What Do Prediction Markets Predict?

by

John Fountain and Glenn W. Harrison[†]

January 2010

ABSTRACT. We show that prediction markets cannot be relied on to always elicit any interesting statistic of aggregate beliefs. Formal derivations of the bets placed in prediction markets can be viewed as demands for state-contingent commodities. We provide derivations for two popular cases, Log utility and Constant Relative Risk Aversion utility, connecting these derivations to familiar scoring rules. We then use these results to demonstrate how the properties of prediction markets depend critically on the assumed homogeneity of participants.

[†] Department of Economics, University of Canterbury, Christchurch, New Zealand (Fountain); and Department of Risk Management & Insurance and CEAR, Robinson College of Business, Georgia State University, USA (Harrison). E-mail contacts: john.fountain@canterbury.ac.nz and gharrison@gsu.edu. Harrison thank the U.S. National Science Foundation for research support under grants NSF/HSD 0527675 and NSF/SES 0616746.

Experimental economists were the first to develop prediction markets, originally defined over presidential elections in the United States: see Forsythe, Nelson, Neumann and Wright [1992]. There has been controversy over the claim that these markets elicit "aggregate beliefs," normally understood to mean the average of beliefs for the population.[1] There is no debate over whether these markets generate good predictions, in the sense that they forecast outcomes well.[2] Instead the debate has been over the additional claim that the prices in these markets recover the mean of aggregate beliefs. In other words, are the observed prices in these markets good estimates or predictors of aggregate beliefs?

Manski [2006] presented a simple, formal model in which a theoretical market did not elicit average beliefs. He built in asymmetry on the buying and selling side of the market in terms of point-mass beliefs, and assumed risk neutral agents with the same wealth (or a finite bet constraint, which amounts to the same thing in terms of the effect on market behavior).

Gjerstad [2005] and Wolfers and Zitzewitz [2005] present formal models in which they relax the assumption of risk neutrality, assume uni-modal beliefs, and find that markets can generate prices that reflect average beliefs.[3] However, their models build in homogeneity on both sides of the market, which is the key criticism of Manski [2006]. Every agent is risk averse, but has the same

---

[1] There is no claim that these markets elicit individual beliefs or that they are free of individual irrationalities: see Forsythe, Rietz and Ross [199] for a survey of field and laboratory evidence on this issue. The claim about aggregate beliefs rests on informal statements about the propensity of "marginal traders" to get market prices right (in some sense).

[2] A separate issue is the ability of prediction markets to reliably elicit beliefs in "informationally complex" environments: see Healy, Ledyard, Linardi and Lowery [2008]. As the state-space of events grows, markets defined over each possible event-combination would be expected to become thin, resulting in unreliable information aggregation. Another issue is whether prediction markets and hypothetical polls are forecasting the same thing: the former pay out depending on the actual outcome when it is realized, and the former are often posed counterfactually as eliciting beliefs if the outcome were to be realized "today." We have no interest in counterfactuals here.

[3] In private correspondence Gjerstad shows that *symmetric*, bi-modal beliefs will also lead to near-perfect aggregate belief elicitation in his model.

attitude to risk. And every agent has the same wealth level, or faces the same maximal bet constraint. So each agent chooses to bet the same amount, on both sides of the market, and one recovers a homogenous bidder model. One common feature of their models is that the Log utility specification emerges as the "poster boy" of aggregate belief elicitation: deviations from Log utility are associated with deviations from belief elicitation, although they each argue that these deviations are small for a wide range of belief distributions.

We consider the bets that agents with Log utility and Constant Relative Risk Averse (CRRA) utility would place, assuming that they are price takers in the prediction market.[4] We provide formal derivations of these bets, and then use these results to evaluate the claims about the beliefs elicited in prediction markets as one relaxes assumptions about homogeneity.

1. Formal Derivations

Bets placed in prediction markets, or stock markets for that matter, are simply purchases of state-contingent commodities. We provide formal derivations of compensated demands for state contingent commodities, $x^c(s)$, for two common utility functions. In each case we represent beliefs about state $s \in S$ by vectors $b(s)>0$, $\sum_s b(s)=1$ and market prices (reports) by vectors $p(s)>0$, $\sum_s p(s)=1$. A unit of s-contingent wealth pays \$1 if state s occurs and \$0 otherwise, for all states s=1,2....n. We also define $\check{s} \in S$, to aid the statement of summation operations.

