whether the market is dominated by risk-lovers or risk-aversers. In a market of risk-aversers the gap between the maximum and minimum prices obtained in the simulations is more or less constant at around 15 percentage points regardless of the true probability of the contract. But in a market of risk-lovers the gap is noticeably wider at the extremities, where the possible combination of extreme wealth levels with extreme risk-loving can severely distort the

market. In our model, at these extremes the thinness of the market can lead to equilibrium prices diverging from the true probabilities by almost 30 percentage points.²¹

These simulations therefore suggest that, even where the population of potential traders holds beliefs that are quite tightly centred on the truth, prediction markets can be seriously inaccurate when the number of traders participating is small. And since we have reported that thin markets are a feature of our data until a few months before the contract's event occurs prediction markets may be of rather limited use for economic forecasting.

5.3 Limited shorting

Table 9 presents the effect on prices when 67% of traders are assumed never to short. The results show that this restriction acts very much like the strong bias for buying the contract: it generally raises the price of contracts, by up to four or five percentage points. One way in which aversion to shorting can have a more dramatic effect is of course when there are thin markets and hence the possibility that there are simply no traders willing to take the short end of the market. It is quite possible that such circumstances occurred in the early life of many of the contracts in our data set and is the reason why, on many days, in those early periods no trades occurred. In slightly less extreme circumstances the absence of traders willing to sell a contract can lead to a price being sharply higher than it would otherwise be. Simulations which assumed thin markets and a 67% aversion to shorting but discarded simulations where no market could be formed produced prices as high as 0.4 for contracts with a true probability of 0.1 and prices as high as 0.98 for contracts with a true probability of 0.9. Maximum prices were even higher where a strong average pro-contract bias was allowed.

²¹ We also carried out simulations of thin markets with a uniform wealth distribution. There was a similar, but less pronounced, effect on the range of prices found in the simulations: differences between the maximum and minimum prices in a risk-loving market were around 10 percentage points at the extremes.

6. SUMMARY AND CONCLUSIONS

Examination of a detailed set of prediction markets relating to the 2008 and 2012 US Presidential Elections reveals highly diverse trading patterns, patterns that are difficult to reconcile with traders being homogeneous in all but their beliefs and wealth. Rather it suggests that trade in such markets takes place between traders with very different attitudes to risk, levels of financial sophistication, motives and participation rates. Incorporating such diversity into simulations of a prediction market produces many of the outcomes observable in the actual data, such as polarised market, persistent contrariness, a skewed profits distribution and favorite-long-shot bias. Indeed, the incorporation of risk-loving traders into the market is *all* that is needed to produce these effects.

The simulations reveal price distortions which have implications for the usefulness of prediction markets for forecasters. Of course there are other reasons why prediction markets might be inaccurate: traders' beliefs may be far from the true probability; and even if they are not, the true probability may change as circumstances change. But our simulations suggest that even if the true probability of an event were to remain constant over time and if the market were, on average, to have a correct assessment of that probability, heterogeneities in traders' motives and attitudes can lead to prediction market prices being inaccurate. In particular we find that despite traders' beliefs being tightly and symmetrically distributed around the true probability, prediction markets containing risk-loving traders will exhibit marked favorite-long-shot bias, and that inherent biases against a contract, an aversion to shorting and thin markets also produce price inaccuracies, some of which are exacerbated, and some tempered, by risk-loving traders.

APPENDIX 1: CONSTRUCTING A CONTRARINESS INDEX

The contrariness measure represents the extent to which an individual consistently trades in the opposite direction to the majority. We assume that an information signal arrives at the start of each trading period and induces trade to take place. At every trading period each trader could belong in the crowd (buying when most traders buy, or selling when most traders sell), or in the group of contrarians which trades against the crowd. Our assumed trading period is a day.

