

Interpreting Prediction Market Prices as Probabilities

Justin Wolfers

Wharton, University of Pennsylvania,
CEPR, IZA and NBER

jwolfers@wharton.upenn.edu
www.nber.org/~jwolfers

Eric Zitzewitz

Stanford GSB

ericz@stanford.edu
<http://faculty-gsb.stanford.edu/zitzewitz>

Abstract

While most empirical analysis of prediction markets treats prices of binary options as predictions of the probability of future events, it has recently been argued that there is little existing theory supporting this practice. We provide relevant analytic foundations, describing sufficient conditions under which prediction markets prices correspond with mean beliefs. Beyond these specific sufficient conditions, we show that for a broader class of models prediction market prices are usually close to the mean beliefs of traders. The key parameters driving trading behavior in prediction markets are the degree of risk aversion and the distribution on beliefs, and we provide some novel data on the distribution of beliefs in a couple of interesting contexts. We find that prediction markets prices typically provide useful (albeit sometimes biased) estimates of average beliefs about the probability an event occurs.

This draft: *January 8, 2007*

Keywords: Prediction markets, information aggregation, belief heterogeneity, binary options
JEL Codes: C53, D4, D84, G13, G14

We would like to thank Christopher Adams, Ray Fair, Brian Galebach, Robin Hanson, Robert Hahn, Ed Kaplan, Brian Knight, Charles Manski, Marco Ottaviani, David Pennock, Erik Snowberg, Peter Norman Sørensen, Betsey Stevenson, Rob Stevenson, Abe Wickelgren and seminar participants at Penn State, Stanford GSB, UNC, Wharton, the LBS Conference on Information Markets and the AEA meetings for useful conversations and insights. Wolfers gratefully acknowledges the support of a Hirtle, Callaghan & Co. – Geewax, Terker and Co. Research Fellowship, Microsoft Research, and the support of the Zull/Lurie Real Estate Center.

1. Introduction

Prediction markets are an increasingly widely used information aggregation device in academic research and public policy discussions, and the prices of contracts tied to events as diverse as the re-election of President Bush, the ouster of Saddam Hussein, next month's non-farm payrolls number, or the success of specific products have been interpreted as though interchangeable with "the market's beliefs". In these markets one typically trades a security that pays \$1 if a specified event occurs, and so the convention of treating the price of this security as a market-generated probability estimate has some intuitive appeal. Yet, as Manski (2006) notes, there exists very little theory describing the link from the underlying distribution of beliefs about the event to the equilibrium prediction market price.

This paper presents some simple theory, calibration results, and field evidence on the mapping from the distribution of beliefs to equilibrium prediction market prices. Our analytic results yield sufficient conditions for prediction market prices to equal mean beliefs.¹ Yet when relaxing some of our restrictive assumptions, we find that the prices may diverge from mean beliefs. Thus we turn to a quantitative exercise, keeping our baseline model simple enough that it is amenable to calibration. We find that empirically relevant parameterizations of beliefs and trading behavior, yield prediction market prices that approximately equal the average belief among traders.

The question we seek to answer is how prices relate to subjective beliefs about probabilities. This question is important for interpreting studies that use a prediction market as an indicator of the beliefs that are impounded in other asset market prices.² It is subtly different from a related question: do prediction market prices provide accurate estimates of event probabilities? This latter question is addressed by numerous recent empirical studies that examine whether the likelihood of a binary prediction market

¹ In concurrent work, Steven Gjerstad (2005) derives a number of related results.

² Examples include Slemrod and Greimel's (1999) analysis of the effect on municipal bond prices of changes in Steve Forbes' election probability; Leigh, Wolfers and Zitzewitz (2003) on the equity market consequences of changes in the probability of an Iraq war; Knight (2006) or Snowberg, Wolfers, and Zitzewitz (2007) on the equity market consequences of shocks to the likelihood of Republican Presidential election victories; and Fisman, Fisman, Galef and Khurana (2006) analyzing the effects on Cheney-connected firms of various political shocks.

contract paying off is equal to its price, which have in general concluded that they are.³ Naturally, our analysis of belief aggregation in prediction markets can help one better understand their predictive performance.

We proceed as follows. The next section sketches a simple model in which prediction market prices do coincide exactly with the mean of the belief among traders. The following section generalizes the model, showing that prediction market prices can deviate from mean beliefs, but that this deviation is typically small. The extent of the deviation depends crucially on how widely dispersed beliefs are, and so in the final section we present field evidence on this point. To preview, our results suggest that while prediction market prices and mean beliefs may diverge, they are typically very close. We interpret our results as providing a microfoundation for the claim that prediction markets (approximately) efficiently aggregate beliefs.

2. A Simple Model

We consider a simple prediction market in which traders buy and sell an all-or-nothing contract (a binary option) paying \$1 if a specific event occurs, and nothing otherwise. There is heterogeneity in beliefs among the trading population, and we denote trader j 's belief that the event will occur as q_j .

These beliefs are uncorrelated with wealth levels ($y \sim G(y)$), and are drawn from a distribution, $F(q)$. Individuals are price-takers and trade so as to maximize their subjectively expected utility. Wealth is only affected by the event via the prediction market, so there is no hedging motive for trading the contract.

We first consider the case where traders have log utility. In deciding how many contracts, x , to buy at a price π , traders solve the following problem:

³ Analyses of prediction market unbiasedness include Berg, et. al. (2003) for the Iowa Market's vote share markets, Wolfers and Zitzewitz (2006) for their binary winner-take-all markets, Wolfers and Zitzewitz (2004) for the Hollywood Stock Exchange's markets, Gürkaynak and Wolfers (2005) for Economic Derivatives, and Tetlock (2004), Borghesi (2006), and Zitzewitz (2006) for TradeSports' sports and financial markets. In addition, Thaler and Ziemba (1994) and Snowberg and Wolfers (2006) provide evidence for horse racing markets, and Zitzewitz (2006) provides evidence for CBOE option spreads that approximate binary options. Some of these studies report evidence of a "favorite-longshot bias", in which securities probabilities close to the extremes are mispriced. Chen et. al. (2006) and Fair (2006) both provide interesting model-based interpretations of Tradesports-generated probabilities.

$$\begin{aligned} \text{Max}_{\{x_j\}} EU_j &= q_j \text{Log}[y + x_j(1 - \pi)] + (1 - q_j) \text{Log}[y - x_j\pi] \\ \text{yielding: } x_j^* &= y \frac{q_j - \pi}{\pi(1 - \pi)} \end{aligned} \quad (1)$$

Thus, individual demand is:

- *Zero when prices equal beliefs:* When traders believe that the event is more likely than the price ($q > \pi$) they have positive demand, and they are net suppliers if $q < \pi$.
- *Linear in beliefs:* For traders with a given wealth level (y), demand increases linearly with their beliefs.
- *Decreasing in risk:* For a given expected return ($q - \pi$), greater risk (π close to $1/2$) yields smaller demand.
- *Homothetic:* Demand for these contracts rises proportionately with initial wealth, y .
- *Unique:* Only for prices between 0 or 1. (We will confine our attention to interior solutions.)

The prediction market is in equilibrium when supply equals demand:

$$\int_{-\infty}^{\pi} \int y \frac{q - \pi}{\pi(1 - \pi)} dG(y) dF(q) = \int_{\pi}^{\infty} \int y \frac{\pi - q}{\pi(1 - \pi)} dG(y) dF(q) \quad (2)$$

If beliefs (q) and wealth (y) are uncorrelated, then this implies:

$$\frac{y}{\pi(1 - \pi)} \int_{-\infty}^{\pi} (q - \pi) f(q) dq = \frac{y}{\pi(1 - \pi)} \int_{\pi}^{\infty} (\pi - q) f(q) dq \quad (3)$$

$$\text{and hence: } \pi = \int_{-\infty}^{\infty} q f(q) dq = \bar{q} \quad (4)$$

The monotonicity of demand in expected returns ($q - \pi$) implies that this is the only price for which aggregate demand will be zero for the given set of beliefs.

Thus, in this simple model, market prices are equal to the mean belief among traders. It is worth noting then that prices depend on the beliefs of all traders, and not simply an unspecified “marginal trader.” To see the relevant intuition, note that when all

traders optimally determine the size of their investment, they must all be marginal with respect to the last contract purchased.

