UNIVERSITY OF MANNHEIM

University of Mannheim / Department of Economics

Working Paper Series

Why prediction markets work:
The role of information acquisition and endogenous weighting

Christoph Siemroth

Working Paper 14-02

Januar 2014

Why prediction markets work:
The role of information acquisition and endogenous weighting*

Christoph Siemroth†
University of Mannheim

January 10, 2014

Abstract

In prediction markets, investors trade assets whose values are contingent on the occurrence of future events, like election outcomes. Prediction market prices have been shown to be consistently accurate forecasts of these outcomes, but we don't know why. I formally illustrate an information acquisition explanation. Traders with more wealth to invest have stronger incentives to acquire information about the outcome, thus tend to have better forecasts. Moreover, their trades have larger weight in the market. The interaction implies that a few well-endowed traders can move the asset price toward the true value. One implication for institutions aggregating information is to put more weight on votes of agents with larger stakes, which improves on equal weighting, unless prior distribution accuracy and stakes are negatively related.

Keywords: Information Acquisition, Information Aggregation, Forecasting, Futures Markets, Prediction Markets.
JEL Classification: D83, D84, G10.

*I am grateful to Pierre Boyer, Hans Peter Grüner, Felix Jarman, Konrad Stahl, Philipp Zahn, and seminar participants in Mannheim for discussion and very helpful comments. This research was supported by the German Science Foundation (DFG) via SFB 884.

†Department of Economics, L7, 3-5, 68131 Mannheim, Germany. E-Mail: christoph.siemroth@gess.uni-mannheim.de.

1 Introduction

In a 2003 forecasting tournament, participants predicted outcomes of football games throughout a season to win prizes. Probability forecasts were rated with a quadratic scoring rule, so only participants with consistently accurate forecasts would be in the top ranks. Two mock entrants simply used the prices from two different prediction markets as their forecasts, and placed 6th and 8th out of almost 2,000 participants (Servan-Schreiber et al., 2004). More generally, prediction markets have been shown to provide better forecasts than polls in political elections (e.g., Forsythe et al., 1992; Berg et al., 2008), expert forecasts in sports (Spann and Skiera, 2009), or sales forecasts in business (Wolfers and Zitzewitz, 2004).

One kind of asset traded in these markets is called winner-take-all (WTA) contract. It pays out 1 iff a pre-specified condition is fulfilled, otherwise it pays out 0. For example, the IEM prediction market for the 2012 US Presidential election traded a Democratic and a Republican contract, which would pay out $1 iff the respective candidate obtained the majority of popular votes cast for the two major parties. Consequently, the price of a WTA contract may be interpreted as the market probability estimate that the respective candidate wins the election. With similar contracts, market-based predictions can be obtained for virtually all areas beyond politics.

Why do these markets predict so accurately? There are no satisfying explanations so far. As Berg and Rietz (2006) state, "exactly how prediction markets become efficient is something of a mystery." The main goal of this paper is to provide and formally illustrate a theory. In what I shall call information acquisition explanation, traders have stronger incentives to acquire information about the unknown outcome the larger their endowment. Consequently, high endowment traders are better informed. Moreover, high endowment traders have larger impact on the market price, because they can buy more assets. This interaction implies that few, but well-endowed traders can move the market price—interpreted as prediction—in the right direction, thereby explaining the observed accuracy. Unlike many financial market models, the explanation does not rely on the existence of rational expectations, nor on the presence of insiders or the ability of traders to infer information from asset prices. Even markets with traders who have systematically biased opinions about the outcomes can produce accurate forecasts, because of effective incentives for information acquisition and endogenous weighting by investment volume.

In my model, traders start out with an initial opinion about the outcome of the election. Based on this prior belief, endowment and asset prices, they decide whether to acquire information, whose accuracy depends on their information acquisition effort. Consequently, informed traders and noise traders (driven by opinion) evolve endogenously, which explains partially where beliefs originate and when beliefs (and market forecasts) tend to be accurate.

I establish that traders with prior beliefs close to the 'market estimate' (price), or with high endowment, have the strongest incentive to acquire information. The interpretation is that traders with extreme opinions about the outcome do not expect to be swayed by evidence, and hence do not acquire it, while traders with opinions close to the market price acquire information, because it might change their investment decision. High endowment traders have more at stake and are willing to acquire costly information in order to safeguard their investment. Comparative statics show that the bias of the market price is usually reduced in response to an endowment increase of all traders, because information acquisition is supported. It also shows that a shift of prior beliefs toward the true outcome usually improves forecasts, but may in rare cases *increase* bias. Numerical examples give an idea about the effect size of information acquisition and exogenous changes in the parameters.

One lesson for institutional design is that giving more weight to votes or investments of high endowment agents might improve information aggregation. If accuracy of prior beliefs and endowment are not negatively related, then high endowment agents tend to have better information, which can be exploited via weighting. An empirically testable implication is that forecasts should be better with weighting by stake rather than equal weighting. Moreover, forecasts should be better if investment per trader is larger in the market.

WTA prediction markets are analytically identical to fixed-odds betting, provided odds are set competitively. Prediction market prices are translated into odds like probabilities. For example, if the price of the Democratic contract is π, and the complement is priced at $1 - \pi$, then the odds of a Democratic victory are $(1 - \pi)/\pi$ in the corresponding betting market (ignoring fees). Thus, the results also apply to betting markets.

While the idea that prediction (or stock) markets provide incentives to search for information is not new (e.g., Servan-Schreiber et al., 2004; Wolfers and Zitzewitz, 2004; Arrow et al., 2008), this is the first paper to formalize it in order to explain prediction market accuracy, and demonstrate its interaction with the endogenous weighting implied by market clearing. Existing models are designed to address different questions (see next section). They typically assume an arbitrary belief distribution or give traders informative signals by default, so that forecast accuracy is a trivial consequence of the accuracy of these primitives. In contrast, quality of information is endogenous in my model, so it is more suitable to explain the accuracy or inaccuracy of prediction markets. In a numerical example, the model with information acquisition has a 10 percentage points smaller forecast bias on average than a model without information acquisition (e.g., Manski, 2006), holding all exogenous parameters constant. Moreover, the model motivates a different view on the interpretation of prediction market prices (e.g., Manski, 2006, Wolfers and Zitzewitz, 2006). Instead of comparing prices to statistics of the belief distribution, the key question is whether beliefs are driven by information rather than opinion (section 3).

1.1 Related literature

I confine attention to the relevant theoretical contributions in this section.[1] An informal explanation of prediction market accuracy was put forth by Forsythe et al. (1992). In what they dub marginal trader hypothesis (MTH), they argue that "prices are determined by the marginal trader" and that marginal traders are "free of judgment bias" (i.e., have accurate beliefs about the outcome). According to their definition (p. 1158 and fn. 21), a marginal trader is a trader who places limit orders within 2 cents of the current price. However, they are vague on how marginal traders "set prices." Why should they have the power to set prices, while biased traders—who may be as convinced of the correctness of their beliefs as unbiased traders—do not? Indeed, the existence of a group with perfect forecast is not sufficient for an accurate market forecast if wealth is bounded. A small subset of informed traders cannot always outbid hordes of biased traders to keep the price at the fair value. Nor is their presence necessary—possible biases of Republicans and Democrats may cancel out to accurately predict the election vote share. Consequently, the MTH does not explain the consistently good performance of prediction markets.

My information acquisition explanation has similarities to the MTH, in that both imply that a small group of well informed traders can influence the market price to the better. Interestingly, Forsythe et al. (1992) find that marginal traders (with good forecasts about the outcome) have higher investment than the rest, which is consistent with the information acquisition explanation. However, the information acquisition explanation differs in that informed traders need not be free from judgment biases, need not have the power to set prices, and need not have beliefs equal to what is implied by the market price.

Instead of explaining predictive accuracy, most of the theoretical literature on betting markets tries to explain the favorite long-shot bias, i.e., that favorites' chances are often underestimated, while long-shots' chances are overestimated. Explanations for the bias go back to at least Ali (1977) for horse race betting. He shows that market probabilities are less extreme than objective probabilities in a market with risk neutral bettors, whose median belief is equal to the objective probability. For parimutuel betting, Ottaviani and Sørensen (2009) show that inaccurate odds until just prior to the "last call" can result from bettors without rational expectations, i.e., bettors do not anticipate the final odds and make no inferences about the true state of the world from them. In a model with heterogeneous priors, Ottaviani and Sørensen (2013) show that the competitive equilibrium price under-reacts to information, and the effect is exacerbated for more spread out priors. Page and Clemen (2013) argue theoretically and empirically that prediction markets, while "reasonably well calibrated" for predictions in the near future, can be less accurate if the outcome to be predicted is far in the future. According to them, the loss of accuracy

[1] A good introduction to prediction markets with examples is provided by Wolfers and Zitzewitz (2004). An overview of the large betting literature in economics can be found in Sauer (1998), and Thaler and Ziemba (1988) provide an introduction to empirical anomalies in betting markets. Tziralis and Tatsiopoulos (2007) give an extensive overview of the prediction market literature with categorization into subfields.

is explained by reduced trading activity due to discounting of future gains. For a more complete overview of explanations for the favorite long-shot bias, consult the references in Ottaviani and Sørensen (2010). In all of these models, the origin of beliefs is either unmodeled or informative signals are obtained by default with exogenous precision.

Another strand of literature focuses on how prediction market prices can be interpreted (cf. Wolfers and Zitzewitz, 2004 for an overview of the different contracts traded in prediction markets). Manski (2006) shows that the equilibrium price in a prediction market trading WTA contracts with risk neutral traders does not generally correspond to the mean (or median) probability estimate among traders, as conventional wisdom holds. In response, Gjerstad (2005) and Wolfers and Zitzewitz (2006) show that the prediction market price is close to the mean belief with risk averse traders. The model presented here proposes another view on the issue (section 3). The question is not whether the price is close to the mean or median belief, but whether it is close to the actual outcome. A distribution of opinions may or may not have the correct mean or median. Instead, the key is whether beliefs and consequently prices are informed or mere opinion, which can only be answered by specifying how beliefs are formed. Endogenous information acquisition is one way to achieve this.