For each trader, k , we construct a net contrariness index, NCI_k , which is increasing in the tendency of the trader to trade against the crowd. The index is made up of two elements. The first is a measure of the number of potential counter-parties for the trader and equals

$\frac{N_{Crowd,i,t}}{N_{Contrarians,i,t}}$. $N_{Crowd,i,t}$ is the number of traders in the i -crowd in period t (that is the number of traders who buy contract i when most traders are buying, or sell when most are selling); and $N_{Contrarians,i,t}$ is the number of i -contrarians in period t . This is a measure of the ‘polarisation’ of the market for contract i during trading period t . The higher it is, the more a contrarian trader is trading against the crowd, and the more ‘conventional’ is a member of the crowd.

The second element in the index is $\frac{NetV_{i,k,t}}{NetV_{i,t}}$, where $NetV_{i,k,t}$ is the net number of shares of contract i bought or sold by trader k in period t ;²² and $NetV_{i,t}$ is the net number of shares of contract i traded in period t (which, of course, must equal the net number of shares bought or sold by contrarians, and sold or bought by the crowd). In the absence of information about unmatched bids and offers, the number of shares traded relative to market volume is a proxy for the strength of the trader’s views or taste for risk.

²² Traders sometimes buy and sell shares within one time period.

Contrarian trader k is assigned contrariness score $Con_{i,k,t}$ for any trading period t in which she has traded in contract i :

$$Con_{i,k,t} = \frac{N_{Crowd,i,t}}{N_{Contrarians,i,t}} \times \frac{NetV_{i,k,t}}{NetV_{i,t}}, \quad (A1.1)$$

where trader k is considered a contrarian if she buys contract i during period t when most traders are selling, or sells when most are buying. If there are equal numbers of sellers and buyers in trading period t , trades in that period are excluded from traders' contrariness scores. A trader who trades with the crowd is assigned the same index but this is preceded by a minus sign in the final aggregation (equation (A1.2)).

The net contrariness index for trader k in contract i , $NCI_{i,k}$, is a weighted average of the individual trading periods' scores for trader k in contract i , the weighting being the market-wide volume of activity in each trading period:

$$NCI_{i,k} = \frac{\sum_t I_{i,k,t} V_{i,t} Con_{i,k,t}}{\sum_t V_{i,t}}, \quad (A1.2)$$

where $I_{i,k,t}$ is an indicator which equals 1 if trader k is contrarian in share i during period t , and -1 if she is in the crowd; and $V_{i,t}$ is the total (gross) number of shares traded in contract i in period t . Each trader k is allocated with a separate index for each contract in which she trades.

Finally, for each trader k we aggregate contract-specific $NCI_{i,k}$ scores into one overall net contrariness index, NCI_k , which weights the contract-specific scores by the trader's trading volume in each contract, as shown in equation (A1.3).

$$NCI_k = \frac{\sum_{\{i_k\}} (NCI_{i,k} V_{i,k})}{\sum_{\{i_k\}} V_{i,k}}, \quad (A1.3)$$

where $\{i_k\}$ denotes the set of contracts in which trader k has traded, and $V_{i,k}$ is the total (gross) number of shares traded in contract i by trader k . For testing the stability of contrariness of traders who traded in both elections, we rank each NCI_k within the distribution of NCI_k s and adjust each set of ranks by dividing them by the number of ranked observations, so that they run from 0 to 1 in each election.

APPENDIX 2: RISK AVERTERS AND FAVORITE-LONG-SHOT BIAS

Assume the market consists of risk-averse traders whose beliefs about the probability of the event occurring are distributed equally and symmetrically around the true probability. For simplicity assume all traders have the same utility function, the same risk aversion parameter, γ , and the same wealth. The expected-utility maximising quantity of contracts held is $x^*(q_k)$, where q_k is the trader k 's belief about the true probability.

Consider the case where traders' beliefs straddle p^0 : for one group of traders, A, $q_A = p^0 + \varepsilon$ and for the other, equally-sized group B, $q_B = p^0 - \varepsilon$. Since, by assumption, beliefs are distributed equally and symmetrically around the true probability, p^0 must equal that probability.