The source of the heterogeneity in beliefs merits some attention. To be clear: the beliefs we analyze in our theoretical model motivate trading at the market price, and hence must be posterior beliefs after observing price. We do not explicitly model any learning, including learning from the trading process, through which traders arrive at these posteriors, but instead seek to identify the market price that is consistent with a given set of posterior beliefs. We restrict attention to the price-posterior belief relationship because our goal is to assess the appropriateness of using prediction market prices as proxies for beliefs impounded in other asset prices. In addition, when we calibrate our model, it will be using field data on beliefs that are likely to be posterior.

It is worth noting the possible sources of belief heterogeneity. One possibility is that trade is motivated by uncommon priors, and Morris (1995) argues that this is particularly likely when the event has been observed with insufficient frequency to suggest a common frequentist prior. This seems a useful interpretation for the types of events often traded in prediction markets.⁴ An alternative would be differences in information that motivate heterogeneous beliefs, but as Sebenius and Geanakoplos (1983) and other no-trade results show, these information asymmetries cannot induce two individuals to bet against each other: As each trader indicates a willingness to bet, she reveals some of her information, leading beliefs to converge to a degenerate distribution (at the market price). Given the active trade that we observe in many prediction and betting markets and the fact that this trade seems unlikely to be fully explained by liquidity, hedging, or entertainment motives, one might be skeptical that full convergence occurs in reality. Aumann (1976) suggests that a behavioral bias—such as anchoring—may lead traders to update their beliefs insufficiently in light of disconfirming evidence, such as the willingness of others to bet. Under this hypothesis, differences in beliefs among imperfect Bayesians persist.⁵ This is consistent not only with the trade we observe, but also anecdotal evidence that many prediction market participants expect to

⁴ For instance, Morris (1995, p.248) suggests “some safe generalizations about where individuals are more likely to agree... and where they are less likely to (political prospects in Iraq...)”.

⁵ Relatedly, Harris and Raviv (1993) suggest that even under common priors with common information, if traders differ in how they interpret news, this yields uncommon posteriors that motivate trade.

earn money, and data we discuss below that reveal substantial heterogeneity of beliefs, even in settings where one is reasonably certain that respondents have observed market prices.⁶

3. Generalizing the Model

We now turn to relaxing some of the assumptions made above. To preview, relaxing the assumption that budgets are independent of beliefs yields the intuitively plausible result that prediction market prices are a wealth-weighted average of beliefs among market traders. And second, the result that the equilibrium price is exactly equal to the (weighted) mean of beliefs reflects the fact that demand for the prediction security is linear in beliefs, which is itself a byproduct of assuming log utility. Calibrating alternative utility functions, we find that prices can systematically diverge from mean beliefs, but that this divergence is typically small. (The advantage of our very simple model is that it is easily amenable to calibration exercises.)

What if Beliefs and Wealth are Correlated?

Consider traders drawn from a distribution $F(q,y)$, where $E[q,y] \neq 0$. Thus, equation (1) continues to describe the demand of individual traders, given their beliefs and wealth levels. However, our equilibrium condition now changes, as traders with specific beliefs will have greater weight in the market. As before, equilibrium requires that supply equals demand:

$$\int y \frac{q - \pi}{\pi(1 - \pi)} dF(q \leq \pi, y) = \int y \frac{\pi - q}{\pi(1 - \pi)} dF(q \geq \pi, y) \text{ and hence: } \pi = \int q \frac{y}{y} F(q, y) \quad (5)$$

⁶ Two subsequent papers model learning explicitly and provide further useful results. Adams (2006) provides a multiperiod version of our model in which agents update their beliefs based on prices revealed in earlier periods; he shows that this leads prices to eventually converge to the mean of prior beliefs. Ottaviani and Sorensen (2006) analyze a fully-revealing rational expectations equilibrium, showing that if traders share similar priors, beliefs converge to the Bayesian posterior aggregating all private signals, and the market price equals this shared posterior. With heterogeneous priors, beliefs no longer converge, but the market price remains a generalized average of the Bayesian posteriors of market participants. Under CARA preferences, this generalized average corresponds to the rational posterior that aggregates all private information, given a “market prior” which in turn is a weighted average of the priors of traders (specifically, the log-likelihood ratios of the market prior are a risk-tolerance-weighted average of the log-likelihoods of individual traders). The interaction of heterogeneous priors and DARA or IARA preferences yields a wealth effect that confounds these neat results, although the quantitative significance of this channel remains unclear.

Thus, as before, the market equilibrium price is an average of beliefs in the trading population, although each belief is weighted according to the average wealth of traders holding that belief, relative to the average wealth level (\bar{y}).

If long-run market forces lead those with a history of accurate evaluation to become wealthier, then this wealth-weighted average may be a more accurate predictor than an unweighted average. Furthermore, if we are interested in using the prediction market price as a proxy for the beliefs of the marginal investor in other asset markets and if the wealthy trade more in these other markets, then the wealth-weighted average may again be closer to the object of interest.

Alternative Utility Functions

The key to the analytic simplicity of the above results is that individual demands in equation (1) are linear functions of each agent's beliefs. This, in turn, reflects the convenient assumption of log utility. As we relax this assumption, simple analytic results will be less easy to obtain.

Before analyzing the general case, it is worth analyzing an interesting special case, involving demand functions and distributions of beliefs that are symmetric. If:

- i) Individual demand for the prediction security is a function of the difference between beliefs and market prices and symmetric around zero:

$$x(q-\pi) = -x(\pi-q).$$

- ii) The distribution of beliefs is symmetric: $f(q - \bar{q}) = f(\bar{q} - q)$

then $f(q)x(q)$ will be symmetric and supply will equal demand if and only if price is the mean of beliefs.

These *dual symmetry* assumptions are sufficient to imply that equilibrium prices are equal to the mean beliefs of traders. Thus, the assumption of log utility is sufficient, but not necessary for this result. These dual symmetry conditions are likely to be met for distributions of beliefs symmetric around $\bar{q} = 1/2$, as long as traders are not affected by

framing issues.⁷ This suggests that we might generally expect prediction markets to be particularly accurate for prices close to $\frac{1}{2}$. (We illustrate this in section 3.)

Beyond these special assumptions though, different utility functions and distributions of beliefs will yield less analytically neat expressions, and raise the possibility that market prices diverge from mean beliefs. Our aim is to evaluate the extent to which plausible parameters yield important divergence.

In the general case, the trader's utility maximization problem becomes:

$$\begin{aligned} \text{Max}_{\{x_j\}} \quad & EU_j = q_j U(y + x_j(1 - \pi)) + (1 - q_j) U(y - x_j \pi) \\ \text{FOC:} \quad & \frac{U'(y + x_j - \pi x_j)}{U'(y - \pi x_j)} = \frac{\pi}{1 - \pi} \frac{(1 - q_j)}{q_j} \end{aligned} \quad (6)$$

which yields an interior solution when the trader is risk averse and a corner solution of betting one's entire wealth when the trader is risk neutral or risk loving. Specific functional forms allow us to solve for individual demand functions, and by specifying the distribution of beliefs, to solve for the relationship between equilibrium prices and mean beliefs. Table 1 shows the individual demand functions for a few widely used functional forms.

Figure 1 plots demand as a function of beliefs for $\pi = \$0.33$ for a range of interesting utility functions. Four features are evident in the chart. First, all involve investors increasing their demand as the divergence between their beliefs and the market price increase. Second, the aggressiveness with which traders respond to perceived profit opportunities is dictated by their risk aversion. Third, all functions involve traders investing nearly all of their wealth as the investment approaches a perceived "sure thing". Fourth, while only log utility yields an exactly linear demand function, almost all of these functions are *approximately* linear over the range in which beliefs are within about 10-20 percentage points of the price. This suggests that for distributions of beliefs that are not too disperse, market prices will be quite close to the mean belief across traders.⁸

⁷ That is, we assume that investors are indifferent between economically equivalent long and short trades, such as buying a security at π with belief q , and shorting at $1 - \pi$ with belief $1 - q$, so that $x(\pi, q) = -x(1 - \pi, 1 - q)$. At a price of $\frac{1}{2}$, this implies $x(\frac{1}{2}, q) = -x(\frac{1}{2}, 1 - q)$ and hence demand is symmetric.

⁸ As Manski notes, given that we are analyzing posterior beliefs, partial updating of priors in light of the market price would lead the distribution of beliefs to become less disperse.