A. Constant Relative Risk Aversion Utility

A risk averse CRRA agent has an EU defined over the state-contingent commodity $x^c(s)$ of

[4] Ottaviani and Sørensen [2007] consider the way in which trade in prediction markets might lead agents with heterogeneous prior beliefs to revise those beliefs in a rational expectations equilibrium.

the form

$$EU = [\sum_s b(s)\ x(s)^{(1-r)}]/(1-r) \tag{1}$$

The corresponding certainty equivalent (CE) wealth is a weighted mean wealth of order 1-r:

$$CE = [\sum_s b(s)\ x(s)^{(1-r)/r}]^{1/(1-r)} \tag{2}$$

This expression for the CRRA agent's CE is in standard CES form, where the elasticity of substitution $\eta = 1/r$ is the reciprocal of the Arrow-Pratt relative risk aversion parameter r, and is a measure of risk tolerance.

Consider the CRRA agent's optimal choice problem as an expenditure minimization problem: minimize $m=\sum_s p(s)x(s)$ subject to a *certainty equivalent* constraint

$$CE^{1-r} = \sum_s b(s)\ x(s)^{(1-r)} \tag{3}$$

The first-order conditions are, for an appropriate Lagrange multiplier λ,

$$p(s) = \lambda(1-r)\ b(s)\ (x(s)^{-r}) \tag{4}$$

We may identify λ by using (4) with the constraint (3), from the following expression:

$$CE^r = \lambda(1-r)\ \{ \sum_{\check{s}} b(\check{s})\ [b(\check{s})/p(\check{s})]^{(1-r)/r} \}^{r/(1-r)} \tag{5}$$

Substituting (5) into (4) yields Hicksian compensated demand functions $x^c(s)$ for wealth in state s:

$$x^c(s) = CE \times \{ [b(s)/(p(s)] / [\sum_{\check{s}} [b(\check{s})/x(\check{s})]^{(1-r)/r}]^{r/(1-r)} \}^{1/r} \tag{6}$$

Multiplying each $x^c(s)$ in (6) by p(s) and summing over all states yields the expenditure function $m=\sum_s p(s)x^c(s)$, or

$$m = CE / \{ \sum_{\check{s}} b(\check{s})\ [b(\check{s})/p(\check{s})]^{(1-r)/r} \}^{r/(1-r)} \tag{7}$$

Substituting (7) into (6) yields the *ordinary demand functions*

$$x^o(s) = \{ m \times [b(s)/(p(s)]^{1/r} \} / \{ \sum_{\check{s}} b(\check{s})\ [b(\check{s})/p(\check{s})]^{(1-r)/r} \} \tag{8}$$

The denominator of (7) is a risk-attitude-adjusted expected rate of return, a power mean of order (1-r)/r of individual contingent asset expected rates of return b(s)/p(s) for this CRRA agent.

B. Log Utility

A Log EU function, based on an increasing, concave utility function u(x)=ln(x), has form:

$$EU = \sum_{s} b(s) \ \ln(x(s)) \tag{9}$$

The corresponding CE is defined by

$$EU = \ln(CE) \tag{10}$$

From the first order conditions for expenditure minimization subject to an EU constraint we have:

$$p(s) = \lambda \, b(s) \, (1/x(s)) \tag{11}$$

for an appropriate Lagrange multiplier λ. Solving for $x^c(s)$ we obtain

$$x^c(s) = \lambda \, [\, b(s)/p(s) \,] \tag{12}$$

To eliminate λ, take the log of both sides of (12) and substitute back into (9) :

$$\ln(CE) = \ln(\lambda) + \sum_{s} b(s) \ln[\, (b(s))/p(s) \,] \tag{13}$$

Exponentiating both sides to eliminate λ, and substituting back into (12):

$$x^c(s) = CE \times \{ \, [\, b(s)/(p(s) \,] \, / \, [\, \prod_{\check{s}} [\, b(\check{s})/p(\check{s}) \,]^{b(\check{s})} \, \} \tag{14}$$