To my knowledge, information acquisition has been incorporated only once in the context of prediction markets. Hanson and Oprea (2009) investigate whether a manipulator can drive the price away from the fundamental value of the asset. Since his presence as well as the strength of the manipulation preference is common knowledge, informed traders may react to the manipulation attempt by obtaining more precise signals, thus raising prediction market accuracy on average. In contrast to my model, they use the quantal response equilibrium concept, all distributions are assumed to be common knowledge, all random variables are normally distributed, and there are no budget constraints.

Among the first to consider information acquisition in economics more generally, Grossman and Stiglitz (1980) show that a fully revealing rational expectations equilibrium with costly information acquisition does not exist. The following literature (e.g., Verrecchia, 1982; Barlevy and Veronesi, 2000; Peress, 2004) focuses on noisy rational expectations equilibria, where the price is affected by noise and only partially revealing to retain the incentive for information acquisition. In contrast to these papers, my model imposes no common knowledge assumptions except for the observability of the price.

2 A model of costly information acquisition

An unknown state of the world (briefly 'outcome') from the set $\{A, B\}$, $B = A^c$, is exogenously given, and will be publicly revealed in the future. For example, suppose a presidential candidate (incumbent) faces a challenger in an upcoming election. Then A represents victory by the incumbent in the election, whereas B means the challenger is victorious. Formally, $\theta = 1$ iff A and $\theta = 0$ iff B, which is the parameter to be predicted.

The economy is populated by a continuum of risk-neutral traders.[2] Traders may be heterogeneous in their endowment $\omega_i \in (0, \overline{\omega}]$, $\overline{\omega} < \infty$, which is distributed with cdf $W(\omega_i)$. Moreover, each trader i is characterized by prior q_i, drawn from a continuous cdf $Q(q_i|\theta, \omega_i)$, which is i's subjective estimate of $\Pr(\theta = 1)$. This is a deviation from the common prior assumption imposed in the better part of the literature, and implies that traders may disagree about prospects. Diversity of opinion allows for interpretations like wishful thinking, where opinions are influenced by preferences (Forsythe et al., 1999). It also makes the explanation stronger, because in spite of many traders having wrong priors about the outcome, the market price may nevertheless be an accurate forecast. The specification of prior belief distribution Q allows for a dependence with endowment ω_i.

Traders do not receive an informative signal about θ by default. Instead, they may acquire a private binary signal, which is costly in terms of effort. The precision of the signal is a function of effort $e_i \geq 0$, $\nu(e_i) = \Pr(s_i = 1|\theta = 1) = \Pr(s_i = 0|\theta = 0)$. The interpretation is that a trader can run Internet searches or talk to experts (signal), which influence his beliefs about the outcome (posterior). But the effort cost may be too high, so a trader may rather rely on his opinion (prior) to make the investment decision. Effort costs are not paid out of the endowment, but enter linearly in the utility function.

Two state-contingent futures contracts—also called winner-take-all contracts—are traded in a prediction market. One A-contract pays 1 iff $\theta = 1$ to the holder, and 0 otherwise. Conversely, one B-contract pays 1 iff $\theta = 0$. The contracts are issued by the market maker. The prediction market is thus a complete one-period Arrow-Debreu security market, like the IEM prediction market described in the previous section. Let the price of the A-contract be π, and the price of the B-contract $1 - \pi$, to rule out arbitrage opportunities. The profit per A-contract held if A occurs is the difference of 1 and the price π; the loss per A-contract if B occurs is π, the price paid.

The timing of trader decisions is shown in Figure 1. Each trader first decides on information acquisition effort e_i, or equivalently signal precision $\nu_i := \nu(e_i)$, for a given price π ($t = 0$). Then he receives a signal with precision ν_i, and computes posterior $p_i := \Pr(\theta = 1|s_i)$ using Bayes' rule. The posterior is equal to the prior if zero effort, i.e., $\nu_i = 1/2$, is chosen. Based on posterior p_i, endowment ω_i, and price π, the trader decides which option to invest in by specifying investment volume $a_i(p_i, \omega_i, \pi), b_i(p_i, \omega_i, \pi)$ ($t = 1$). Finally, the outcome θ is revealed and the assets pay out ($t = 2$).

The equilibrium price π^* can be viewed as the market's probability estimate that A occurs. The more traders believe that A is going to occur, the more invest in A, thus raising

[2]As an alternative to risk neutrality, it is not clear whether agents should be modeled as risk averse or risk loving in the context of betting. Almost all papers estimating risk preferences from betting data find that traders are risk loving (e.g., Jullien and Salanié, 2000 or Snowberg and Wolfers, 2010), whereas people exhibit risk aversion in most other contexts. Risk neutrality, then, is a robust compromise. The aggregate behavior induced by risk neutrality has the desirable trait of strictly decreasing aggregate demand if beliefs are unchanged, which is not guaranteed under strong risk aversion (e.g., Quah, 2003). The qualitative results carry over for slight risk aversion or risk loving preferences.

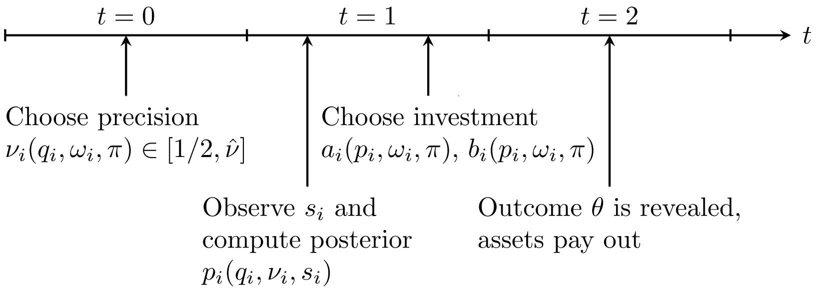

Figure 1: Timing of decisions for trader i and payout.

the price (forecast) π^*. The bias of the market forecast is $|\pi^* - \theta|$, so the best forecast is for the market price to equal the fundamental value of the asset (i.e., 1 or 0). In competitive equilibrium, the price equates the aggregate demand for A and B contracts, so that money is redistributed from losers to winners, and the market operates at zero profit. Modeling the market with the possibility to buy either of two complementary contracts is identical to having just one contract with supply and demand side.

The competitive equilibrium concept used here requires that price π^* induces information acquisition, which leads to beliefs and asset demands clearing the market for that same price.

Definition 1. *A competitive equilibrium with endogenous information acquisition requires*

1. *a precision level function $\nu_i(q_i, \omega_i, \pi)$ for all i, which maximizes expected utility at $t = 0$ anticipating optimal behavior at $t = 1$,*

2. *posterior beliefs $p_i(q_i, \nu_i, s_i)$ for all i, computed via Bayes' rule,*

3. *investment functions $a_i(p_i, \omega_i, \pi)$, $b_i(p_i, \omega_i, \pi)$ for all i, which maximize expected utility subject to $a_i + b_i \leq \omega_i$ at $t = 1$, and*

4. *an equilibrium price π^*, which induces information acquisition $\nu_i(q_i, \omega_i, \pi^*)$ at $t = 0$, leading to beliefs $p_i(q_i, \nu_i, s_i)$ and investments $a_i(p_i, \omega_i, \pi^*), b_i(p_i, \omega_i, \pi^*)$ clearing the asset market at $t = 1$, i.e.,*

$$\int_0^{\overline{\omega}} \int_0^1 a_i(p_i, \omega_i, \pi^*)/\pi^* dQ(q_i|\theta, \omega_i) dW(\omega_i) = \int_0^{\overline{\omega}} \int_0^1 b_i(p_i, \omega_i, \pi^*)/(1-\pi^*) dQ(q_i|\theta, \omega_i) dW(\omega_i).$$

The difference to an Arrow-Debreu equilibrium is that the equilibrium price must simultaneously induce information acquisition levels ν_i and clear the market for the resulting investment functions a_i, b_i. Consequently, the equilibrium concept yields a single equilibrium price, even though in principle the sequential decisions of the traders (Figure 1) allow for different prices at $t = 0$ and $t = 1$. To motivate this equilibrium notion, suppose a Walrasian auctioneer announces an initial price π^0 at $t = 0$. Traders make their information acquisition decision, update their beliefs and trade at $t = 1$, which leads to market clearing price $\pi^1 \neq \pi^0$ if π^0 does not fulfill the above definition. Thus, if we were to repeat the procedure with π^1 as initial price, traders might make different information acquisition decisions and ultimately different investment decisions, possibly leading to yet another price. If this tâtonnement procedure stops for some $\pi^0 = \pi^1$, then $\pi^0 = \pi^*$ as defined above.

Many financial market models use rational expectations equilibria to explain the informativeness of asset prices. The involved assumptions are strong. As Radner (1979) writes, traders require "'models' or 'expectations' of how equilibrium prices are related to initial information," in order to infer and pool information from the observable price. In contrast, this model provides an explanation that does not rely on the strict common knowledge assumptions of rational expectations equilibria. As in actual prediction markets, traders do not know the beliefs or endowments of other traders, and do not know if trades are motivated by hedging or manipulation motives, i.e., they do not have the necessary 'models' to extract information from the price. But the price may nonetheless be correlated with the true state of the world, because signals are informative. The explanation of why the market price can be a good forecast relies on the realistic assumptions of endowment asymmetries, endogenous information acquisition, and a competitive equilibrium, which approximates the double auction mechanism used in asset markets.