Consider aggregate demand if the market price were equal to p^0 . Dropping the cumbersome *, a Taylor expansion gives:

$$x(q_A) = x(p^0 + \varepsilon) = x(p^0) + \varepsilon x'(p^0) + 0.5\varepsilon^2 x''(p^0) + \text{terms of order } \varepsilon^3 \text{ or less.} \quad (\text{A2.1})$$

$$\therefore x(q_A) + x(q_B) \approx 2x(p^0) + \varepsilon^2 x''(p^0). \quad (\text{A2.2})$$

However, $x(p^0)$ must equal 0, since a risk-averse trader whose beliefs equal the market price will not wish to trade.

$$\text{Hence } x(q_A) + x(q_B) \approx \varepsilon^2 x''(p^0). \quad (\text{A2.3})$$

Consequently, when $x''(p^0) > 0$ there will be excess demand and the equilibrium price will be pushed above the true probability. When $x''(p^0) < 0$ there will be excess supply and equilibrium price must be below the true probability. And, for reasons of symmetry,

whatever effect is observed when $p^0 < 0.5$ will be the opposite of the effect observed when $p^0 > 0.5$.

For the utility function in this paper, $x_k = W_{0,k} \frac{\theta_k^{1/\gamma} - 1}{1 + p(\theta_k^{1/\gamma} - 1)}$ where $\theta_k \equiv \frac{q_k(1-p)}{p(1-q_k)}$,

$$x''(p^0) = \frac{(1-\gamma)(1-2p^0)}{\gamma^2(p^0)^2(1-p^0)^2}.$$

Therefore when $p^0 < 0.5$, $x(p^0 + \varepsilon) + x(p^0 - \varepsilon)$ will be positive if $0 < \gamma < 1$. There will be excess demand and the price will be pushed above the true probability, creating favorite-long-shot bias (FLSB). Conversely, $x(p^0 + \varepsilon) + x(p^0 - \varepsilon)$ will be negative if $\gamma > 1$, pushing the price down: the opposite of FLSB.

Similar arguments apply when $p^0 > 0.5$. For example, $x(p^0 + \varepsilon) + x(p^0 - \varepsilon)$ will be negative if $0 < \gamma < 1$. There will be excess supply and the price will be pushed below the true probability, again creating favorite-long-shot bias (FLSB).

In summary, for the utility function in our paper:

If $0 < \gamma < 1$ then equidistant beliefs around the truth create favorite-long-shot bias (FLSB).

If $\gamma > 1$ then equidistant beliefs around the truth create the opposite of FLSB.

In both cases the FLSB and its opposite will be second-order effects.

REFERENCES

- Berg, J., Forsythe, R., Nelson, F. and T. Rietz, 2008, "Results from a Dozen Years of Election Futures Markets Research," in *The Handbook of Experimental Economic Results*, Vol. 1, C. Plott and V.L. Smith, eds., Elsevier Science B.V., Amsterdam, The Netherlands.
- Berkowitz, J.P., Depken, C.A. and J.M. Gandar, 2016, "A Favorite-Longshot Bias in Fixed-Odds Betting Markets: Evidence from College Basketball and College Football", *Quarterly Review of Economics and Finance*, forthcoming (doi:10.1016/j.qref.2015.11.011)
- Borghesi, R., 2006, "Underreaction to New Information: Evidence From an Online Exchange", University of Texas mimeo
- Brown, A., 2014, "Information Processing Constraints and Asset Mispricing", *Economic Journal*, 124, 245-268
- Brown, A. and F. Yang, 2016, "Limited Cognition and Clustered Asset Prices: Evidence from Betting Markets", *Journal of Financial Markets*, forthcoming, (doi.org/10.1016/j.finmar.2015.10.003)
- Harris, M. and A. Raviv, 1993, "Differences of Opinion Make a Horse Race", *Review of Financial Studies*, 6(3), 473-506.
- Kahneman, D. and A. Tversky, 1979, "Prospect Theory: An Analysis of Decision Under Risk", *Econometrica*, 47(2), 263-291.
- Lakonishok, J., Shleifer, A., and R.W. Vishny, 1992, "The Impact of Institutional Trading on Stock Prices", *Journal of Financial Economics*, 32, 23-43.
- Leigh, A. and J. Wolfers, 2006, "Competing Approaches to Forecasting Elections: Economic Models, Opinion Polling and Prediction Markets", IZA Discussion Papers No. 1972, <http://hdl.handle.net/10419/33137>.
- Manski, C.F., 2006, "Interpreting the Predictions of Prediction Markets", *Economics Letters*, 91, 425-429
- Ottaviani, M. and P.N. Sørensen, 2008, "The Favorite-Longshot Bias: An Overview of the Main Explanations", In Hausch, D.B. and W. Ziemba, eds, *Handbook of Sports and Lotteries Markets* (Elsevier BV), Ch 6, 83-101.
- Paravasini, D., Rappoport, V., and E. Ravina, 2016, "Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios", *Management Science*, forthcoming. <http://dx.doi.org/10.1287/mnsc.2015.2317>
- Rothschild, D., 2009, "Forecasting Elections: Comparing Prediction Markets, Polls, and their Biases", *Public Opinion Quarterly*, 73(5), 895-916.
- Rothschild, D., and R. Sethi, 2016, "Trading Strategies and Market Microstructure: Evidence from a Prediction Market." *The Journal of Prediction Markets*, 10(1), 1-29.