The only exception to this approximate local linearity is the risk-neutral investor, who always invests her entire wealth whenever market prices diverge from her beliefs. Interestingly, this is the only case considered by Manski (2006), adapting the example in Ali (1977).⁹ Not surprisingly, these strong assumptions yield strong implications. Specifically for \$0.33 to be an equilibrium requires twice as many sellers as buyers at this price (because risking one's entire wealth allows each buyer to purchase twice as many contracts as a seller). Noting that a risk-neutral investor switches from buying to selling when the price falls below her belief, this implies that a price of \$0.33 corresponds to the 67th percentile of the distribution of beliefs. Manski shows that this logic extends through the price distribution, and hence in that case, the equilibrium price corresponds not with the mean of beliefs, but rather with the $100(1-\pi)$ th percentile of the distribution of beliefs.

As noted, mean beliefs and prices exactly coincide for the log utility case. While we will solve for equilibrium prices below, we can first use the demand curves shown in Figure 1 to provide some intuition about the sign of any divergence between prices and mean beliefs. Specifically, the chart shows that traders with low risk aversion ($\gamma < 1$) trade particularly aggressively on longshots. This suggests that if beliefs are distributed symmetrically that prices will be biased toward $\frac{1}{2}$. Further, if the mean beliefs are accurate predictors, this implies a "favorite-longshot bias", in which longshots are priced above their objective probability. Analogously, the demand functions for more risk-averse investors ($\gamma > 1$) suggest that investors with extreme beliefs have a greater effect on prices than in the log utility case. As such, prices will tend to be biased toward zero or one, and a "reverse favorite-longshot bias" may occur.

Mapping Prices to Probabilities

Given the individual demand functions derived above, all that is required to solve for equilibrium prices is a distribution of beliefs. We start by assuming that beliefs are drawn from a uniform distribution with a range of 10 percentage points, and solve for the mapping between mean beliefs and prices implied by each of the utility functions shown in Figure 1. (We rescale beliefs outside the (0,1) range to 0 or 1.)

⁹ Similarly, risk-loving investors also invest their entire wealth, although given that they value the opportunity to gamble, the discontinuity in their demand function can occur at a different point.

Figure 2 shows that for moderately dispersed beliefs, prediction market prices tend to coincide fairly closely with the mean beliefs. While there is some divergence, it is typically within a percentage point, although the risk neutral model yields larger differences. Greater risk-aversion leads to a bias toward more extreme prices, while lesser risk aversion leads prices to be biased toward \$0.50. The divergence between prices and average beliefs is greatest for prices closest to \$0 and \$1, although behavior at the extremes partly reflects the distribution of beliefs becoming increasingly skewed as the mass point at 0 or 1 grows.

Figure 3 shows the mapping from prices to probabilities when beliefs are more disperse (in this case the standard deviation and range were doubled). As the dispersion of beliefs widens, the number of traders with extreme beliefs increases, and hence the non-linear response to the divergence between beliefs and prices is increasingly important. As such, the biases evident in Figure 2 become even more evident as the distribution of beliefs widens. Even so, for utility functions with standard levels of risk-aversion, these biases are small.

Finally, in Figure 4 we show that alternative utility functions yield similar implications. We have also experimented with uniform, beta and log-normal distributions of beliefs, and these results are also similar.

Cumulatively, figures 2-4 show six main patterns. First, under all conditions, log utility yields prices that coincide exactly with mean beliefs. Second, for other utility functions the divergence between prices and mean beliefs is generally quite small. Third, this divergence is zero when beliefs are symmetric around a 50% probability¹⁰ and generally very small when prices are in the \$0.20-\$0.80 range.¹¹ Fourth, the sign of the deviation between prices and beliefs varies with the assumed utility function. Fifth, increasingly disperse beliefs yield a larger gap between prices and mean beliefs. And sixth, the extent of the divergence between prices and mean beliefs depends on the specific assumptions adopted about the utility function of traders, and the distribution of

¹⁰ For a formal proof, see theorem 2 in Gjerstad (2004).

¹¹ This generalization may not hold for particularly disperse beliefs, and especially when beliefs are bimodal. For instance, Manski considers maximally disperse distributions so as to establish bounds on mean beliefs implied by a price. Appendix A expands on this analysis.

beliefs (particularly when beliefs are close to zero or one). Manski's model consistently delivers the most extreme results.

We now turn to trying to extract some empirical evidence on the most relevant parameters, from field data.

4. Field Evidence on the Distribution of Beliefs

We begin with a very simple, but salient, example. Throughout 2003 and 2004, TradeSports ran a prediction market in a security that paid \$1 if President Bush were re-elected, and nothing otherwise. The price of this security on Election eve was \$0.55. At the same time, pre-election polls suggested that 62 percent of the population thought that President Bush was more likely to win than John Kerry.^{12,13}

The prediction market price and the poll result place restrictions on the distribution of beliefs. For these data to be useful, we must assume that these survey data are representative of beliefs among traders, and specifically their posterior beliefs in light of the information embodied in prediction market prices. For different utility and belief distribution functional forms, we can derive the belief distribution that matches these two facts and examine how the observed prediction market price relates to central moments of the implied belief distribution. (In the absence of contrary evidence, we continue to assume that wealth is uncorrelated with beliefs.)

If the perceived probability of a Bush victory is q , then the poll result and the prediction market price respectively imply that:

¹² We draw this number from various polls. Specifically, the following proportions thought Bush more likely to win in the final pre-election poll: CBS/NYT: 60% (n=920 adults polled 10/28-10/30); Gallup/CNN/USA Today: 61% (n=1013 adults polled 10/14-10/16); ABC: 62% (n=3617 adults polled 10/27-10/30); Marist: 60% (n=1300 registered voters polled 10/31); Pew: 64% (n=2804 registered voters polled 10/27-10/30); Princeton: 64% (n=1117 registered voters polled 10/27-10/39); Fox: 61% (n=1000 likely voters polled 10/17-10/18). Allocating non-respondents 50-50 (instead of dropping them) yields proportions predicting Bush that are usually around 2 percentage points lower. The Gallup question is roughly representative, asking: "Regardless of whom you support, and trying to be as objective as possible, who do you think will win the (presidential) election in November (2004) – John Kerry or George W. Bush?" The only real divergence was Fox, who asked: "Imagine you were given \$100 dollars to place a bet on the outcome of the upcoming (2004) presidential election. Which candidate – (George W.) Bush or (John) Kerry--would you put your money on to win this November?"

¹³ While we focus on jointly explaining prediction market prices and polling data on voter *forecasts*, Fair (2006) describes instead how to rationalize prediction market prices with polling data on voter *intentions*.

$$\int_{0.5}^1 f(q) dq = 0.62 \quad [62\% \text{ thought Bush more likely to win}]$$

$$\int_0^{0.55} X(q) f(q) dq = \int_{0.55}^1 X(q) f(q) dq \quad [\$0.55 \text{ was a prediction market equilibrium}]$$

In Table 2, we examine various two-parameter functions for the distribution of beliefs, and a range of utility functions, and solve for the implied mean belief. We report the mean of the implied distribution of beliefs (and display the parameters underlying these two-parameter distributions in parentheses).

Recalling that the market price of this security was \$0.55, Table 2 shows that this price is a good approximation to the mean belief under any of the specific assumptions that we examined.¹⁴ Note that this occurs despite the fact that in some cases the belief distributions needed to reconcile the market price and poll results are highly asymmetric. Naturally this robustness partly reflects the tendency hinted at earlier that prices close to \$0.50 are typically fairly accurate.¹⁵

Deriving distributions of beliefs from two data points and a distributional assumption may not be particularly satisfying, so we would like a setting where we observe beliefs directly. Unfortunately data surveying expectations about the likelihood of specific events for which there are prediction or other financial market prices is rather rare.

For this reason we turn to two rather unique datasets. The first was provided to us by Probability Football, an advertising-supported free contest that requires players to estimate the probability of victory in every NFL game in a season.¹⁶ Including the pre-season and playoffs, this yields 259 games in the 2000 and 2001 seasons and 267 in 2002 and 2003. On average we observe the probability assessments of 1320 players in each game, for a total sample size of 1.4 million observations. Contestants are scored using a

¹⁴ Note from Table 1 that with CARA utility, risk aversion does not affect the shape of the demand function, just its slope. The same is true of quadratic utility and y^{\max} . Since aggregate supply is zero, multiplying all trader's demand by a constant does not affect market prices, we report results for only one parameter value for CARA and quadratic utility. Likewise, HARA utility and CRRA utility yield betting functions of the same shape (again, allowing for a difference in slope), so we do not report results separately for HARA.

¹⁵ Interestingly, the table also shows that – for a given set of assumptions – market prices can also reveal the dispersion of beliefs in the population. That said, unlike the inferences about the mean beliefs, inferences about the underlying dispersion of beliefs are quite sensitive to the specific assumptions adopted.