2.1 Investment and information acquisition decision

In this section, I determine the optimal individual information acquisition and investment decisions for a given price π. In the next section, individual decisions will be aggregated to determine the equilibrium price π^*. Going backwards on the time line, taking posterior p_i and price π as given, the investment problem of the risk neutral trader at $t = 1$ is

$$\max_{a_i, b_i \geq 0} p_i[(1-\pi)a_i/\pi - b_i] + (1-p_i)[\pi b_i/(1-\pi) - a_i] \text{ s.t. } a_i + b_i \leq \omega_i.$$

That is, the trader believes A occurs with probability p_i, yielding a profit of $(1-\pi)a_i/\pi$ on the a_i investment, and a loss of all B investment b_i. Payoffs for B follow similarly. The linear utility function yields a corner solution, which is $a_i = \omega_i,\ b_i = 0$ if $p_i > \pi$ and $a_i = 0,\ b_i = \omega_i$ if $p_i < \pi$. Henceforth, I will use the short-hand $\alpha_i(\pi) := \omega_i/\pi$ and $\beta_i(\pi) := \omega_i/(1-\pi)$ to denote the maximum amount of A or B-contracts bought, respectively.

Anticipating these investment decisions, the trader decides how much costly effort to spend, which determines the precision of the signal. Note that the effort choice at $t = 0$ has to induce investment behavior which depends on the signal at $t = 1$ ('discriminating effort level'), otherwise the effort cost is incurred for no benefit. For example, if the resulting posterior is $p_i(s_i = 0, q_i, e_i) < p_i(s_i = 1, q_i, e_i) < \pi$, then exerting effort $e_i > 0$ would not make a difference in the investment decision—for either realization of the signal i invests in

B—and therefore cannot be optimal. The minimum discriminating effort level is

$$\tilde{e}_i = \min_e \{e : p_i(s_i = 1, q_i, e) \geq \pi \geq p_i(s_i = 0, q_i, e)\}.$$

The expected utility from choosing a discriminating signal precision *before* observing the signal ($t = 0$) is

$$EU(e_i \geq \tilde{e}_i) = q_i \nu(e_i)(1-\pi)\alpha_i - q_i(1-\nu(e_i))\omega_i + (1-q_i)\nu(e_i)\pi\beta_i - (1-q_i)(1-\nu(e_i))\omega_i - e_i.$$

That is, from his prior perspective, trader i anticipates that he will invest in A iff $s_i = 1$ and in B iff $s_i = 0$, that the signal will be correct with probability $\nu(e_i)$, and wrong with probability $1 - \nu(e_i)$, and he weighs each case according to his prior belief q_i.

From the definition of $\tilde{e}_i$, if $q_i > \pi$, then $p_i(s_i = 1, q_i, \tilde{e}_i) > \pi$, consequently $p_i(s_i = 0, q_i, \tilde{e}_i) = \pi$, which implies the minimum discriminating precision level

$$\tilde{\nu}_i := \nu(\tilde{e}_i) = \frac{q_i/\pi}{q_i/\pi + (1-q_i)/(1-\pi)}.$$

Similarly, if $q_i < \pi$, then $p_i(s_i = 1, q_i, \tilde{e}_i) = \pi$ and $p_i(s_i = 0, q_i, \tilde{e}_i) < \pi$, which implies

$$\tilde{\nu}_i = \frac{(1-q_i)/(1-\pi)}{q_i/\pi + (1-q_i)/(1-\pi)}.$$

The expected utility of not acquiring information is

$$EU(e_i = 0) = \begin{cases} q_i(1-\pi)\alpha_i - (1-q_i)\pi\alpha_i = (q_i - \pi)\alpha_i & \text{if } q_i > \pi, \\ (1-q_i)\pi\beta_i - q_i(1-\pi)\beta_i = (\pi - q_i)\beta_i & \text{if } q_i < \pi. \end{cases} \quad (1)$$

A trader prefers information acquisition, i.e., a positive effort level, only if the benefits from the more informed investment decision exceed the effort costs. Hence, positive effort, assuming $\nu_i := \nu(e_i) \geq \tilde{\nu}_i$, is incentive compatible iff for $q_i > \pi$

$$q_i\nu_i(1-\pi)\alpha_i - q_i(1-\nu_i)\pi\alpha_i + (1-q_i)\nu_i\pi\beta_i - (1-q_i)(1-\pi)(1-\nu_i)\beta_i - e_i \geq (q_i - \pi)\alpha_i. \quad (2)$$

Solving the first order condition ($\nu_{ee} < 0$) for the LHS, the unconstrained optimal positive effort level is

$$e_i^* = \nu_e^{-1}\left(\frac{1}{q_i\alpha_i + (1-q_i)\beta_i}\right) > 0,$$

where ν_e is the partial derivative with respect to effort. The optimal level exists under the Inada conditions $\nu_e(e) \to 0$ as $e \to \infty$ and $\nu_e(e) \to \infty$ as $e \to 0$.

For explicit solutions, I assume a specific form for the effort-precision function, $\nu(e) = \min\{\frac{1}{2}(\sqrt{e} + 1), \hat{\nu}\}$, $\hat{\nu} < 1$, so that $\nu(e = 0)$ is normalized to 1/2 (an uninformative signal), and a unique unconstrained optimum exists. Because the square root function is unbounded, I impose an upper bound $\hat{\nu} < 1$, so that traders cannot learn the true state of the world perfectly by investing a lot of effort. In this specification, the optimal unconstrained precision is

$$\nu_i^* := \nu(e_i^*) = \frac{q_i\alpha_i + (1-q_i)\beta_i + 4}{8}.$$

Formally, a marginal precision increase makes a correct investment (i.e., profits) in either state of the world more likely, and a wrong investment (i.e., losses) less likely. Adding the benefits of more likely gains and less likely losses, weighted by prior probability, gives the marginal benefit $q_i\alpha_i + (1-q)\beta_i$. The marginal effort cost is $8\nu - 4$, and the optimal precision equates the two.

The upper and lower bound of incentive compatible precisions $\underline{\nu}_i := \underline{\nu}(q_i, \omega_i, \pi)$ and $\overline{\nu}_i := \overline{\nu}(q_i, \omega_i, \pi)$ if $q_i > \pi$ and $\nu_i \geq \tilde{\nu}_i$ are the solutions to the quadratic equation (2),

$$(\underline{\nu}_i, \overline{\nu}_i) = \frac{q_i\alpha_i + (1-q_i)\beta_i + 4}{8} \pm \sqrt{\frac{(q_i\alpha_i + (1-q_i)\beta_i + 4)^2}{64} - 1/4(1 + q_i\alpha_i)}. \tag{3}$$

For $q_i < \pi$, the RHS (expected utility without information acquisition) of (2) changes, resulting in solutions

$$(\underline{\nu}_i, \overline{\nu}_i) = \frac{q_i\alpha_i + (1-q_i)\beta_i + 4}{8} \pm \sqrt{\frac{(q_i\alpha_i + (1-q_i)\beta_i + 4)^2}{64} - 1/4(1 + (1-q_i)\beta_i)}. \tag{4}$$

In short, i acquires information iff $[\underline{\nu}_i, \overline{\nu}_i] \cap [\tilde{\nu}_i, \hat{\nu}] \neq \varnothing$, and chooses ν_i^* if it is in the intersection. The following proposition characterizes the information acquisition decision.

Proposition 1.

i. Incentive compatible precision levels always discriminate, i.e., $\underline{\nu}(q_i, \omega_i, \pi) \in \mathbb{R} \implies \underline{\nu}(q_i, \omega_i, \pi) \geq \tilde{\nu}(q_i, \pi)$.

ii. The incentive compatibility constraint to acquire information becomes less stringent as $|q_i - \pi|$ *decreases. If* $q_i \leq \pi \leq 1/2$ *or* $1/2 \leq \pi \leq q_i$*, then* $\frac{\partial \nu_i}{\partial |q_i - \pi|} \leq 0$*. Moreover, if prior beliefs in the economy are distributed according to a density* q *with full support on* $[0, 1]$*, then there exists a positive probability mass of traders with prior* q_i *around* π *who acquire information.*

iii. The incentive compatibility constraint becomes less stringent as ω_i *increases for sufficiently large* ω_i*. If, for some* $\omega_i > 0$*, informative signals are acquired, then larger* ω_i *strictly increases signal precision until* $\nu_i = \hat{\nu}$*.* ($\frac{\partial \nu_i}{\partial \omega_i} \geq 0$)

Proof. See Appendix. □

Traders who have a prior belief close to the price π are more likely to acquire information (*ii.*). To understand the intuition, note that traders compute the expected utility based on their prior belief when deciding about acquiring information. Whenever the prior deviates