Snowberg, E. and J. Wolfers, 2008, “Examining Explanations of a Market Anomaly: Preferences or Perceptions?”, in Hausch, D.B. and W. Ziemba, eds, *Handbook of Sports and Lotteries Markets* (Elsevier BV), Ch 7, 103-136

Snowberg, E. and J. Wolfers, 2010, “Explaining the Favorite-Long Shot Bias: Is it Risk-Love or Misperceptions?”, *Journal of Political Economy*, 118(4), 723-746

Snowberg, E., J. Wolfers and E. Zitzewitz, 2013, “Prediction Markets for Economic Forecasting”, in Elliott, G. and A. Timmerman, eds, *Handbook of Economic Forecasting* (Elsevier BV), Vol 2A, Ch 11, 657-687

Tetlock, P., 2004, “How Efficient Are Information Markets? Evidence from an Online Exchange” University of Texas mimeo.

Wei, K.D., Wermers, R. and T. Yao, 2015, “Uncommon Value: The Characteristics and Investment Performance of Contrarian Funds”, *Management Science*, 61(10), 2394-2414.

Wolfers, J. and Zitzewitz, E., 2007, “Interpreting Prediction Market Prices as Probabilities”, NBER Working Paper No. 12200 dated 2006; latest draft dated January, 2007 available at <http://users.nber.org/~jwolfers/research.php#predictionmarkets>

Zitzewitz, E., 2006, “Price Discovery Among The Punters: Using New Financial Betting Markets to Predict Intraday Volatility”, University of Stanford mimeo.

TABLES

Table 1 Contracts

		No. of shares traded	
		2008	2012
To win Presidential election			
By party:	Democratic Party candidate	156,109	138,751
	Republican Party candidate	166,917	113,398
By name:	Barack Obama (2008, 2012)	1,244,892	4,150,159
	Mitt Romney (Republican candidate 2012)		3,437,100
	John McCain (Republican candidate 2008)	1,489,323	
‘Other’ to win:		161,310	112,650
States (Note 1)	Arizona	13,611	4,956
	Colorado	9,026	42,726
	Connecticut	1,577	622
	Florida	35,937	188,508
	Indiana	63,030	2,963
	Iowa	10,022	17,083
	Michigan	4,991	15,865
	Minnesota	7,129	6,118
	Missouri	42,419	6,964
	Montana	14,733	1,213
	Nevada	9,464	14,259
	New Hampshire	8,075	13,920
	New Mexico	8,267	4,032
	North Carolina	42,439	73,908
	Ohio	28,185	53,127
	Oregon	4,440	1,396
	Pennsylvania	19,963	31,413
	South Carolina	1,987	797
	Texas	1,752	414
	Virginia	36,628	69,410
	Wisconsin	7,733	34,227