¹⁶ Levitt (2004) analyzes a related sample from a different source. The advantage of our data is that they also include a measure of participants' beliefs. We are grateful to Brian Galebach for sharing these data.

quadratic scoring rule; they receive $100 - 400(w - q)^2$, points where w is an indicator variable for whether the team wins and q is the stated probability assessment. Truthfully reporting probabilities yields the greatest expected points, a fact that is explicitly explained to contestants.

The top three players receive cash prizes. While these rank-order incentives potentially provide an incentive to add variance to one's true beliefs, it turns out that given the number of games in a season, this incentive is small. For instance, in 2003, two mock entrants to this contest that simply used prices from TradeSports and the Sports Exchange (a sports-oriented play-money prediction market run by NewsFutures.com) as their probabilities placed seventh and ninth out of almost 2,000 entrants.¹⁷ We simulated strategies that took these prediction market prices and added or subtracted noise, finding that adding or subtracting 1 percentage point to the market price yielded the highest probability of winning a prize. Even so, quite a few players appear to believe that more variance is optimal; and about 40 percent of players report zero or one for at least ten percent of their games. This is a losing strategy, despite comprising 40 percent of players, they account for only 5 percent of those who make the top 5 percent.¹⁸ Since we are interested in learning about the distribution of beliefs, we drop all probability reports from these players. (Qualitatively this doesn't much affect our results.)

Figure 5 reports the distribution of all probability for games in which prediction market prices are close to \$0.33. Even after cleaning the data, there are still mass points at zero and one and some clustering at focal numbers. Beyond this, the distribution appears roughly normal.

Figure 6 examines how this distribution of beliefs varies with the (real-money) prediction market price. In general, prices closely approximate the mean (the dashed line) or median of beliefs. To the extent there is divergence, prices diverge away from \$0.50. This relationship is more consistent with what one would expect under high risk aversion ($CRRA > 1$). Recall that under log utility ($CRRA=1$), betting functions are linear, and the equilibrium price is the mean of beliefs. In contrast, with greater risk aversion, traders respond aggressively only when prices deviate substantially from their

¹⁷ Servan-Schreiber, Wolfers, Pennock and Galebach (2004) used the data collected from this to compare the predictive power of real and play-money markets, finding that they were roughly equal.

¹⁸ Probability Football has added strategy advice to its website that makes this point to players.

beliefs. As such, those with extreme beliefs have the most significant effect on prices (and indeed in the limit as risk aversion approaches infinity, prediction market prices are determined by the numbers of traders who are certain that the event will occur relative to those certain it will not). These forces push prices away from mean beliefs, towards zero for longshots, and one for favorites, as seen in Figure 6.

Figure 7 formalizes this intuition. Specifically we apply our model to our empirical data on beliefs, varying parameters of the utility functions. The figure plots the implied equilibrium prediction market price against the actual mapping. From this graph, it appears that predicted prediction market prices most closely approximate actual prediction market prices for a utility function with $CRRA = 5$.

Our second dataset of probability beliefs that can be matched with prices comes from the Michigan Survey of Consumers, which asks respondents for the probability a \$1000 investment in a diversified stock mutual fund will increase by 10 percent or more in the next year.¹⁹ The security corresponding to this event would be a binary equity index option with an expiry date one year from today and a strike price 10 percent above the current index level. While this exact option usually does not exist, we can estimate its price using the prices of related options.²⁰

Figure 8 plots our estimated binary option price along with prediction market prices that we simulate using beliefs from the Michigan survey and different assumptions about preferences.²¹ Beliefs in the Michigan survey are quite disperse, with a mean of 31

¹⁹ We are grateful to Charles Manski for providing this data. The survey is described in more detail in Dominicz and Manski (2004).

²⁰ For the purposes of this exercise, we interpret the pre-dividend returns of the S&P 500 as a proxy for the returns on a “diversified mutual fund.” This implicitly assumes that the dividend yield and mutual expense ratios are roughly equal, which they are for our time period (the average S&P 500 dividend yield was 1.78% per the *Ivy OptionMetrics* dataset and the average diversified stock fund expense ratio was 1.53% for year-end 2002 per the *CRSP Mutual Funds* dataset). For each day in our sample, we estimate the price of a binary option paying \$1 if the S&P 500 is more than 10 percent above its current level in one year. We do so by interpolating an estimated implied volatility for our hypothetical option using the implied volatilities for the CBOE S&P 500 index options with the nearest strike prices and expiry dates (obtained from the *OptionMetrics* dataset). We then calculate the binary option price using the derivative of Black’s (1976) pricing formula for options on futures with respect to strike price.

²¹ In simulating the prediction market prices, we correct for the fact that risk averse investors will have higher marginal utility in low-wealth states of the world. From 1980-2005, conditional on returning more than 10 percent in twelve consecutive calendar months, the S&P 500 has returned 22.6 percent on average. Conditional on returning less than +10 percent, it has returned -3.2 percent. This implies that an investor with a portfolio-average beta of 0.5 would have wealth that is 13 percent higher if our “S&P 500 up 10% or more” security pays off. If such an investor had log utility, she would value a security with an objective payoff probability of 0.3 at approximately 0.275. For each level of risk aversion, we account for this

percent and a standard deviation of 25 percent. Risk preferences therefore have a significant effect on our simulated prediction market price. The graph suggests that a CRRA of approximately 1 best reconciles beliefs and binary option prices.²²

While these results are interesting, they are obviously not a particularly robust way to measure risk aversion. Rather, we take the results as simply suggesting that models with moderate risk aversion are roughly consistent with our data on prediction market prices and the distribution of beliefs.

5. Conclusion

An old joke about academics suggests that we are often led to ask: “We know it works in practice, but does it work in theory?” This paper arguably follows that model. As discussed above, a variety of field evidence across several domains suggests that prediction market prices appear to be quite accurate predictors of probabilities. This paper suggests that this evidence is easily reconcilable with a theory in which traders have heterogeneous beliefs that are correct on average.

All of the models we have explored yield a monotonic mapping between prediction market prices and the mean of beliefs. Moreover, we have provided several sets of sufficient conditions under which prediction market prices exactly coincide with the mean of beliefs. More generally there can be a wedge between the two, but for most practical purposes, our simulations suggest that it is likely to be small. As such, we believe that this provides a logical rationale for our earlier assertion that “markets aggregate opinions” (Leigh, Wolfers and Zitzewitz, 2003).

Manski (2006) presented a specific example “under special assumptions that may constitute a best-case scenario” in which this wedge between prices and average beliefs was large. By contrast, our analysis endogenizes the decision as to whether and how much to trade, and we find that Manski’s special case is in fact a worst-case scenario.

difference when simulating the prediction market prices in Figure 8, although making the correction only materially affects the curves for CRRAs greater than 2. Given the much lower correlation between outcomes and wealth for the political or football prediction markets discussed above, we do not make this adjustment for those markets.

²² The Michigan Survey also asks respondents for a probability that the mutual fund will increase in value in general. For this question, beliefs are distributed roughly symmetrically around 50 percent, and the binary option price we calculate is also close to 50 percent. As a result, simulated prediction market prices approximate binary option prices for a wide range of risk preferences.