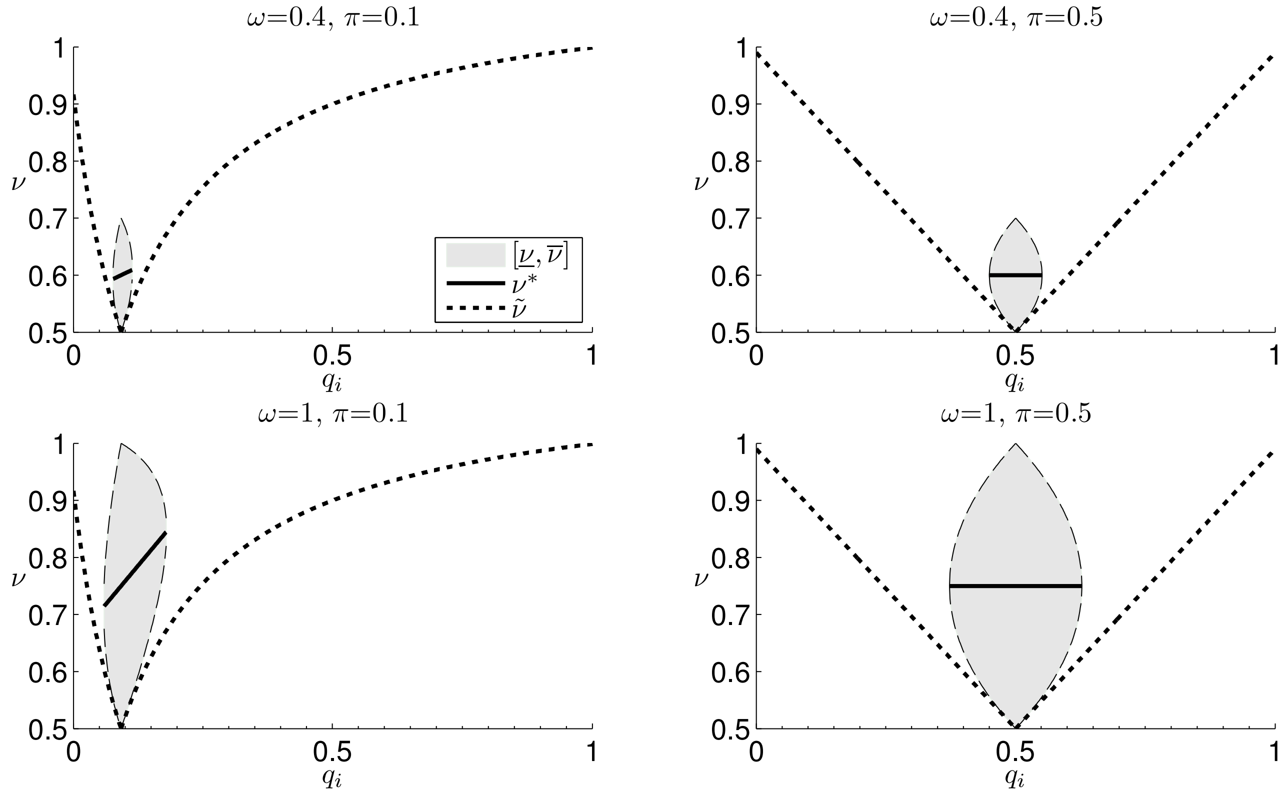

Figure 2: Information acquisition decision: optimal (solid line), incentive compatible (shaded area) and minimum discriminating (dotted line) precision levels, with varying endowment ω and price π, depending on prior belief q_i.

considerably from the price, the trader expects large gains $|q_i - \pi|$ per contract held without information acquisition (see (1)). This can be interpreted as a trader having a strong opinion about the outcome, who does not expect to be swayed by evidence and hence does not acquire it. Conversely, the trader expects only small gains if q_i is close to π, so his opinion about the outcome is not as strong, and he is willing to acquire information and possibly revise his beliefs if there is evidence contradicting his prior.

All else equal, higher endowment makes information acquisition more likely, or increases information seeking effort, if endowment is sufficiently large (*iii.*). The intuition is that, while endowment (and thus potential gains and losses) scale up, the cost of acquiring information remains the same. Thus, due to higher stakes, traders want to be more certain that their investment decision will be the right one. Peress (2004) obtains this comparative static for the same reason in his rational expectations model.

Notice that the motivation of the traders is not to make the best prediction possible. If that were the case, traders would also want to obtain information if the prior is far from the price. Instead, the motivation is to maximize utility, and that might require saving on effort costs and accepting inferior information.

These results are illustrated in Figure 2, in which the optimal precision level ν^*, the range of incentive compatible precision levels $[\underline{\nu}, \overline{\nu}]$ (conditional on being discriminating) and the minimum discriminating precision level $\tilde{\nu}$ for different values of ω, π, q_i are plotted.

It illustrates how a higher endowment increases the range of priors where information acquisition is incentive compatible, and increases the chosen precision, until ν^* reaches upper bound $\hat{\nu}$ (which is not included in the figure for visibility). The figure also suggests that the range of types q_i investing effort is smaller if the market price deviates from $\pi = 1/2$.

2.2 Competitive equilibrium

How does the possibility of information acquisition affect the forecast of the prediction market? The following example illustrates how information acquisition can improve the forecast compared to a model without, as for example used in Gjerstad (2005), Wolfers and Zitzewitz (2006), or Manski (2006).

Example. Suppose $\theta = 1$, each trader has endowment $\omega = 1$ and priors q_i are uniformly distributed, i.e., $q_i \sim U[0, 1]$. Without information acquisition, the equilibrium price is $\pi_0^* = 1/2$, because at that price half of the population is willing to invest in A and half in B, thus clearing the market. The equilibrium price divides the continuum in traders who always invest in A ($q_i > \pi_0^*$), and traders who always invest in B ($q_i < \pi_0^*$). I shall call these 'categorical' traders, because their decision is based solely on their prior, and therefore independent of the true state of the world.

The previous section showed that information acquisition is incentive compatible only for traders with prior close to the price. When allowing information acquisition (keeping the price constant), traders with $q_i < \pi_0^*$ just left of 1/2 turn from categorical B traders into discriminatory traders acquiring information (who trade contingent on signal), and the share of those investing in the correct option A improves from 0 to $\nu_i(q_i, \omega, \pi_0^*) > 1/2$. Conversely, traders with $q_i > \pi_0^*$ just right of 1/2 turn from categorical A traders into discriminatory traders, and the share of those investing in A decreases from 1 to $\nu_i(q_i, \omega, \pi_0^*) > 1/2$. The mass of traders turned into discriminatory traders is equal for both sides about 1/2 if $\pi = 1/2$ (Figure 2), so the mass of traders in the population willing to invest in the correct option A increases. As a consequence, the equilibrium price must change. Indeed, if demand is weakly decreasing in the price, it must increase toward the true value $\theta = 1$, because a larger mass than 1/2 is willing to invest in A at price $\pi = 1/2$. □

As the example illustrates, the possibility of information acquisition can sway traders with wrong opinion to invest in the correct outcome instead, thus improving the forecast.

More formally, let $\nu(q_i, \omega_i, \pi)$ denote the precision level resulting from the endogenously chosen information acquisition effort. For readability, I will omit the conditioning set of cdf $Q(q_i|\theta, \omega_i)$ in the following. Following definition 1, the equilibrium price π^* is implicitly defined as the fixed point of ($\mathbf{1}\{.\}$ is the indicator function)

$$\int_0^{\overline{\omega}}\int_0^1 a_i(p_i,\omega_i,\pi^*)/\pi^* dQ(q_i)dW(\omega_i) = \int_0^{\overline{\omega}}\int_0^1 b_i(p_i,\omega_i,\pi^*)/(1-\pi^*)dQ(q_i)dW(\omega_i)$$
$$\iff \int_0^{\overline{\omega}} \omega_i\Bigg(\int_0^1 \mathbf{1}\{\nu(q_i,\omega_i,\pi^*)=1/2\}\mathbf{1}\{q_i \geq \pi^*\} + \mathbf{1}\{\nu(q_i,\omega_i,\pi^*)>1/2\}$$
$$[\theta\nu(q_i,\omega_i,\pi^*) + (1-\theta)(1-\nu(q_i,\omega_i,\pi^*))]dQ(q_i) - \pi^*\Bigg)dW(\omega_i) = 0,$$

where $\mathbf{1}\{\nu(\pi^*,q_i) = 1/2\}\mathbf{1}\{q_i \geq \pi^*\}$ indicates categorical A traders, who do not acquire information and always invest in A, and the remaining term is the contribution of discriminatory traders to A demand, who invest according to their information. They invest in A only in $\nu(q_i,\omega_i,\pi^*)$ of the cases if the true state of the world is A, otherwise in $1-\nu(q_i,\omega_i,\pi^*)$ of the cases. The equivalence is shown in the proof of Proposition 2 (*ii.*). Rewriting the share of categorical A investors in the population at π^*,

$$\int_0^1 \mathbf{1}\{\nu(q_i,\omega_i,\pi^*)=1/2\}\mathbf{1}\{q_i \geq \pi^*\}dQ(q_i) = \int_{\pi^*}^1 \mathbf{1}\{\nu(q_i,\omega_i,\pi^*)=1/2\}dQ(q_i) = 1-Q(t_u(\pi^*)),$$

where t_u (t_l) is the upper (lower) threshold for incentive compatible priors. Explicit expressions for the thresholds, which depend on π and ω_i, are derived in the appendix. Since the precision is bounded by $\hat{\nu}$, for $\theta = 1$

$$\int_0^1 \mathbf{1}\{\nu(q_i,\omega_i,\pi^*)>1/2\}\nu(q_i,\omega_i,\pi^*)dQ(q_i) = \int_{t_l}^{t_u} \min\left\{\nu^*(q_i,\omega_i,\pi^*),\hat{\nu}\right\}dQ(q_i).$$

Hence, the equilibrium price π^* if $\theta = 1$ fulfills

$$\int_0^{\overline{\omega}}\left(1-Q(t_u) + \int_{t_l}^{t_u}\min\left\{\frac{q_i\alpha_i + (1-q_i)\beta_i + 4}{8},\hat{\nu}\right\}dQ(q_i) - \pi^*\right)\omega_i dW(\omega_i) = 0. \quad (5)$$

The following proposition shows that an equilibrium always exists, and can be shown to be unique in most cases.

Proposition 2.

i. There exists an equilibrium price π^* fulfilling definition 1,

ii. which is implicitly defined as the fixed point of (5).

iii. Assuming homogeneous endowment ($\omega_i = \omega$), the equilibrium is unique for ω such that $\hat{\nu}$ is binding,

iv. and unique for ω such that $\hat{\nu}$ is non-binding if ω is small.