Multiplying (14) by p(s) and summing over all states yields

$$m = CE \, /[\, \prod_{\check{s}} [\, b(\check{s})/p(\check{s}) \,]^{b(\check{s})} \tag{14}$$

where the denominator is a weighted geometric mean of the individual contingent asset expected rates of return b(s)/p(s). Substituting (15) into (14) we obtain the ordinary demand functions

$$x^o(s) = m \times [\, b(s)/(p(s) \,] \tag{16}$$

2. Implications for Prediction Markets

It is easiest to demonstrate the implications of these formal derivations for the properties of prediction markets by simulation, simply because the objective is to incrementally relax analytically simplifying assumptions about homogeneity of participants.[5]

Consider the Log utility model initially. Assume 1,000 simulated agents with beliefs distributed according to some parametric form. Focus initially on a unimodal Normal distribution with mean 0.30 and standard deviation 0.10, and truncate at 0 the very few random draws that are negative. Let wealth be distributed uniformly between 1 and 100. Assume initially that there is no discounting, even though a key feature of prediction markets is that one places bets today for payoff in the future. We also allow for "free range betting," so that agents can rationally decide if they want to place a bet or not; in the absence of discounting they always want to place a bet, but when discounting enters they may not.[6]

Given their preferences, wealth, and beliefs, each agent evaluates their ordinary demand and decides on an optimal bet on the event or its complement. We undertake this evaluation for each agent, using (8) or (16), and then select the equilibrium price as the unique price at which demand equals supply. We evaluate all prices between 0.001 and 0.999 in increments of 0.001. Figure 1 shows the happy result that in this homogenous world prediction markets predict well and indeed recover the mean of aggregate beliefs perfectly.

Figure 2 relaxes some of these assumptions. The top left panel is a baseline, that repeats the assumptions from Figure 1 for reference. The top right panel allows for discounting behavior, and free-range betting. Thus agents can determine if the *present value* of the certainty-equivalent of their optimal bet is less than their current wealth, which is non-stochastic, and decide not to participate in the market. We select discount rates uniformly between 1% and 25%. Again we find that prediction

[5] The simulation is implemented using version 11 of *Stata*, and the command file is available on request.

[6] It is possible to use prediction markets in a setting in which there is "forced feed betting," where individuals have tokens to bet with and that are useless unless used to place a bet. Such settings arise in laboratory experiments and some corporate, in-house applications of predictions.

markets do just fine.

But then we introduce some bias, in the bottom left panel of Figure 2, in favor of the "optimistic" side of the market. We define the optimistic side as those agents that have beliefs which are greater than the average belief. In this case we allow them to have greater wealth, and to be more patient, than the baseline assumptions for them.[7] These assumptions move the equilibrium price in favor of their beliefs, and away from the average of beliefs. The bottom right panel combines discounting and bias, for the same qualitative result. The difference between equilibrium price and average belief is not quantitatively large, but it is clear.

Figure 3 repeats this exercise, but substituting CRRA utility for the log utility specification. We draw CRRA values from a uniform distribution defined over the open interval (0, 1), to correspond to modest amounts of risk aversion; log utility, of course, is the special case of CRRA when the CRRA parameter tends to 1, and risk neutrality is the case of the CRRA parameter being 0. The notion of bias is extended to also include a lower CRRA coefficient for those with optimistic beliefs, implying that they are less risk averse and hence more inclined to place a bet in support of their beliefs.[8] The results are virtually the same as for log utility in Figure 2, but the inaccuracy of the equilibrium price is increasing compared to the Log utility case.

Figure 4 repeats the simulation from Figure 2, using CRRA utility, but with an asymmetric, bi-modal distribution generated as a linear average of two normal distributions. We now see even larger differences between the equilibrium price and the average belief. Figure 5 then considers a completely diffuse distribution of beliefs, generated uniformly between 0.01 and 0.99, and shows a

---

[7] Specifically, the have 100 units of wealth more than the baseline, and their discount rates are cut in half.

[8] Specifically, that the CRRA coefficient is lower by 0.25, as long as this would not make them risk-loving.

similarly difference between equilibrium price and average belief.