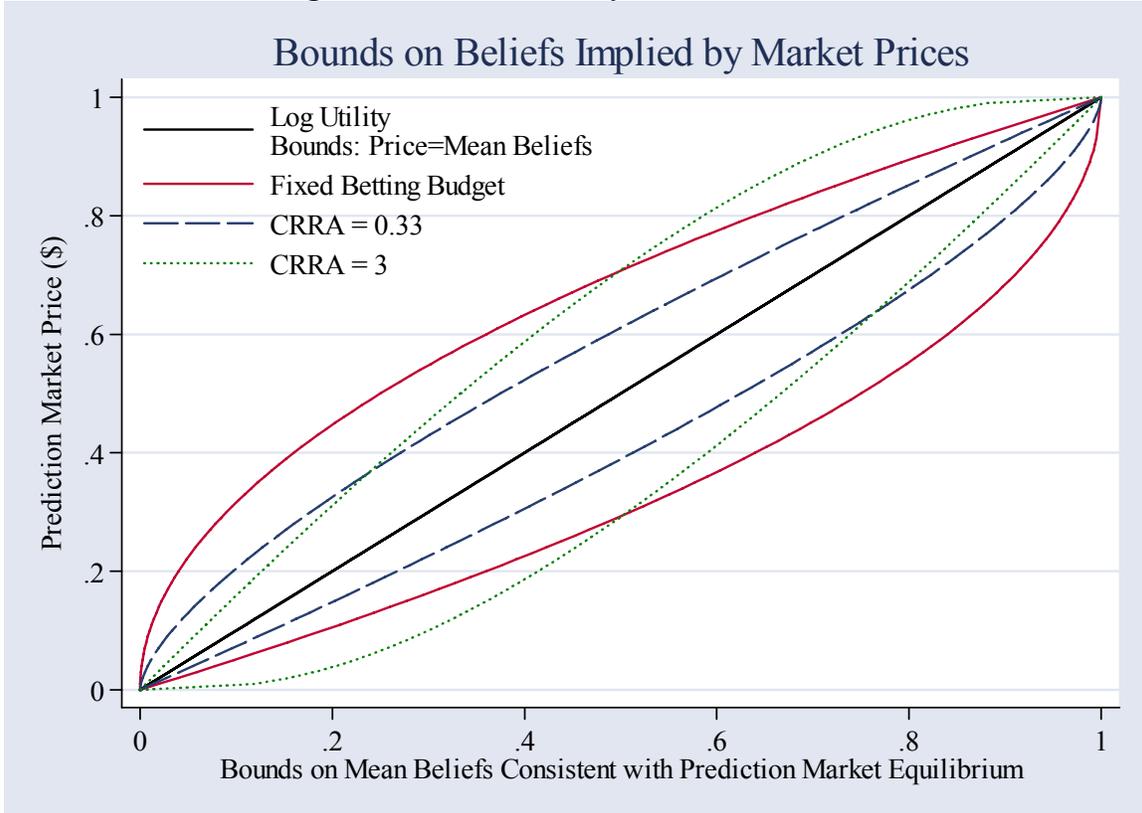
Moreover, while his worst-case scenario is at odds with observed field data, our model is consistent with observed data on the distribution of beliefs and prediction market prices.

Finally, we conclude with some guidance for practitioners. In most cases we find that prediction market prices aggregate beliefs very well. Thus, if traders are typically well-informed, prediction market prices will aggregate information into useful forecasts. The efficacy of these forecasts may however be undermined somewhat for prices close to \$0 or \$1, when the distribution of beliefs is either especially disperse, or when trading volumes are somehow constrained, or motivated by an unusual degree of risk-acceptance.

Appendix A: Bounds on Mean Beliefs

Manski considers the bounds on mean beliefs implied by a market price based on his fixed bet size model. Specifically, recall the result that the market price, π coincides with the $1-\pi^{\text{th}}$ percentile of the distribution of beliefs. Thus this price is consistent with many different distributions of beliefs: the lower bound reflects the two-point distribution $f(0)=1-\pi$ and $f(\pi+\varepsilon)=\pi$; the upper bound is generated by the distribution $f(\pi-\varepsilon)=\pi$ and $f(1)=1-\pi$ (for $\varepsilon \rightarrow 0$). As such, mean beliefs are bounded by $(\pi^2, 2\pi-\pi^2)$. These bounds are shown as solid lines in Figure A1.

Figure A1: Bounds Analysis in Four Models



We extend the spirit of Manski's bounds analysis to our other models. That is, we solve:

$$\text{Lower bound} = \text{Min}_{\{f(q)\}} \int_0^1 qf(q) dq \quad \text{Upper bound} = \text{Max}_{\{f(q)\}} \int_0^1 qf(q) dq$$

$$\text{subject to prediction market equilibrium: } \int_0^p x(q)f(q) dq = - \int_p^1 x(q)f(q) dq$$

It is easy to show that these bounds must be generated by two-point distributions. We map these bounds for each prediction market price between \$0 and \$1, for each utility function. Log utility yields prices that always coincide with mean beliefs. For other utility functions these two-point distributions yield some divergence between prices and probabilities, but the divergence is typically smaller than in the fixed bet size model. (Increasing risk aversion substantially makes extreme prices somewhat less informative.) By construction these bounds reflect extreme bimodal distributions of beliefs; for results with more standard distributions, see Figures 2-4.

6. References

- Adams, Christopher (2006), "Does Learning in Prediction Markets Lead to Information Aggregation?", *mimeo*, Federal Trade Commission.
- Ali, Mukhtar (1977), "Probability and Utility Estimators for Racetrack Bettors", *Journal of Political Economy*, 85(4), 803-815.
- Aumann, Robert J. (1976), "Agreeing to Disagree", *The Annals of Statistics*, 4(6).
- Berg, Joyce, Robert Forsythe, Forrest Nelson and Thomas Rietz (2003) "[Results from a Dozen Years of Election Futures Markets Research](#)," forthcoming in Charles Plott and Vernon Smith (Eds) *Handbook of Experimental Economic Results*.
- Black, Fisher (1976), "Studies of Stock Price Volatility Changes," in *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association*, 171-181.
- Borghesi, Richard (2006), "Underreaction to New Information: Evidence from an Online Exchange", *mimeo*, Texas State University.
- Chen, Keith, Jonathan E. Ingersoll and Edward H. Kaplan (2006), "Do States Swing Together? Evidence from Political Prediction Markets", *mimeo*, Yale University.
- Dominitz, Jeff and Charles Manski (2004), "How Should We Measure Consumer Confidence," *Journal of Economics Perspectives* 18(2), 51-66.
- Fair, Ray (2006), "Interpreting the Predictive Uncertainty of Presidential Elections", *mimeo*, Yale University.
- Fisman, David, Ray Fisman, Julia Galef and Rakesh Khurana (2006), "Estimating the value of connections to Vice-President Cheney", *mimeo*, Columbia University.
- Gjerstad, Steven (2005), "Risk Aversion, Beliefs, and Prediction Market Equilibrium", *mimeo* University of Arizona.
- Gürkaynak, Refet and Justin Wolfers (2005), "Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts, Uncertainty and Risk", *NBER International Seminar on Macroeconomics Annual*, 2005.
- Harris, Milton and Artur Raviv (1993), "Differences of Opinion Make a Horse Race", *The Review of Financial Studies*, 6(3).
- Knight, Brian (2006), "Are Policy Platforms Capitalized into Equity Prices? Evidence from the Bush/Gore 2000 Presidential Election", *Journal of Public Economics*, 90(4-5).
- Leigh, Andrew, Justin Wolfers and Eric Zitzewitz (2003), "What Do Financial Markets Think of War in Iraq?", NBER Working Paper #9587.

- Levitt, Steven (2004), "Why Are Gambling Markets Organised So Differently from Financial Markets?", *Economic Journal*, vol. 114.
- Manski, Charles (2006), "Interpreting the Predictions of Prediction Markets", *Economic Letters*, 91(3), 425-429.
- Milgrom, Paul and Nancy Stokey (1982), "Information, Trade, and Common Knowledge", *Journal of Economic Theory*, 26(1).
- Morris, Stephen (1995), "The Common Prior Assumption in Economic Theory", *Economics and Philosophy*, 11.
- Morris, Stephen (1996), "Speculative Investor Behavior and Learning", *Quarterly Journal of Economics*, 111(4).
- Ottaviani, Marco and Peter Norman Sørensen (2006), "Aggregation of Information and Beliefs in Prediction Markets", *mimeo*, London Business School.
- Sebenius, James and John Geanakoplos (1983), "Don't Bet on It: Agreements with Asymmetric Information", *Journal of the American Statistical Association*, 78(382).
- Servan-Schreiber, Emile, Justin Wolfers, David Pennock and Brian Galebach (2004) "Prediction Markets: Does Money Matter?", *Electronic Markets* 14(3).
- Slemrod, Joel and Timothy Greimel (1999), "Did Steve Forbes Scare the Municipal Bond Market?", *Journal of Public Economics*, 74(1), 81-96.
- Snowberg, Erik and Justin Wolfers (2006), "Understanding the Favorite-Longshot Bias: Risk Preferences versus Misperceptions", *mimeo*, University of Pennsylvania.
- Snowberg, Erik, Justin Wolfers and Eric Zitzewitz (2007), "Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections", *Quarterly Journal of Economics*, 122(2).
- Tetlock, Paul (2004), "How Efficient Are Information Markets? Evidence from an Online Exchange", *mimeo*, University of Texas.
- Thaler, Richard H., and William T. Ziemba (1988), "Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries," *Journal of Economic Perspectives*, 2(2).
- Wolfers, Justin and Eric Zitzewitz (2004), "Prediction Markets", *Journal of Economic Perspectives*, 18(2), 107-126.
- Wolfers, Justin and Eric Zitzewitz (2006), "Five Open Questions About Prediction Markets", in Robert Hahn and Paul Tetlock (eds) *Information Markets: A New Way of Making Decision in the Public and Private Sectors*, (Washington DC: AEI-Brookings Press), 2006.