Proof. See Appendix. □

In most cases, A-demand decreases in response to a price increase, because categorical A traders with q_i close to t_u are replaced by discriminatory traders, who only invest in A in $\nu_i < 1$ of the cases, and discriminatory traders with q_i close to t_l are replaced by categorical B traders, who never invest in A. Let us call this the price effect. Still, the demand for the A-contract need not necessarily be decreasing in the price. Since the precision chosen by discriminatory traders $\nu^*(q_i, \omega_i, \pi)$ increases for large π, a price increase may result in more investments in A if $\theta = 1$. However, this effect is strong enough to reverse the price effect only if π is large and the distribution of priors is very unsmooth.[3] Since this is a rather artificial situation, in the following I assume demand to be weakly decreasing in the price. The validity of this assumption is confirmed in all numerical examples.

In order to be more specific on how the distribution of priors affects the equilibrium outcome, I consider two parametric forms of Q. First, priors may be distributed according to normal distribution $\mathcal{N}(\mu, \sigma^2)$. For computational ease, instead of truncating at 0 and 1, any probability mass that would have been truncated is assumed to be located at the corners.[4] Second, the triangular distribution is a computationally simpler distribution that also allows to shift probability mass closer to or farther away from the truth. This property is important to answer how a change of opinion in the economy influences the market prediction π^* or whether a more informed public generates better forecasts. The triangular distribution is easier to handle analytically and there are no boundary issues. It is distributed on the interval $[0, 1]$ with the density peak at $d \in [0, 1]$ and density function

$$q(q_i) = \begin{cases} \frac{2}{d} q_i & \text{if } q_i \leq d, \\ \frac{2(1-q_i)}{1-d} & \text{if } q_i > d. \end{cases}$$

The following proposition gives some comparative statics regarding the market price if $\theta = 1$. This condition is without loss of generality, since outcomes A and B can be redefined. The results are sufficient conditions for a shift of priors toward the true value to reduce bias (increase the equilibrium price), and for an endowment increase to reduce bias. Assumption $\pi^* \geq 1/2$ means that ω, Q are such that the market correctly predicts the favorite A, i.e., π^* is closer to $\theta = 1$ than to 0, prior to the exogenous parameter change.

[3]For example, a distribution with small density for q_i around t_u and around t_l and a large density for q_i between t_l and t_u leads to little loss of A-demand via price effect when the price increases, while all q_i in the interior increase their precision.

[4]In the comparative statics analysis, we are interested in how the market price changes if the prior belief distribution changes. However, the cdf of the truncated normal distribution can be decreasing in μ at some $x \in (0, 1)$, because of truncation. This means the equilibrium price in the model without information acquisition (Manski, 2006) can *decrease* in response to an *increase* of μ, so truncation introduces additional effects depending on the parameters, which interacts with the comparative static effects to be analyzed. For example, a negative response to a μ-increase could be due to changes in truncation, or due to reduced information acquisition effort. To solve this problem, I assume any truncated mass is located at the corners, so that the prior belief distribution in the interior is just the cdf of a normal distribution, which is always decreasing in response to an increase of μ.

Proposition 3. *Comparative statics if demand is non-increasing in the price, ω is homogeneous and $\theta = 1$.*

i. *Any change in the prior distribution from Q to R such that $Q(q_i) > R(q_i)$ $\forall q_i \in (0,1)$ increases the market price π^* if $\hat{\nu}$ is binding, unless all agents in the Q-economy already acquire information.*

ii. *With priors $q_i \sim \mathcal{N}(\mu, \sigma^2)$ and $\omega < (1-\pi^*)4$, an upward shift of μ increases the equilibrium price π^* for each μ, σ if $\pi^* \leq 1/2$.*

iii. *With priors $q_i \sim \mathcal{N}(\mu, \sigma^2)$ and endowment $\omega < (1-\pi^*)4$ sufficiently small, an upward shift of μ increases the equilibrium price π^* for each μ, σ and π^*.*

iv. *With priors from the triangular distribution, an increase in the mode value d leads to an increase of the market prediction π^* if $d < t_l$ or $d > t_u$.*

v. *Higher endowment increases information in the economy as measured by*

$$D_A = \int_{t_l}^{t_u} \min\{\nu^*(q_i), \hat{\nu}\} dQ(q_i),$$

i.e., more traders obtain information of higher precision.

vi. *For any distribution of priors with non-increasing density q the market prediction π^* increases in response to a non-binding endowment increase if $\pi^* \geq 1/2$.*

vii. *With priors $q_i \sim \mathcal{N}(\mu, \sigma^2)$ the market prediction π^* increases in response to an endowment increase for $\pi^* \geq 1/2$, $\hat{\nu}$ non-binding and $|\mu - t_l| < |\mu - t_u|$.*

Proof. See Appendix. □

The market price increases if probability mass is shifted towards higher values of q_i and if the precision is unaffected by price (*i.*). In particular, this includes an increase of d for the triangular distribution or an increase of μ for the normal distribution. However, if the precision cap is not reached, this need not always hold. If $\pi > 1/2$, then the optimal precision $\nu^*(q_i, \pi)$ is decreasing in q_i, because the profitable state B (price only $1-\pi < 1/2$) is considered less likely. Thus, if the mass of traders with high q_i increases, then it can decrease the mass of agents investing in A for a given price π. In particular, this occurs for some parameter values using the triangular distribution with $t_l < d < t_u$, or for the normal distribution, $\pi > 1/2$ and σ small. It is therefore possible that a population with *more accurate* opinions produces *worse* forecasts, because it invests less in acquisition effort. Still, in most cases the comparative static holds, because more weight on high values of q_i implies more categorical A traders and more weight on discriminatory traders, who favor A. Moreover, small non-monotonicity is usually a local problem and vanishes for sufficiently large discrete changes of the exogenous variable. The price always increases in response to increases of ω or μ in the numerical example using discrete changes (section 2.3).

Parts (*ii.*) and (*iii.*) state that a shift of μ toward the true value $\theta = 1$ increases the market price if endowments are sufficiently low or if $\pi^* < 1/2$. Intuitively, if endowment is low, then the mass of discriminatory traders is low and so the fact that the precision of discriminatory traders may be reduced has little weight.

Parts (*v.*)-(*vii.*) consider an endowment increase in the economy. While higher endowment always increases information in the economy as measured by investment from discriminatory traders in the correct outcome A, this need not imply that the bias always decreases in response to an endowment increase. After all, if opinions happen to be very accurate, then correct opinions are replaced with information that is correct in only $\nu_i^* < 1$ of the cases. Still, as with a shift of priors toward the correct outcome, this comparative static holds in all numerical cases (section 2.3).

2.3 Numerical example

2.3.1 Preliminaries

The purpose of this section is threefold. First, it demonstrates that this model produces better predictions than the model without information acquisition (Manski, 2006). Second, it quantifies the impact of exogenous variables on π^*. And third, it shows the comparative statics are less ambiguous for discrete changes than the theoretical results may let on.

I calculate the equilibrium price π^* numerically for various parameter values. Priors are assumed to be normally distributed, but any probability mass $\int_1^\infty q(q_i)dq_i$ is bunched at 1 (similarly all mass below 0 is bunched at 0). This is preferable to truncation, as the truncated mass depends on parameters μ, σ, which hinders identification of comparative static effects (see footnote 4). Consequently, $\mu \neq 1/2$ is not actually the mean of the prior distribution, but merely a position parameter.

The equilibrium price with information acquisition is defined in (5), where (Φ denotes the cdf of the standard normal distribution, ϕ its density function)

$$Q(q_i) = \Phi\left(\frac{q_i - \mu}{\sigma}\right), \; q(q_i) = \frac{1}{\sigma}\phi\left(\frac{q_i - \mu}{\sigma}\right) \; \forall q_i \in (0, 1).$$

The equilibrium price without information acquisition, where traders rely solely on their priors, is given by (e.g., Manski, 2006)

$$\int_0^{\overline{\omega}} (1 - Q(\pi_0^*) - \pi_0^*)\omega_i dW(\omega_i) = 0.$$

I fixed $\hat{\nu} = 0.95$ as maximal signal precision, i.e., no trader can perfectly infer the

Table 1: Equilibrium price with information acquisition π^* for various parameter values.

ω	μ	σ	θ	π^*	D_B	C_B	π_0^*
0.5	0.1	0.1	0	0.18	0.13	0.69	0.19
0.5	0.1	0.3	0	0.27	0.07	0.66	0.28
0.5	0.5	0.1	0	0.49	0.29	0.22	0.5
0.5	0.5	0.3	0	0.49	0.1	0.41	0.5
0.5	0.7	0.1	0	0.65	0.23	0.12	0.66
0.5	0.7	0.3	0	0.61	0.09	0.3	0.61
1	0.1	0.1	0	0.15	0.32	0.52	0.19
1	0.1	0.3	0	0.24	0.17	0.59	0.28
1	0.5	0.1	0	0.43	0.53	0.03	0.5
1	0.5	0.3	0	0.46	0.24	0.3	0.5
1	0.7	0.1	0	0.6	0.39	0.01	0.66
1	0.7	0.3	0	0.58	0.22	0.2	0.61
5	0.1	0.1	0	0.04	0.77	0.18	0.19
5	0.1	0.3	0	0.09	0.51	0.4	0.28
5	0.5	0.1	0	0.13	0.87	0	0.5
5	0.5	0.3	0	0.24	0.69	0.07	0.5
5	0.7	0.1	0	0.28	0.72	0	0.66
5	0.7	0.3	0	0.37	0.6	0.02	0.61