3. Conclusion

Taking all these results together, and in conjunction with those of the earlier literature, we conclude that prediction markets can be expected to do a good job recovering the average of aggregate beliefs under certain circumstances: unimodal distributions of beliefs, with no *a priori* reason to expect heterogeneity on either side of the market. Indeed, this environment *might* characterize many interesting settings, such as political elections or closed prediction markets in which there is minimal sample selection into the market (on the basis of beliefs, preferences and endowments). But the result is not general, and it is easy to construct examples in which prediction markets do a predictably poor job of recovering average beliefs. This unfortunate result is not an artefact of assuming extreme heterogeneity, such as one might incorrectly conclude from the example of Manski [2006].

References

Forsythe, Robert; Nelson, Forrest; Neumann, George R., and Wright, Jack, "Anatomy of an Experimental Political Stock Market," *American Economic Review*, 82(5), December 1992, 1142-1161.

Forsythe, Robert; Rietz, Thomas A., and Ross, Thomas W., "Wishes, Expectations and Actions: A Survey on Price Formation in Election Stock Markets," *Journal of Economic Behavior & Organization*, 39, 1999, 83-110.

Gjerstad, Steven, "Risk Aversion, Beliefs, and Prediction Market Equilibrium," *Unpublished Manuscript*, Economic Science Laboratory, University of Arizona, January 2005.

Healy, Paul J.; Ledyard, John O.; Linardi, Sera, and Lowery, J. Richard, "Prediction Market Alternatives for Complex Environments," *Working Paper*, Department of Economics, Ohio State University, 2008.

Manski, Charles F., "Interpreting the Predictions of Prediction Markets," *Economics Letters*, 91(3), 2006, 425-429.

Ottaviani, Marco, and Sørensen, Peter Norman, "Aggregation of Information and Beliefs in Prediction Markets," *Working Paper*, London Business School, May 2007.

Wolfers, Justin, and Zitzewitz, Eric, "Interpreting Prediction Market Prices as Probabilities," *Unpublished Manuscript*, The Wharton School, University of Pennsylvania, February 2005.

Figure 1: Log Utility Agents in a Homogeneous Prediction Market

Solid vertical line is mean of true beliefs. Dashed vertical line is equilibrium price.
Simulations with 1000 agents and assuming free range betting.
Fraction of agents with CRRA utility is 0; fraction with log utility is 1.
Those with optimistic beliefs are similar to those with pessimistic beliefs.