Zitzewitz, Eric (2006), “Price Discovery Among the Punters: Using New Financial Betting Markets to Predict Intraday Volatility”, *mimeo*, Stanford University.

Table 1: Utility Functions and Demand for Prediction Securities

Utility Function	Utility	Demand
Log utility (CRRA with $\gamma = 1$)	$u(y) = \ln(y)$	$\frac{y}{\pi(1-\pi)}(q-\pi)$
Constant relative risk aversion (CRRA) ($\gamma > 0, \gamma \neq 1$)	$u(y) = \frac{y^{1-\gamma}}{1-\gamma}$	$\frac{y}{\pi} \cdot \left(\frac{\pi \{ [\frac{q(1-\pi)}{\pi(1-q)}]^{\frac{1}{\gamma}} - 1 \}}{1 + \pi \{ [\frac{q(1-\pi)}{\pi(1-q)}]^{\frac{1}{\gamma}} - 1 \}} \right)$
Constant absolute risk aversion (CARA)	$u(y) = -e^{-\gamma y}$	$r^{-1} \cdot [\ln(\frac{q}{1-q}) - \ln(\frac{\pi}{1-\pi})]$
Quadratic utility	$u(y) = -\frac{1}{2}(y^{\max} - y)^2$	$(y^{\max} - y) \frac{q - \pi}{q(1-\pi) - \pi(q-\pi)}$
Hyperbolic absolute risk aversion (HARA)*	$u(y) = \frac{1-\gamma}{\gamma} \left(\frac{ay}{1-\gamma} + b \right)^\gamma$	$\frac{y + b(1-\gamma)a^{-1}}{\pi} \cdot \left(\frac{\pi \{ [\frac{\pi(1-q)}{q(1-\pi)}]^{\frac{1}{\gamma-1}} - 1 \}}{1 + \pi \{ [\frac{\pi(1-q)}{q(1-\pi)}]^{\frac{1}{\gamma-1}} - 1 \}} \right)$

* The HARA utility function nests the others as special cases. (For log utility $\gamma \rightarrow 0$; risk neutral: $\gamma \rightarrow 1$; quadratic: $\gamma = 2$; CRRA: $\gamma < 1$ and $b=0$; CARA: $\gamma \rightarrow -\infty$ and $b > 0$).

Table 2. Distribution of beliefs about the probability that Republicans would win the 2004 election consistent with market price and poll results under different utility functions and distributional assumptions

	Normal [μ, σ]	Beta (α, β)	Uniform (q_L, q_H)
Implied Distribution of Beliefs			
Fixed bet size (Limit; $\gamma \rightarrow 0$)	0.578 [0.584, 0.278]	0.571 [2.112, 1.589]	0.586 [0.229, 0.942]
CRRA; $\gamma = 1/3$	0.560 [0.561, 0.201]	0.558 [3.370, 2.675]	0.575 [0.252, 0.897]
Log Utility ($\gamma = 1$)	0.550 [0.550, 0.163]	0.550 [4.640, 3.804]	0.550 [0.342, 0.758]
CRRA; $\gamma = 3$	0.546 [0.546, 0.149]	0.547 [5.337, 4.432]	0.549 [0.343, 0.755]
CRRA; $\gamma = 20$	0.544 [0.544, 0.144]	0.546 [5.640, 4.707]	0.548 [0.345, 0.752]
CARA; $\rho = 3$.544 [0.544, 0.144]	0.545 [5.692, 4.754]	0.547 [0.351, 0.743]
Quadratic; $y^{\max} = 3$.542 [0.542, 0.138]	0.542 [6.568, 5.553]	0.546 [0.351, 0.742]

Notes: Table shows mean of distribution. [Parameters of the belief distribution shown in parentheses]

Source: Authors' calculations. Note that beliefs outside (0,1) were treated as $\lim. q \rightarrow 0$ or 1, respectively.