ω	μ	σ	θ	π^*	D_A	C_A	π_0^*
0.5	0.1	0.1	1	0.2	0.12	0.08	0.19
0.5	0.1	0.3	1	0.28	0.07	0.21	0.28
0.5	0.5	0.1	1	0.51	0.29	0.22	0.5
0.5	0.5	0.3	1	0.51	0.1	0.41	0.5
0.5	0.7	0.1	1	0.67	0.25	0.42	0.66
0.5	0.7	0.3	1	0.62	0.09	0.53	0.61
1	0.1	0.1	1	0.22	0.22	0.01	0.19
1	0.1	0.3	1	0.3	0.17	0.13	0.28
1	0.5	0.1	1	0.57	0.53	0.03	0.5
1	0.5	0.3	1	0.54	0.24	0.3	0.5
1	0.7	0.1	1	0.71	0.54	0.17	0.66
1	0.7	0.3	1	0.65	0.22	0.43	0.61
5	0.1	0.1	1	0.44	0.44	0	0.19
5	0.1	0.3	1	0.45	0.45	0.01	0.28
5	0.5	0.1	1	0.87	0.87	0	0.5
5	0.5	0.3	1	0.76	0.69	0.07	0.5
5	0.7	0.1	1	0.93	0.93	0	0.66
5	0.7	0.3	1	0.85	0.66	0.19	0.61

Description: The table displays the equilibrium price with information acquisition π^* and without (π_0^*). Priors are drawn from distribution $\mathcal{N}(\mu, \sigma^2)$, and endowment is ω. The true state of the world is θ, so the bias is $\pi^* - \theta$. D_A is the probability mass of discriminatory bettors betting on A in equilibrium, while C_A is the probability mass of categorical bettors betting on A. The upper bound of the signal precision is $\hat{\nu} = 0.95$.

true state of the world by investing sufficient information acquisition effort. Endowment is homogeneous, with ω as a parameter to be varied. Finally, I confirmed numerically that the equilibrium is unique.

Table 1 displays the equilibrium price with information acquisition (π^*), and without (π_0^*), as well as the corresponding parameter values $\mu, \sigma, \omega, \theta$. The bias of the market prediction $\pi^* - \theta$ can be easily computed. It also shows D_j, $j = A$ iff $\theta = 1$, $j = B$ iff $\theta = 0$, which is the probability mass of discriminatory traders investing in the correct outcome in equilibrium. This is a measure of information in the economy. A high value of D_j either means many traders acquire information, or that information is good (i.e., signals with high precision are acquired). For $\theta = 1$, $D_A = \int_{t_l}^{t_u} \min\{\nu^*(\pi^*, q_i), \hat{\nu}\} dQ(q_i)$. C_j is the probability mass of categorical traders investing in the correct outcome, i.e., the mass of traders that does not acquire information in equilibrium, but has prior q_i that makes them invest in the correct option. For $\theta = 1$, $C_A = 1 - Q(t_u)$. By definition of equilibrium, $C_A + D_A = \pi^*$.

2.3.2 Numerical results

A higher endowment ω leads to smaller bias $|\pi^* - \theta|$. The reason is that, first, more endowment leads to a higher mass of traders acquiring information. And second, traders who do acquire information choose more precise information (i.e., higher precision). This is reflected by D_j, the probability mass of discriminatory traders investing in the correct outcome, which is increasing in endowment. Some correctly investing categorical traders are turned into discriminating traders, who choose the wrong outcome in $1 - \nu_i$ of the cases, but the net effect as represented by π^* is nevertheless unambiguously bias reducing. As endowment rises, the market price is more and more driven by information rather than prior beliefs. For example, at $\omega = 5$, $\theta = 1$ and $\sigma = 0.1$, the probability mass of categorical A traders C_A is 0 (rounded), so it is almost exclusively discriminatory traders investing in A and supporting the equilibrium, while C_A is positive for lower endowment.

Further, μ closer to the true state θ decreases the absolute value of the bias. This not surprising, since it means there is a larger probability mass of categorical traders for each given price π willing to invest in the right choice, all else equal. And although the results in Proposition 3 seemed to suggest that π^* may not always be monotone in μ, these examples show that for plausible cases it is.

The effect of σ is ambiguous. Generally, the parameter shifts probability mass closer to or farther away from μ. If μ is small, then a higher σ means there is a larger probability mass at high values of q_i. Consequently, a higher σ increases bias for low μ and $\theta = 0$, or for high μ and $\theta = 1$. That is, there are more wrongly investing categorical traders, which drive the market price in the wrong direction. Similarly, a larger σ improves the prediction if μ is far from θ, because it increases the mass of correctly investing categorical traders. The effect is more pronounced with lower ω, where the prior has more influence on the market price. With larger σ, the amount of information is reduced (with the exception of $\theta = 1, \omega = 5, \mu = .1$). This is because there is less probability mass for priors around the market price, which is usually close to μ.

The prediction in the model without information acquisition π_0^* is unaffected by endowment, because it does not influence beliefs. Moreover, the prediction is the same independent of the state of the world, because investment decisions are not correlated with θ. Finally, the forecasts in the model with information acquisition are never worse, and usually strictly better than in the model without. Averaging over the 36 cases in Table 1, the forecast bias in the Manski model is larger by 0.1, i.e., about 10.5 percentage points, given the same endowment and prior distribution. Therefore, allowing for information acquisition goes a long way in explaining forecast accuracy. But note that the model without information acquisition may produce better forecasts for very unsmooth prior distributions Q, as the following example shows.

Example. Take any ω and Q with non-increasing aggregate demand. Without loss of generality, consider $\theta = 1$. Comparing the model without (no) information acquisition (NIA) with equilibrium price π_0^* and the model with acquisition (IA), bias in IA is smaller ($\pi^* > \pi_0^*$) if and only if A-investments by discriminatory traders in IA more than replace investments by categorical A traders in NIA that are discriminatory traders in IA at $\pi = \pi_0^*$. That is, if and only if, evaluated at π_0^*,

$$\int_{t_l}^{t_u} \min\left\{\frac{q_i\alpha_i + (1-q_i)\beta_i + 4}{8}, \hat{\nu}\right\} dQ(q_i) > Q(t_u) - Q(\pi_0^*). \tag{6}$$

In this case, price π_0^* leads to excess demand for A in IA, so, because demand is non-increasing, only $\pi^* > \pi_0^*$ clears the market. To construct a counterexample, consider the piecewise uniform density function

$$q(q_i) = \begin{cases} m, & \text{for } q_i \in [0, 1/4] \cup (1/2, 3/4] \\ 2-m, & \text{for } q_i \in (1/4, 1/2] \cup (3/4, 1], \end{cases}$$

where $2 > m > 1$. The equilibrium price without information acquisition is $\pi_0^* = 1/2$, since $1 - Q(1/2) = 1/2$. Just right of $1/2$ is a higher density than just left of $1/2$. By symmetry, $|t_l - 1/2| = |t_u - 1/2|$ for $\pi = 1/2$. Now, ω small enough such that $t_u \leq 3/4$ implies $Q(t_u) - Q(1/2) = (t_u - 1/2)m$, while the LHS of (6) is $2\nu^*(t_u - 1/2)$ (ν^* is constant for all $q_i \in [0, 1]$ at $\pi = 1/2$). Hence, the market price is larger (bias is smaller) in IA if and only if $\nu^*(\omega, \pi = 1/2) > m/2$. □

Intuitively, if there is a large probability mass of categorical A traders in NIA, who are discriminatory traders with low signal precision in IA, while only a small probability mass of categorical B traders is affected, then demand for the A-contract at $\pi = 1/2$ decreases when allowing for information acquisition. By the demand monotonicity assumption, this implies $\pi^* < \pi_0^*$, i.e., greater bias given $\theta = 1$. Nevertheless, for large ω or smoother prior distributions information acquisition improves the market forecast.

2.4 Endogenous weighting: endowment heterogeneity

In elections or votes, each "voice" has equal weight. For purposes of preference aggregation, this may be just, but if voting is meant to determine an objectively correct state, then equal weighting need not yield the best outcome. Indeed, if pooling of information is not possible, the optimal weighting is to give full weight to the agent with the best information about the state. However, it is usually not obvious who that agent is. In the asset market, this dilemma is solved endogenously by weighting each bet with its wager, or each trade with its volume ("weighting effect"). Combining with the earlier result that higher endowment induces more information acquisition ("incentive effect"), the market endogenously gives higher weight to traders with better forecasts, all else equal.

This interaction of the incentive effect and the weighting effect is my main explanation why prediction markets work. It is not necessary that all traders in the economy have accurate beliefs about the outcome. Instead, it is sufficient to have a few high endowment agents, who choose to be informed because of large stakes, and can drive the market price in the right direction due to their large weight in the market. In conventional financial markets, these high endowment agents may be hedge funds and investment banking divisions, who can move millions of dollars and have access to research into potential investment opportunities that smaller investors cannot afford.

In the following, I assume for simplicity there are two endowment groups (ω_1, Q_1) and (ω_2, Q_2), represented by endowment and prior distribution. The share of traders of group 1 in the economy is $1 > \gamma > 0$. The equilibrium price in a hypothetical market populated only by group 1 is π_1^*, and the price in the hypothetical low endowment market is π_2^*, which is a way of expressing prior distribution and endowment of the group in terms of a forecast. Lemma 4 in the Appendix shows that the market price in the economy with both groups is strictly between π_1^* and π_2^* if demand is non-increasing in the price, i.e., $\min\{\pi_1^*, \pi_2^*\} < \pi^* < \max\{\pi_1^*, \pi_2^*\}$. An immediate consequence of this result is the following.