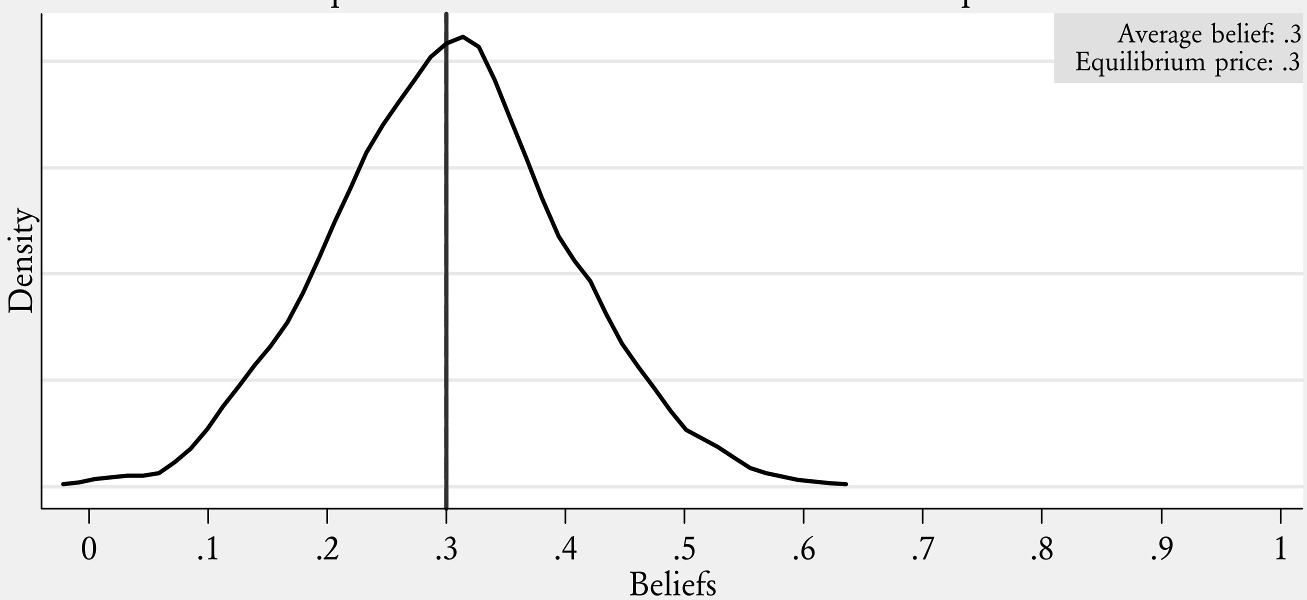

Parameterization for all agents: wealth distributed uniformly between 1 and 100.
risk attitudes derived from log utility function;
no discounting of future payoffs assumed.

Figure 2: Prediction Markets with Log Utility

Solid vertical line is mean of true beliefs. Dashed vertical line is equilibrium price.

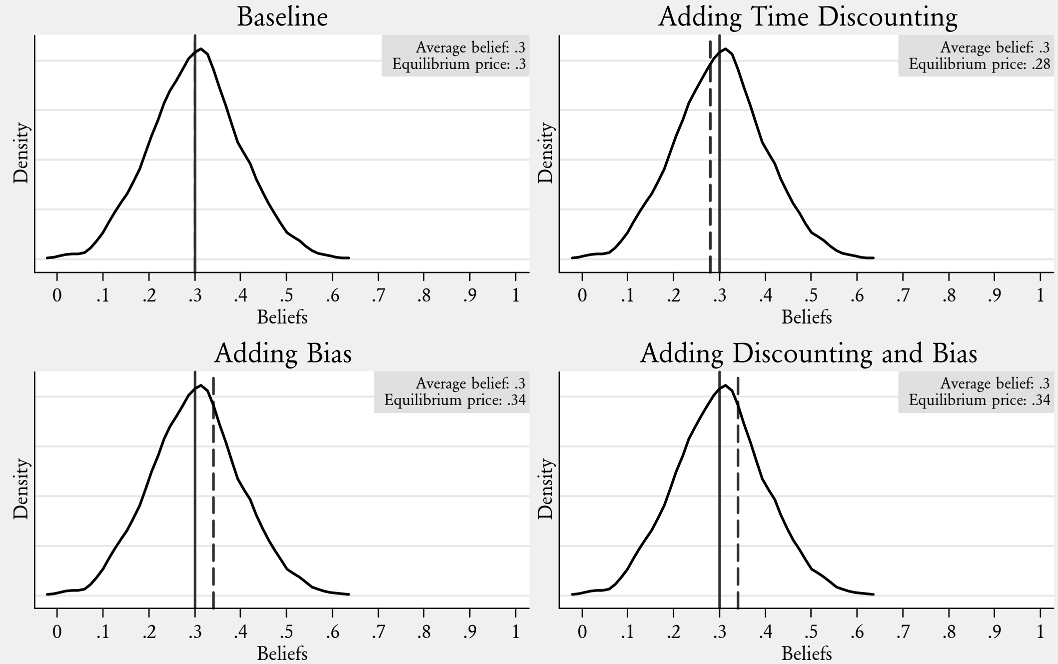

Figure 3: Prediction Markets with CRRA Utility

Solid vertical line is mean of true beliefs. Dashed vertical line is equilibrium price.