Notes

1. For each State there are three contracts: Republican candidate to win; Democratic candidate to win; other to win.

Table 2 Trading statistics per trader

	Mean	Std dev	Min	25% quartile	Median	75% quartile	Max
2008 election							
Net profit, after transaction costs (\$) (Note 1)	-45	32,143	-2,269,598	-45	2	84	297,382
Total shares traded (Note 2)	1,307	11,783	1	17	60	272	548,224
Total trades	65	904	1	2	6	21	47,549
Number of different contracts traded	3	4	1	1	1	3	58
Maximum investment (\$) (Note 3)	3,126	40,345	0	42	142	498	2,269,598
Mean daily investment (\$) (Note 3)	1,268	15,311	0	21	84	257	941,414
Mean holding period (days)	52	108	0	4	16	49	1,381
2012 election							
Net profit, after transaction costs (\$) (Note 1)	0	90,853	-7,122,061	-42	40	569	867,059
Total shares traded (Note 2)	2,695	39,583	1	29	142	746	2,122,441
Total trades	100	1,576	1	3	10	35	104,170
Number of different contracts traded	3	4	1	1	1	2	60
Maximum investment (\$) (Notes 3, 4)	6,028	91,915	0	93	482	2,234	6,690,572
Mean daily investment (\$) (Note 3)	2,354	9,825	0	50	275	1,399	403,216
Mean holding period (days)	59	134	0	1	7	49	1,464

Notes

1. We account for trading fees charged on a per-share basis but not for the fixed monthly desk fees charged from the start of 2011. Trading fees for trades related to the 2012 election are therefore minimal.
2. Accounting for round trips as separate trades. That is, a share which is held to expiry counts as one share traded, while one which is bought and sold during the period counts as two.
3. Intrade operated a 100% margining system to protect against default risk, each trader maintaining a (non-interest bearing) margin account equal to the trader's maximum eventual loss. So, after a one-off trade of Q shares at price \$P the buyer of the shares deposited \$QP with Intrade, while the seller deposited \$Q(10-P). 'Investment' is the amount held in the margin account.
4. The maximum investment in the 2012 election is lower than the biggest loss, as some losses were made in trades executed on the expiry day and we measure investment as the margin account balance at the end of the previous day.

Table 3 Polarisation of Intrade markets

	2008			2012		
Number of contract-days on which there was no imbalance	140			175		
Number of contract-days on which there was an imbalance	1,359, of which ... 803 had excess buyers, 556 excess sellers			1,906, of which ... 1,241 had excess buyers, 665 excess sellers		
	2008			2012		
	Excess	Buyer excess	Seller excess	Excess	Buyer excess	Seller excess
Mean imbalance	8	9	6	9	11	5
Minimum imbalance	1	1	1	1	1	1
Median imbalance	4	5	4	4	5	3
Maximum imbalance	133	133	83	447	447	92

Notes

1. The table summarises the daily imbalances between buyers and sellers for each contract, on days on which at least 10 traders traded in the contract.
2. The lower portion of the table relates only to contract-days where there was an imbalance. 'Excess' is the absolute difference between numbers of buyers and numbers of sellers.
3. The 'Buyer (Seller) excess' columns relate to days on which the number of buyers (sellers) exceeded the number of sellers (buyers).

Table 4 Polarisation of simulated markets

$\pi \backslash \bar{\gamma}$	-2	-1.5	-1	-0.5	0	0.5	1	1.5	2
0.1	1.94	2.13	2.45	3.17	4.56	3.76	1.56	1.00	0.89
0.2	1.56	1.63	1.86	2.23	3.55	2.57	1.27	1.04	0.96
0.3	1.33	1.38	1.44	1.86	2.13	1.86	1.17	1.04	0.96
0.4	1.13	1.17	1.22	1.33	1.38	1.38	1.08	1.00	1.00
0.5	1.00	1.00	1.00	1.00	0.96	1.00	1.00	1.04	1.00
0.6	0.89	0.85	0.82	0.75	0.67	0.75	0.92	1.04	1.00
0.7	0.75	0.72	0.67	0.56	0.47	0.54	0.85	1.04	1.04
0.8	0.64	0.61	0.54	0.45	0.28	0.39	0.75	1.04	1.04
0.9	0.52	0.47	0.41	0.33	0.22	0.25	0.64	1.04	1.13

Notes

1. Each cell in the table shows the tendency of markets to polarise. It shows the average ratio (over all simulations) of sellers to buyers.
2. π , shown at the beginning of each row, is the assumed true probability.
3. $\bar{\gamma}$, at the top of each column, is the assumed mean value of the distribution of traders' risk aversion parameters; the assumed range of the γ_k s is $\bar{\gamma} \pm 1$.