Figure 1: Individual Demands Under Alternative Utility Functions

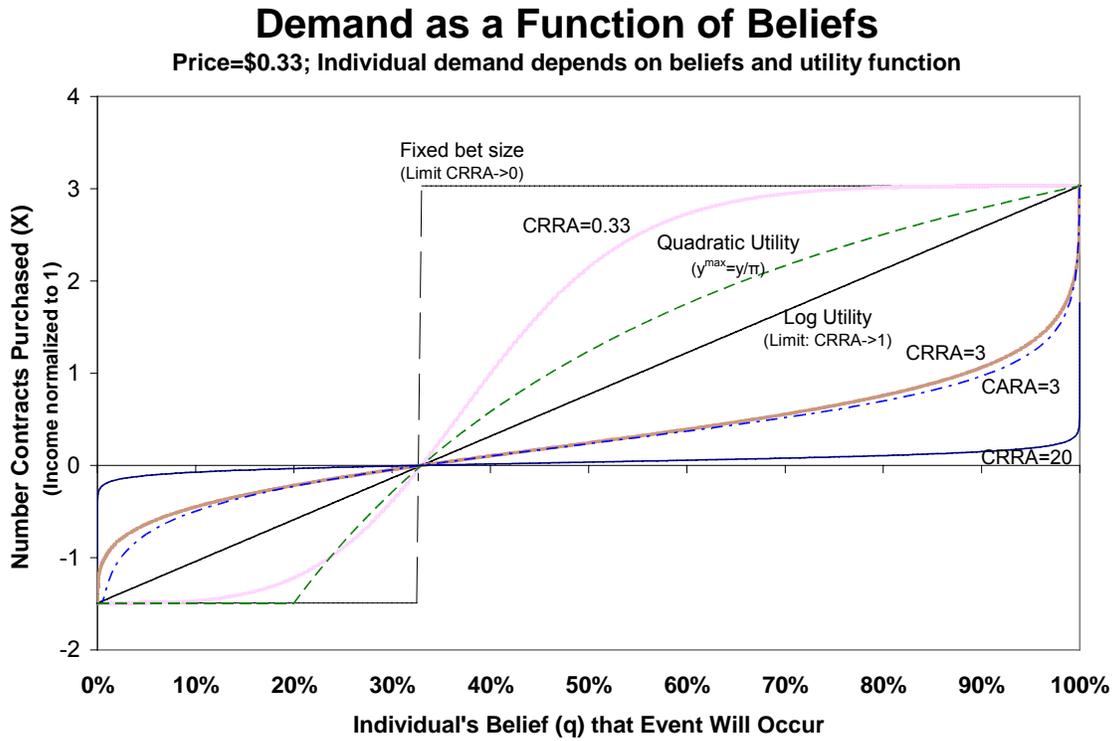


Figure 2: Mapping Prices to Probabilities

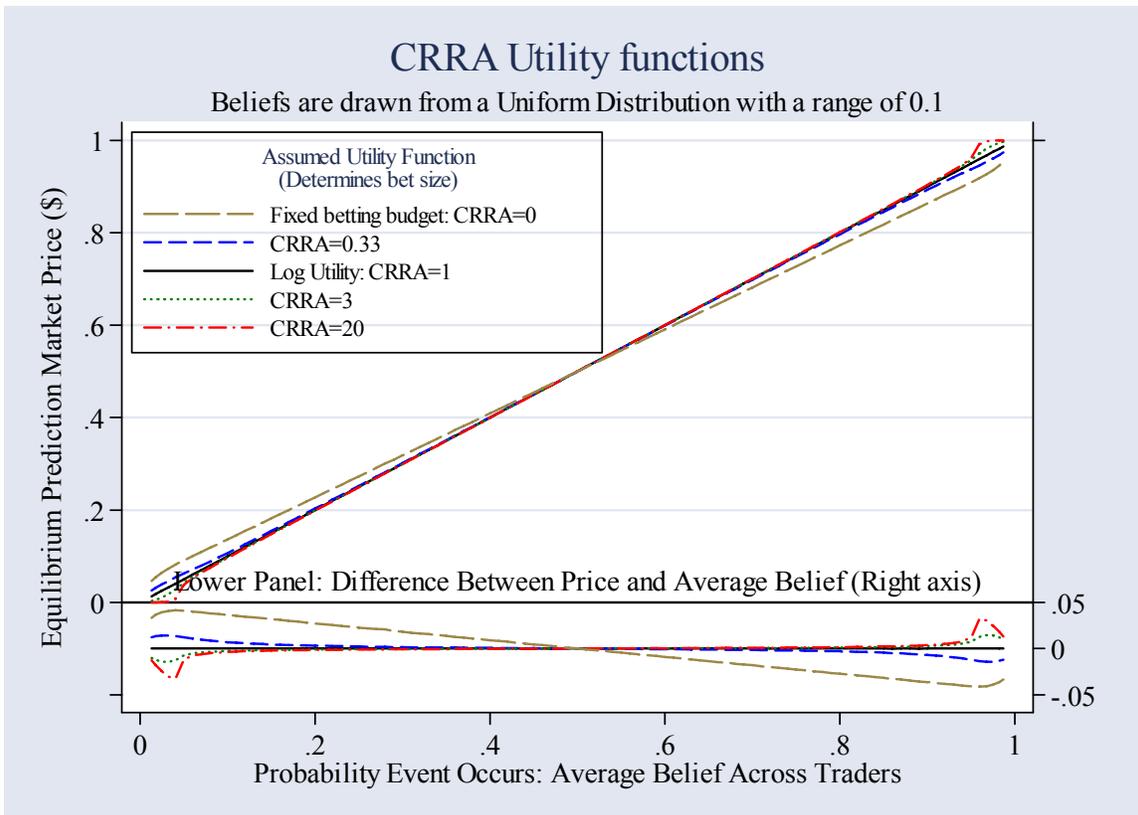


Figure 3: Increasing Dispersion in Beliefs

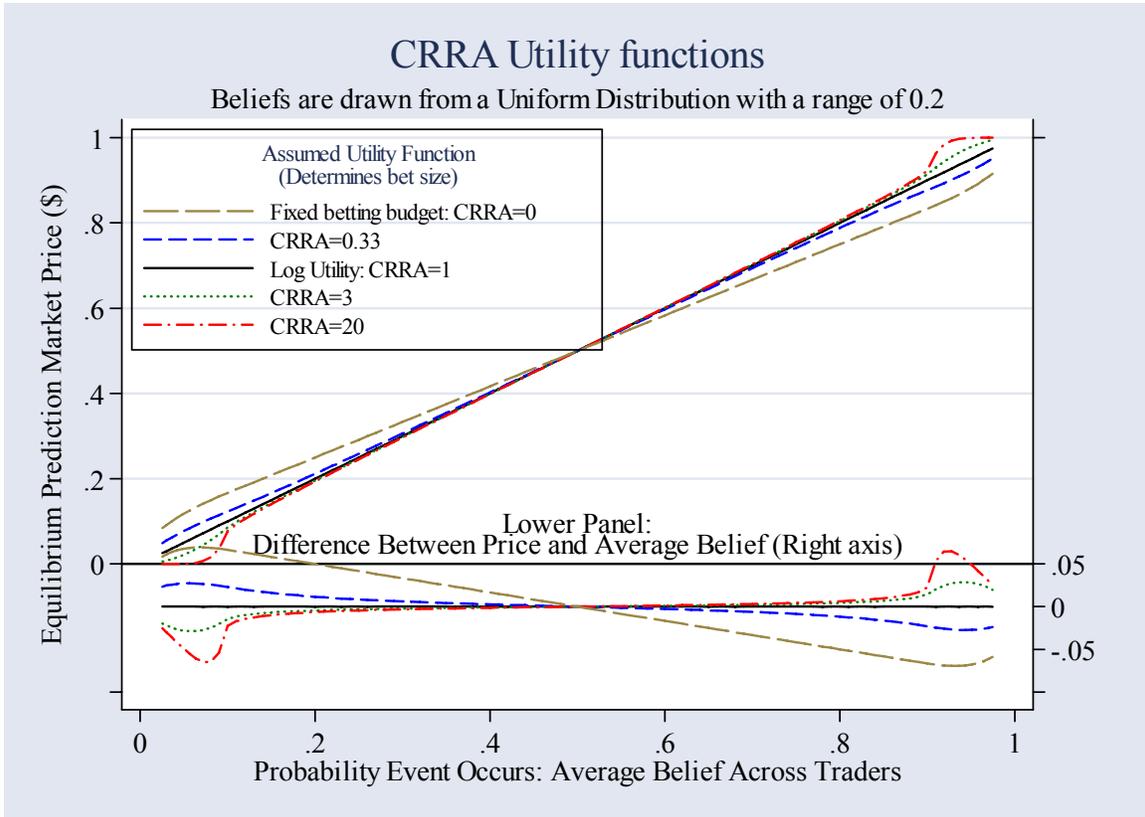


Figure 4: Alternative Utility Functions

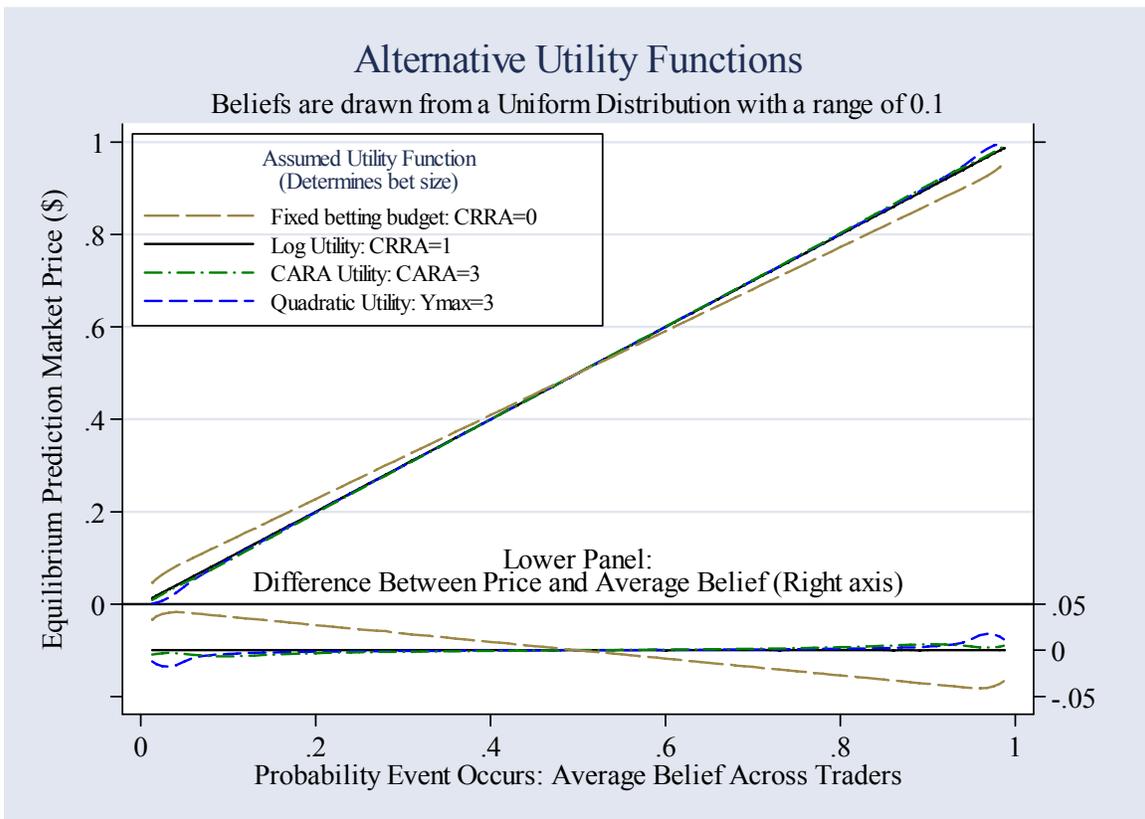