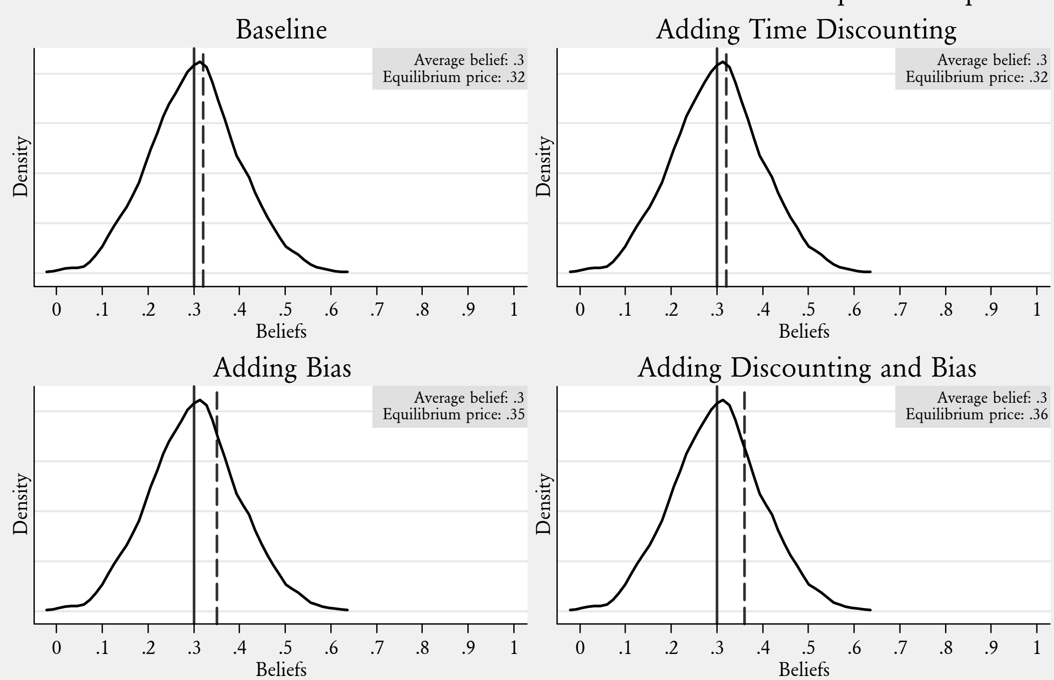

Figure 4: Prediction Markets with CRRA Utility
And Asymmetric, Bi-Modal Aggregate Beliefs

Solid vertical line is mean of true beliefs. Dashed vertical line is equilibrium price.

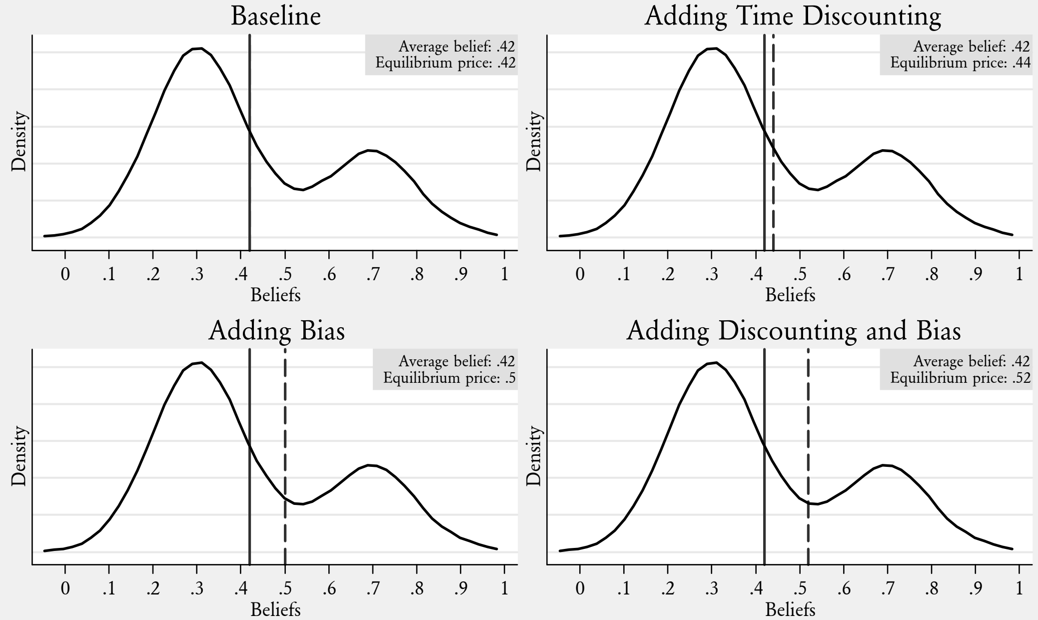

Figure 5: Prediction Markets with CRRA Utility
And Diffuse Aggregate Beliefs

Solid vertical line is mean of true beliefs. Dashed vertical line is equilibrium price.