Table 5 Contrariness and risk aversion

$\pi \backslash \bar{\gamma}_j$	$\bar{\gamma} = -2$					$\bar{\gamma} = 0$					$\bar{\gamma} = 2$				
	-2.78	-2.40	-2.01	-1.60	-1.19	-0.80	-0.41	0.04	0.45	0.82	1.20	1.62	2.02	2.40	2.80
0.1	62	53	35	19	3	57	31	2	0	0	46	45	50	49	43
0.2	69	58	40	23	6	67	38	5	1	1	51	50	54	54	54
0.3	67	58	42	30	18	67	43	18	16	17	52	51	54	54	59
0.4	60	56	45	39	34	60	47	33	33	35	55	55	57	57	55
0.5	53	53	53	53	53	52	50	52	49	47	55	56	55	55	56
0.6	60	51	49	42	32	62	47	35	31	28	56	57	55	55	57
0.7	68	54	46	32	17	69	43	18	14	13	56	56	53	54	57
0.8	70	54	43	25	5	70	38	4	0	0	53	53	49	49	55
0.9	62	49	39	19	2	60	30	1	0	0	48	49	44	44	50

Notes

1. The table shows the frequency with which traders in the j 'th risk aversion quintile acted as contrarians over 100 simulations. $j = 1, \dots, 5$.
2. Each column relates to traders within quintile j , with the mean risk aversion parameter of traders in that quintile, $\bar{\gamma}_j$, shown at the head of the column.
3. π , shown at the start of each row, is the assumed true probability.
4. For each of the three blocks in the table $\bar{\gamma}$ is the assumed mean value of the entire distribution of traders' risk aversion parameters, the γ_k s; the assumed range of the γ_k s is $\bar{\gamma} \pm 1$.

Table 6 Risk aversion and profits, ‘000s

$\pi \backslash \bar{\gamma}_j$	$\bar{\gamma} = -2$					$\bar{\gamma} = 0$					$\bar{\gamma} = 2$				
	-2.78	-2.40	-2.01	-1.60	-1.19	-0.80	-0.41	0.04	0.45	0.82	1.20	1.62	2.02	2.40	2.80
0.1	-12.4	-7.9	-0.4	6.5	14.2	-4.0	-1.3	2.0	1.7	1.4	-0.0	-0.0	0.0	-0.0	0.0
0.2	-11.9	-6.9	-0.1	6.1	12.8	-4.1	-0.7	2.2	1.4	1.0	0.0	0.0	-0.0	-0.0	-0.0
0.3	-5.6	-3.4	0.2	3.1	5.8	-7.1	-1.4	4.4	2.7	1.5	0.0	0.0	-0.0	-0.0	-0.0
0.4	-1.1	-1.3	0.3	0.7	1.3	-2.6	-1.0	1.7	1.0	0.5	0.0	-0.0	-0.0	0.0	0.0
0.5	0.1	-0.1	0.1	0.1	-0.1	0.2	-0.0	-0.0	-0.0	0.1	-0.0	0.0	0.0	0.0	0.0
0.6	-1.2	-0.5	-0.6	0.8	1.6	1.0	-0.4	-0.2	-0.2	-0.2	0.0	0.0	-0.0	-0.0	0.0
0.7	-6.8	-3.6	-0.8	3.2	8.0	-5.5	-0.8	3.1	2.1	1.4	-0.0	0.0	-0.0	0.0	-0.0
0.8	-12.1	-5.7	-1.7	5.9	13.6	-6.2	-1.3	0.5	2.4	1.7	0.0	-0.0	0.0	-0.0	-0.0
0.9	-15.3	-8.4	-2.8	8.4	18.2	-9.6	-2.3	4.9	3.8	3.1	0.2	-0.1	0.1	-0.0	-0.1

Notes

1. The table shows the mean per-trader profits of traders in the j 'th risk aversion quintile. $j = 1, \dots, 5$.
2. Each column relates to traders within quintile j , with the mean risk aversion parameter of traders in that quintile, $\bar{\gamma}_j$, shown at the head of the column.
3. π , shown at the start of each row, is the assumed true probability.
4. For each of the three blocks in the table $\bar{\gamma}$ is the assumed mean value of the entire distribution of traders' risk aversion parameters, the γ_k s; the assumed range of the γ_k s is $\bar{\gamma} \pm 1$.