Corollary 5. *If the economy consists of two endowment groups $\omega_1 > \omega_2$, and demand is non-increasing in the price for both groups, then the market price π^* has a smaller bias than the price in a market consisting only of the low endowment group π_2^* if and only if the bias in the high endowment economy is smaller ($|\theta - \pi_1^*| < |\theta - \pi_2^*|$).*

This corollary implies that if one could add a group of high endowment traders or traders with better information to the market, then the market prediction improves as a result. To analyze the weighting effect, consider an equal weight price I^*, defined in the general case as

$$\int_0^1 \mathbf{1}\{a_i(q_i, \omega_i, I^*) > b_i(q_i, \omega_i, I^*)\} dQ(q_i) = I^*.$$

The price is determined by equating the share of the number of investments in A (instead of the share of dollars in A) with the price I^* that induces these investments. This method gives equal weight to every trader, but it is not a market clearing price whenever it diverges from π^*. Thus, the market maker might have to use his own funds to pay the winners, or he might make a profit. If endowment is homogeneous or if $\pi_1^* = \pi_2^*$, then $\pi^* = I^*$.

Proposition 6. *Suppose the economy consists of two endowment groups $\omega_1 > \omega_2$, and demand is non-increasing in the price for both groups. If $\pi_1^* > \pi_2^*$, then $\pi^* > I^*$. Similarly, $\pi_1^* < \pi_2^*$ implies $\pi^* < I^*$.*

Proof. See Appendix. □

Corollary 7. *If $|\pi_1^* - \theta| < |\pi_2^* - \theta|$, then endogenous weighting compared to equal weighting is unambiguously bias reducing, i.e., $|\pi^* - \theta| < |I^* - \theta|$, but the improvement is smaller than for optimal weighting, i.e., $|\pi_1^* - \theta| < |\pi^* - \theta|$.*

As shown in the previous sections, a larger endowment leads to smaller bias $|\theta - \pi^*|$. Thus, the endogenous weighting implied by market clearing is bias reducing if prior beliefs and endowment are independent in the economy, or if higher endowment groups tend to have more accurate priors. An evolutionary argument makes the latter condition plausible: traders with better priors have better forecasts, so they increase their endowment by making the right investments. Hence, over time high endowment traders are the ones with better priors (e.g., Servan-Schreiber et al., 2004). Conversely, traders with bad priors "die out," be it because they go broke or because they recognize their inability to forecast and stop investing. Endogenous weighting increases bias only if the low endowment group has considerably better priors to compensate for the inferior information due to the incentive effect, so that π_2^* is closer to θ than π_1^*. In summary, the endogenous weighting of the market improves on the effect of information acquisition by giving more weight to better forecasts.

3 Discussion: Interpreting prediction market prices

Several authors have asked what prediction market prices represent. Contrary to common interpretation, Manski (2006) shows that an equilibrium price with risk neutral traders, interpreted as probability forecast, may be far from the mean belief in the economy. Gjerstad (2005) and Wolfers and Zitzewitz (2006) demonstrate that this disparity is smaller for risk averse agents. They use a non-Bayesian set-up with arbitrary distribution of beliefs, which may or may not be close to the actual outcome. Consequently, the mean belief is a convenient summary statistic, but it may not be a good predictor of the outcome. What we are instead interested in is how price π^* relates to the outcome θ. The key question is whether the price is driven by opinion, which may be close or far from the truth, or by information and evidence, which is correlated with the outcome and therefore more reliable. Even if initial beliefs about the outcome are off, the prediction market may generate a good forecast if agents have sufficient incentives to seek out information and revise their beliefs and investments. Of course, it is possible to interpret the belief distribution of Manski et al. as posterior distribution after information acquisition. Any information acquisition equilibrium based on prior distribution Q can be reached in the model without information acquisition and distribution R such that (assuming homogeneous endowment)

$$1 - R(\pi^*) = 1 - Q(t_u) + \int_{t_l(\pi^*)}^{t_u(\pi^*)} \min\left\{\frac{q_i\alpha_i(\pi^*) + (1-q_i)\beta_i(\pi^*) + 4}{8}, \hat{\nu}\right\} dQ(q_i) = \pi^*.$$

Yet, merely assuming such a distribution without modeling its generation does not give any insights into its informational content. Moreover, meaningful comparative static analysis is not possible, as R is fixed for ω and Q, and changes in information acquisition (e.g., incentive compatibility) would not be captured.

To understand the role of information in the price, suppose all traders acquire signals of equal precision $\nu > 1/2$, and invest according to their signal, i.e., in A iff $s_i = 1$ and in B iff $s_i = 0$. Then the unique equilibrium price for the case $\theta = 1$ is $\pi^* = \nu$, because exactly the fraction ν of the population invests in A. Hence, the equilibrium price represents the quality of information in this economy and the prediction market favorite would perfectly forecast the outcome. But traders do not always invest based on solid information. In the general model with equilibrium (5), the price represents a mixture of information and opinion. Taking only the investments of informed traders, D_A/D_B reveals the true state of the world perfectly in the continuum ($D_A/D_B > 1$ iff $\theta = 1$), but in anonymous markets these can hardly be told apart from investments driven by opinion. Thus, if neither endowments ω_i nor the distribution of priors Q in the economy are known, the price could be anything from an aggregation of wrong opinions to a very precise forecast. If at least endowments were known, then an outside observer would infer, ceteris paribus, that the price is a better predictor the larger the endowments.

A market observer, who knows that accuracy of priors and endowment are not negatively correlated, can improve on the market forecast by giving more weight to large investments. In theory, he could perfectly infer the true state of the world by observing the difference of investment behavior (conditional on price) of high endowment and low endowment traders as the number of traders becomes large, because investments from high endowment traders are more likely to be correct. However, analyzing investment behavior of traders, Rothschild and Sethi (2013) note that motives other than maximizing financial returns may play a role in practice. Hence, inferring information in favor of A from a large investment in A may not be valid. For example, Rothschild and Sethi (2013) speculate that large investments in the Romney contract in the 2012 US presidential election may have been a manipulation attempt to keep the price from plummeting. This is likely not a large concern for accuracy, otherwise predictions would not be as accurate as we observe. Nevertheless, there is more to a prediction market price than the mean belief of traders.

4 Conclusion and future research

Costly information acquisition explains the existence of uninformed traders, who do not acquire information and rely only on their opinion when investing, and of informed traders. In my model, noise traders, whose existence is usually just assumed to avoid the implications of no trade theorems, and informed traders evolve endogenously from an initial distribution of opinions and endowment. Thus, good forecasts are not explained merely by the existence of insiders in the market (and bad ones by their absence). Rather, the information acquisition explanation implies that accurate initial beliefs as well as large endowment—improving incentives for information acquisition—can lead to good forecasts, but low stakes and inaccurate beliefs may also lead to bad ones. It also implies that larger weight on bets or votes of high endowment agents may improve information aggregation, at least as long as pooling of information is difficult and other motives such as manipulation (Rothschild and Sethi, 2013) play a minor role. Endogenous weighting has a bias reducing effect, unless endowments and priors are negatively related.

There appears to be only one empirical paper that attempts to shed light on the role of money as incentive in futures markets. Servan-Schreiber et al. (2004) compare predictions of a play money prediction market and a real stakes prediction market. They find no significant difference in forecast accuracy, but stress that the play money market has stronger selection of good forecasters, because large 'play endowment' can only be obtained with a record of correct predictions. Moreover, the top traders could redeem play money for prizes, thus providing different (rank order) incentives for information acquisition. A direct test of the information acquisition explanation, e.g., by endowing traders with play money randomly to rule out selection effects, is yet to be done.

An evolutionary explanation is not just compatible, indeed it complements the information acquisition explanation. Successful investing increases endowment, which improves incentives for information acquisition and thereby future success, along with weight in the market. Recently, Blume and Easley (2006) asked whether agents with more accurate beliefs survive over time in the market, whereas agents with inaccurate beliefs vanish. In an infinite horizon consumption model, they show that risk averse agents in complete markets indeed survive only if they have correct beliefs, given that someone else does. If no one has accurate beliefs, then those with beliefs closest to the truth survive, assuming homogeneous discount factors. However, their results may not be directly applicable to prediction markets. For instance, they do not model entry of new traders. And if there is continuous entry of new, biased traders, then the population will never be completely free of biased beliefs. Moreover, selection with respect to beliefs is retroactive, i.e., survivors have had correct forecasts. Consequently, selection need not guarantee accurate forecasts in the future if novel events come up for prediction. A question for empirical research is whether there is a forecasting ability, so that past accuracy is positively related to future accuracy, or whether selection just favors those who "got lucky," in which case there would be a regression-toward-the-mean effect.