Figure 5: Distribution of Beliefs

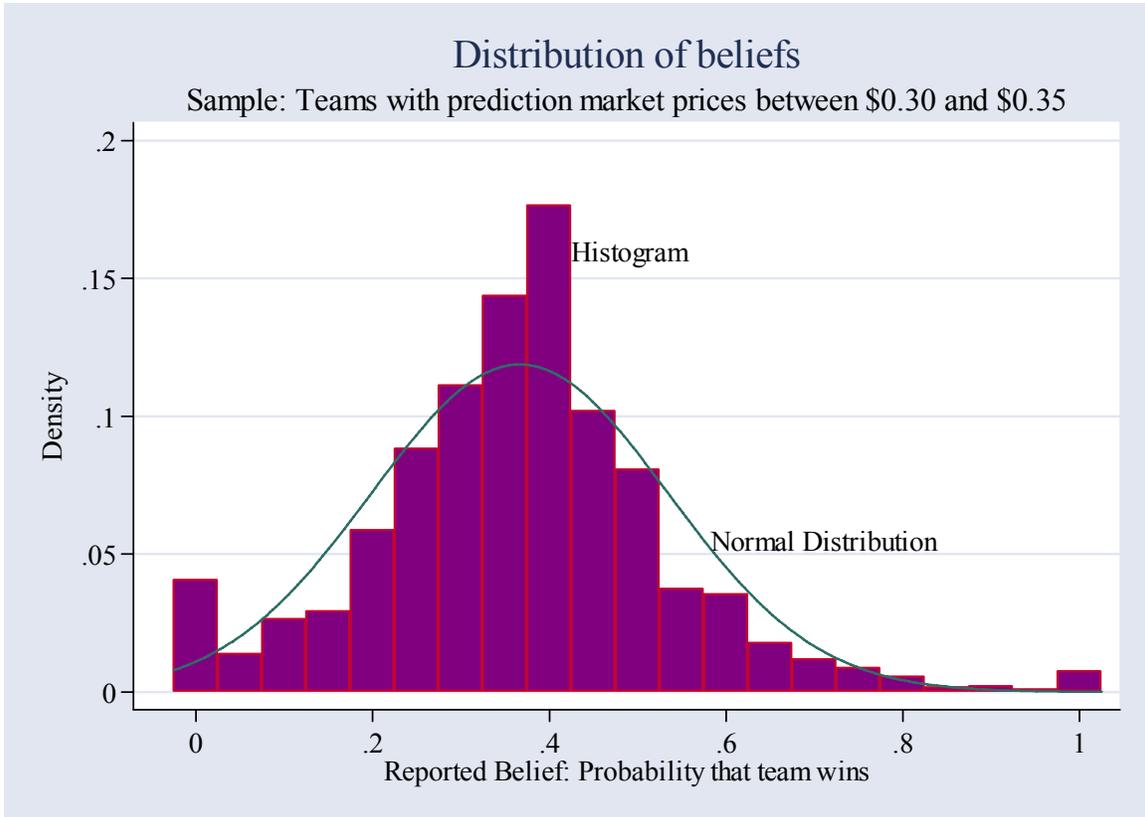


Figure 6: Distribution of Beliefs and Prediction Market Prices

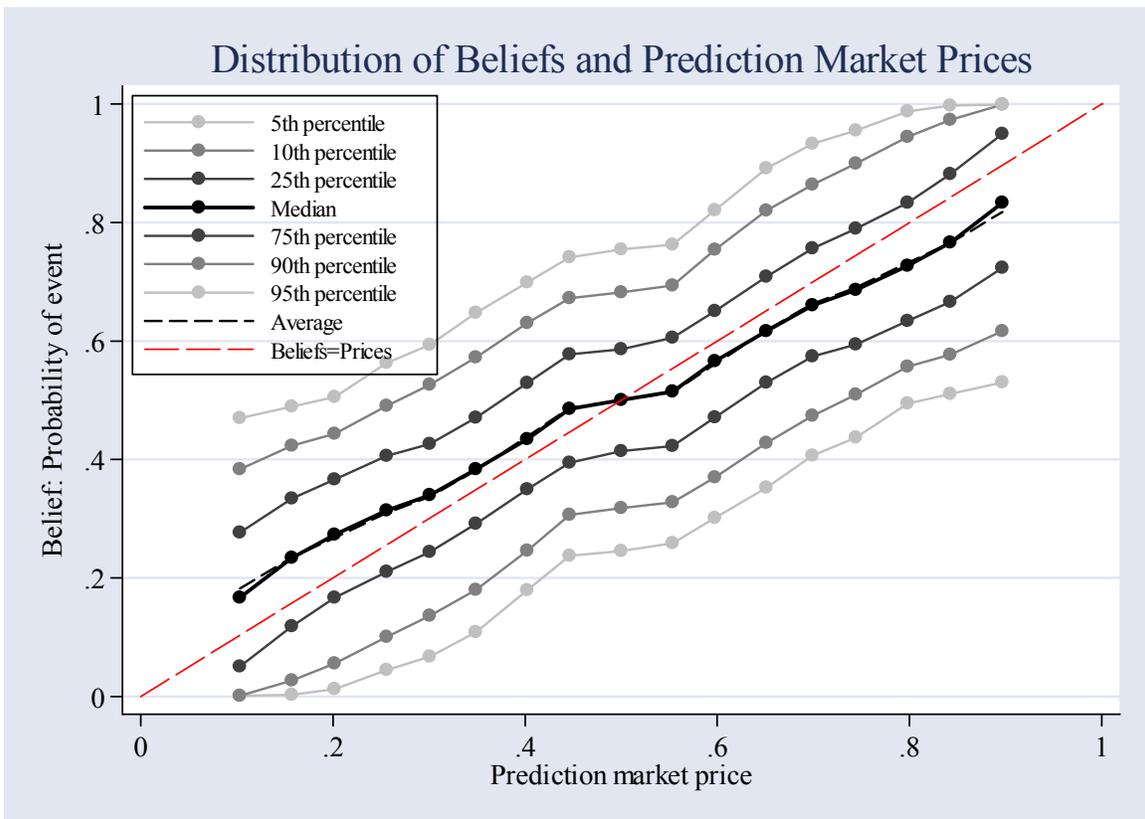


Figure 7: Predicting Prediction Market Prices: Different Models

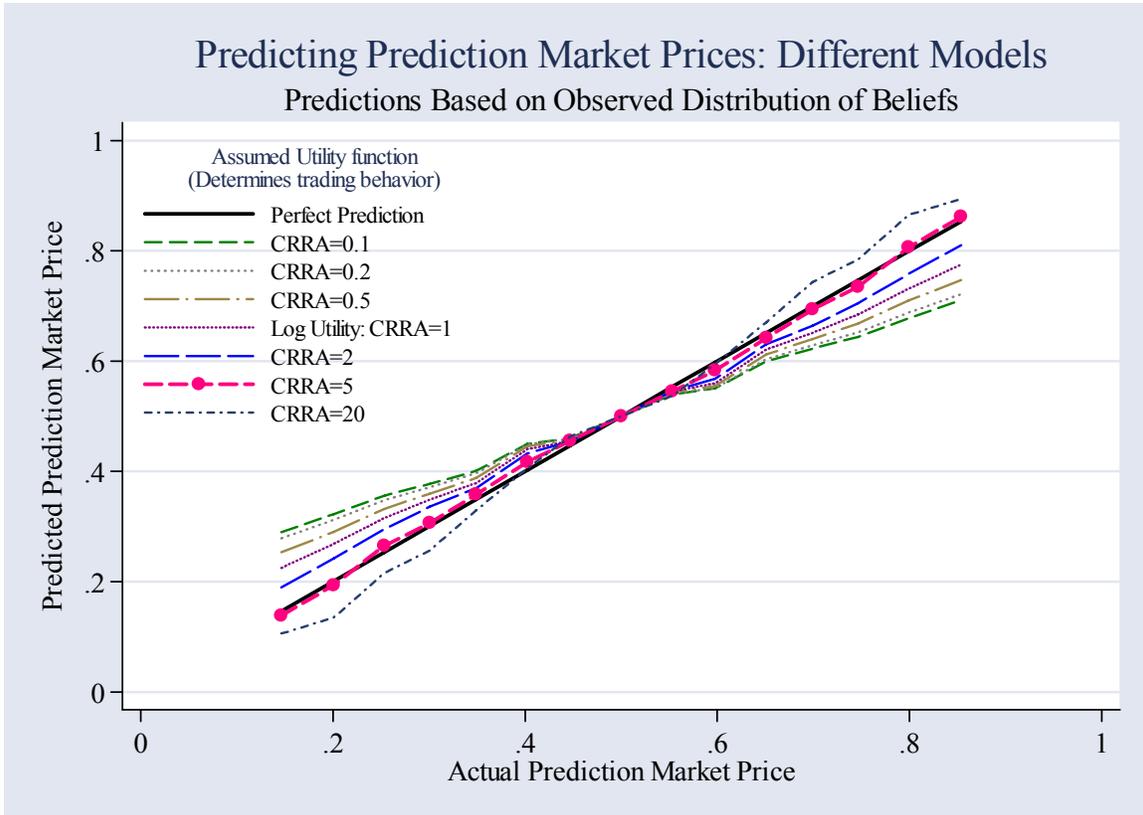


Figure 8: Probability of a 10 percent Stock Market Rise