Table 7 Equilibrium market prices with heterogeneous rates of risk aversion

$\pi \backslash \bar{\gamma}$	-2	-1.5	-1	-0.5	0	0.5	1	1.5	2
0.1	0.34	0.32	0.29	0.25	0.21	0.17	0.12	0.10	0.09
0.2	0.39	0.38	0.35	0.32	0.29	0.25	0.21	0.20	0.20
0.3	0.43	0.42	0.41	0.38	0.36	0.34	0.31	0.30	0.30
0.4	0.47	0.46	0.45	0.44	0.43	0.42	0.40	0.40	0.40
0.5	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50	0.50
0.6	0.53	0.54	0.55	0.56	0.57	0.58	0.60	0.60	0.60
0.7	0.57	0.58	0.59	0.62	0.64	0.66	0.69	0.70	0.70
0.8	0.61	0.63	0.65	0.68	0.71	0.75	0.79	0.80	0.80
0.9	0.66	0.68	0.71	0.75	0.79	0.84	0.88	0.90	0.91

Notes

1. The table shows the mean equilibrium price from 100 simulations, as described in the text.
2. π , shown at the beginning of each row, is the assumed true probability.
3. $\bar{\gamma}$, at the top of each column, is the assumed mean value of the distribution of traders' risk aversion parameters; the assumed range of the γ_k s is $\bar{\gamma} \pm 1$.

Table 8 Equilibrium market prices with heterogeneous rates of risk aversion and biased preferences**Panel A All traders are risk averse**

$\pi \backslash \bar{\gamma}$	Bias against buying			No bias			Bias for buying		
	1	1.5	2	1	1.5	2	1	1.5	2
0.1	0.09	0.06	0.05	0.12	0.10	0.09	0.16	0.19	0.23
0.2	0.17	0.13	0.10	0.21	0.20	0.20	0.28	0.34	0.40
0.3	0.25	0.20	0.16	0.31	0.30	0.30	0.38	0.46	0.52
0.4	0.33	0.28	0.23	0.40	0.40	0.40	0.48	0.56	0.61
0.5	0.42	0.36	0.30	0.50	0.50	0.50	0.57	0.64	0.70
0.6	0.51	0.45	0.39	0.60	0.60	0.60	0.66	0.73	0.77
0.7	0.61	0.55	0.48	0.69	0.70	0.70	0.75	0.80	0.84
0.8	0.72	0.66	0.60	0.79	0.80	0.80	0.83	0.87	0.90
0.9	0.83	0.81	0.77	0.88	0.90	0.91	0.90	0.94	0.95

Panel B All traders are risk-lovers

$\pi \backslash \bar{\gamma}$	Bias against buying			No bias			Bias for buying		
	-2	-1.5	-1	-2	-1.5	-1	-2	-1.5	-1
0.1	0.33	0.30	0.27	0.34	0.32	0.29	0.36	0.33	0.31
0.2	0.38	0.36	0.34	0.39	0.38	0.35	0.41	0.39	0.37
0.3	0.43	0.41	0.39	0.43	0.42	0.41	0.44	0.43	0.42
0.4	0.46	0.45	0.44	0.47	0.46	0.45	0.48	0.47	0.46
0.5	0.49	0.49	0.49	0.50	0.50	0.50	0.51	0.51	0.51
0.6	0.52	0.53	0.54	0.53	0.54	0.55	0.54	0.55	0.56
0.7	0.56	0.57	0.58	0.57	0.58	0.59	0.57	0.59	0.61
0.8	0.59	0.61	0.63	0.61	0.63	0.65	0.62	0.64	0.66
0.9	0.64	0.67	0.69	0.66	0.68	0.71	0.67	0.70	0.73

Notes

1. The table shows the mean equilibrium price from 100 simulations, as described in the text. In each Panel the central block is the same as the equivalent columns shown in table 7, the baseline results.
2. π , shown at the beginning of each row, is the assumed true probability.
3. $\bar{\gamma}$, at the top of each column, is the assumed mean value of the distribution of traders' risk aversion parameters; the assumed range of the γ_k s is $\bar{\gamma} \pm 1$.
4. Preference for or against buying a contract is modelled by adding a multiple of traders' wealth to their demand functions. The multiple is uniformly distributed between 0 and 1.


Table 9 Equilibrium market prices with heterogeneous rates of risk aversion and limited shorting

$\pi \backslash \bar{\gamma}$	-2	-1.5	-1	-0.5	0	0.5	1	1.5	2
0.1	0.37	0.35	0.32	0.29	0.25	0.20	0.14	0.12	0.12
0.2	0.41	0.40	0.38	0.35	0.33	0.28	0.23	0.22	0.22
0.3	0.45	0.44	0.43	0.41	0.39	0.37	0.33	0.32	0.32
0.4	0.48	0.48	0.47	0.47	0.46	0.45	0.43	0.43	0.42
0.5	0.51	0.52	0.52	0.52	0.53	0.53	0.53	0.52	0.52
0.6	0.55	0.56	0.57	0.59	0.60	0.61	0.62	0.62	0.63
0.7	0.58	0.60	0.62	0.65	0.67	0.70	0.72	0.72	0.73
0.8	0.63	0.65	0.69	0.72	0.75	0.78	0.81	0.83	0.83
0.9	0.69	0.72	0.76	0.80	0.84	0.87	0.91	0.93	0.93

Notes

1. The table shows the mean equilibrium price from 100 simulations, as described in the text. In these simulations 67% of traders were assumed to choose not to trade rather than take the short side of the contract. This tendency was randomly assigned.
2. π , shown at the beginning of each row, is the assumed true probability.
3. $\bar{\gamma}$, at the top of each column, is the assumed mean value of the distribution of traders' risk aversion parameters; the assumed range of the γ_k s is $\bar{\gamma} \pm 1$.


Figure 1 Holdings of highest loss-making trader



Notes

The figure shows the cumulative shareholdings for the 2012 market of the trader who made the biggest total losses in the 2008 and 2012 election markets (the pattern of holdings for this trader in 2008 is similar). 'Rep candidates' denotes shares of Republican-to-win contracts and 'Dem candidates' denotes shares of Democrat-to-win contracts.


Figure 2 Number of traders per day during campaign periods



Notes

The figure shows the number of traders operating on Intrade each day during the last 250 days of the two election campaigns.


Figure 3 Equilibrium prices for varying degrees of risk aversion



Notes

The figure shows a subset of the data given in table 7, giving the mean equilibrium price generated by 100 simulations. For each set of simulations $\bar{\gamma}$ is the assumed mean value of the distribution of γ_k s and the assumed range of the γ_k s is $\bar{\gamma} \pm 1$. The thin line which bisects the graph is the 45° line.

Figure 4 Equilibrium market prices with heterogeneous rates of risk aversion and biased preferences




Notes


The figure shows a subset of the data presented in Panel A of table 8. The top part of the graph shows mean prices when (risk-averse) traders have a strong bias for buying the contract. The bottom part shows mean prices when (risk-averse) traders have a strong bias against buying the contract.

Figure 5 Thin markets

Panel A Market consists of risk-averse traders



Panel B Market consists of risk-loving traders



Notes

1. The figure shows mean, minimum and maximum equilibrium prices generated by 100 simulations of a market containing only 10 traders.
2. In Panel A all the traders are risk averse, with $\bar{\gamma}=2$ and all γ_k s in the range $\bar{\gamma} \pm 1$. In Panel B they are all risk-lovers with $\bar{\gamma}=-2$ and all γ_k s in the range $\bar{\gamma} \pm 1$.
3. Wealth is log-normally distributed around 500 with a Gini coefficient of 0.8.