Appendix A: Upper and lower threshold of incentive compatible priors

The goal is to find t_u, the upper threshold of prior belief q_i, where information acquisition is just incentive compatible. Note that, for large ω_i, precision levels ν_i^* may be greater 1 and therefore violate the axioms of probability. Moreover, for large ω_i, there is no type $q_i \in [0, 1]$ for which no information acquisition is preferable (i.e., there is no real solution for q_l, q_u below). This requires case distinctions. Ignoring constraints, the type that is indifferent between information acquisition and relying on his prior knowledge is found by setting the square root term in (3) equal to zero, which gives q_u, and including constraints we obtain

$$t_u = \begin{cases} q_u(\pi) = \frac{-y-\sqrt{y^2-4xz}}{2x} & \text{if } \nu^*(q_u) \leq \hat{\nu},\ q_u \leq 1 \text{ and } q_u \in \mathbb{R}, \\ q : \underline{\nu}(q) = \hat{\nu} \wedge q \geq \pi & \text{otherwise,} \end{cases}$$

where $x = (\alpha - \beta)^2$, $y = -2\beta^2 - 8\alpha - 8\beta + 2\alpha\beta$, $z = \beta^2 + 8\beta$. q_u is the smaller of the two solutions of the U-shaped quadratic equation. Only the case $q_i \geq \pi$ has to be considered here, because at $q_i = \pi$ information acquisition is always incentive compatible (Proposition 1), so the upper threshold must be larger than the price. For the lower threshold type q_l, the square root term in (4) is similarly set to zero, resulting in

$$t_l = \begin{cases} q_l(\pi) = \frac{-y+\sqrt{y^2-4xz}}{2x} = 1 - q_u(1-\pi) & \text{if } \nu^*(q_l) \leq \hat{\nu},\ q_l \geq 0 \text{ and } q_l \in \mathbb{R}, \\ q : \underline{\nu}(q) = \hat{\nu} \wedge q \leq \pi & \text{otherwise,} \end{cases}$$

where $x = (\alpha - \beta)^2$, $y = -2\beta^2 + 8\alpha + 8\beta + 2\alpha\beta$, $z = \beta^2 - 8\beta$. This is the larger of the two solutions to the U-shaped quadratic equation. The following properties of the threshold functions are used in the proofs. With respect to the price,

$$\frac{\partial q_u}{\partial \pi} > 0,\ \frac{\partial q_l}{\partial \pi} > 0, \text{ and } \frac{\partial q_u}{\partial \pi} > \frac{\partial q_l}{\partial \pi} \iff \pi < 1/2,\ \frac{\partial q_u}{\partial \pi} < \frac{\partial q_l}{\partial \pi} \iff \pi > 1/2.$$

Using symmetry,

$$\frac{\partial q_u}{\partial \pi} = (1 - q_l(1-\pi))' = q_l'(1-\pi) = q_u'(\pi),\ q_l'(1/2) = q_u'(1/2).$$

Moreover,

$$\frac{\partial q_u}{\partial \omega} > 0,\ \frac{\partial q_l}{\partial \omega} < 0, \text{ and } \frac{\partial q_u}{\partial \omega} > -\frac{\partial q_l}{\partial \omega} \iff \pi < 1/2,\ \frac{\partial q_u}{\partial \omega} < -\frac{\partial q_l}{\partial \omega} \iff \pi > 1/2.$$

Appendix B: Proofs

Proposition 1.

i. Incentive compatible precision levels always discriminate, i.e., $\underline{\nu}(q_i, \omega_i, \pi) \in \mathbb{R} \implies \underline{\nu}(q_i, \omega_i, \pi) \geq \tilde{\nu}(q_i, \pi)$.

ii. The incentive compatibility constraint to acquire information becomes less stringent as $|q_i - \pi|$ *decreases. If* $q_i \leq \pi \leq 1/2$ *or* $1/2 \leq \pi \leq q_i$*, then* $\frac{\partial \nu_i}{\partial |q_i - \pi|} \leq 0$*. Moreover, if prior beliefs in the economy are distributed according to a density* q *with full support on* $[0, 1]$*, then there exists a positive probability mass of traders with prior* q_i *around* π *who acquire information.*

iii. *The incentive compatibility constraint becomes less stringent as ω_i increases for sufficiently large ω_i. If, for some $\omega_i > 0$, informative signals are acquired, then larger ω_i strictly increases signal precision until $\nu_i = \hat{\nu}$.* ($\frac{\partial \nu_i}{\partial \omega_i} \geq 0$)

Proof. i. Setting $q_i = \pi$, both $\underline{\nu}$ and $\tilde{\nu}$ take value 1/2, the global minimum. The former increases faster as q_i deviates, hence the result follows.

ii. From *i.*, $\underline{\nu}_i \geq \tilde{\nu}$ for q_i close to π for any $\omega_i > 0$. Since $\underline{\nu}_i(\pi = 1/2) = 1/2$ and $\underline{\nu}_i(\pi)$ is continuous with finite slope, there exists a neighborhood around $q_i = \pi$ such that $\underline{\nu}_i < \hat{\nu}$ whenever $\hat{\nu} > 1/2$.

Setting $q_i = \pi$, $\overline{\nu}_i - \underline{\nu}_i = \sqrt{\omega_i}/2 > 0$. Since $\overline{\nu}_i, \underline{\nu}_i$ are continuous in q_i, there exists a neighborhood around $q_i = \pi$ such that $[\underline{\nu}_i, \overline{\nu}_i] \cap [\tilde{\nu}_i, \hat{\nu}] \neq \varnothing$. Because the density of q_i is strictly positive everywhere, there is a positive probability mass of traders with q_i such that information acquisition is incentive compatible.

The term within the square root of (3) or (4) is strictly decreasing in q_i if $q_i > \pi$ and strictly increasing if $q_i < \pi$. If the term is negative, then information acquisition is not incentive compatible. Thus, the incentive compatibility constraint becomes less stringent as q_i approaches π.

Because ν_i^* is increasing in q_i whenever $\pi < 1/2$ and decreasing whenever $\pi > 1/2$, the chosen precision is increasing in q_i for $q_i \leq \pi \leq 1/2$ and decreasing for $1/2 \leq \pi \leq q_i$. Together with the effect on the stringency of the IC, $\frac{\partial \nu_i}{\partial |q_i - \pi|} \leq 0$ if $1/2 \leq \pi \leq q_i$ or $q_i \leq \pi \leq 1/2$.

iii. I first show that the range of incentive compatible effort levels is nondecreasing for sufficiently large $\omega_i > 0$. Incentive compatible positive effort levels exist iff the term within the square root in (3) or (4) is nonnegative. This term becomes positive for sufficiently large ω_i, because all squared ω_i terms have positive sign, while one (linear) ω_i term has a negative sign.

To be shown: choice set $[\underline{\nu}_i, \overline{\nu}_i] \cap [\tilde{\nu}_i, \hat{\nu}]$ is nondecreasing and non-empty as $\omega_i \to \infty$. Since *i.* states $\tilde{\nu}_i \leq \underline{\nu}$, it remains to be shown that sufficiently large ω_i implies $\underline{\nu}_i < \hat{\nu}$.

Claim: $\lim_{\omega \to \infty} \underline{\nu}_i \to \tilde{\nu}$. To see this,[5] apply a Taylor approximation for $\underline{\nu}_i$ (see (3), (4)), which is of the form $y - \sqrt{y^2 - z} \approx z/(2y)$ and holds for $y >> z$. Then, as $\omega_i \to \infty$,

$$\frac{z}{2y} = \frac{1 + q_i \omega_i / \pi}{\omega_i [q_i/\pi + (1 - q_i)/(1 - \pi)] + 4} \to \frac{q_i/\pi}{q_i/\pi + (1 - q_i)/(1 - \pi)} = \tilde{\nu}_i.$$

The above approximation is obtained by Taylor expanding $\sqrt{1 - x}$ at 1, $\sqrt{1 - x} \approx 1 - 1/2x$ (omitting higher order terms, as these vanish asymptotically for x small).

[5] I am grateful to Ron Gordon for this idea.

Then, using $x = z/y^2$,

$$y - \sqrt{y^2 - z} = y - y\sqrt{1 - z/y^2} \approx y - y\left(1 - \frac{z}{2y^2}\right) = \frac{z}{2y}.$$

The Taylor approximation becomes arbitrarily accurate as $\omega_i \to \infty$, since $1 - z/y^2 \to 1$. It is easy to verify that $\underline{\nu}_i$ converges from above. Therefore, the range of incentive compatible effort levels is nondecreasing for sufficiently large ω_i, so the choice set does not decrease. Consequently, ν_i increases as ω_i increases until $\nu_i = \hat{\nu}$, because ν_i^* is strictly increasing in ω_i. □

Proposition 2.

i. *There exists an equilibrium price π^* fulfilling definition 1,*

ii. *which is implicitly defined as the fixed point of* (5).

iii. *Assuming homogeneous endowment ($\omega_i = \omega$), the equilibrium is unique for ω such that $\hat{\nu}$ is binding,*

iv. *and unique for ω such that $\hat{\nu}$ is non-binding if ω is small.*

Proof. i. The maximization problem for each i is continuous, with maximizer a_i, b_i from the compact set $\omega_i \geq a_i + b_i \geq 0$. Applying Berge's maximum theorem, the demand correspondence is upper hemi-continuous (uhc) in the parameters, non-empty and compact-valued. This implies that aggregate demand is uhc in the parameters (Aumann, 1976). Given risk-neutrality, demand is multi-valued if and only if $p_i = \pi$. In this case, demand is the budget line, which is a convex set. Writing (5) as mapping $[0,1] \to [0,1]$,

$$\int_0^{\varpi} \left(1 - Q(t_u) + \int_{t_l}^{t_u} \min\{\nu^*(q_i, \pi), \hat{\nu}\} dQ(q_i)\right) \omega_i dW(\omega_i) / \int_0^{\varpi} \omega_i dW(\omega_i) = \pi,$$

Kakutani's fixed point theorem guarantees the existence of π fulfilling definition 1.

ii. In equilibrium, the number of contracts demanded for either outcome must be identical, so that investments from the losers are identical to earnings of the winners. Without loss of generality, suppose $\theta = 1$. Then, denoting the mass of categorical traders buying A-contracts at price π^* by C_A and the mass of discriminatory A traders by D_A (similarly for B), the equilibrium price fulfills

$$\begin{aligned} \int_0^{\varpi} \omega_i \frac{C_A + D_A}{\pi^*} dW(\omega_i) &= \int_0^{\varpi} \omega_i \frac{C_B + D_B}{1 - \pi^*} dW(\omega_i) \\ \iff \int_0^{\varpi} \omega_i (C_A + D_A)(1 - \pi^*) - \omega_i (C_B + D_B)\pi^* dW(\omega_i) &= 0 \\ \iff \int_0^{\varpi} \omega_i (C_A + D_A - \pi^*) dW(\omega_i) &= 0, \end{aligned}$$

because $C_A + C_B + D_A + D_B = 1$, and because risk neutral agents invest the entire endowment ω_i. Replacing the C_A, D_A terms, this is the desired expression.

iii. Rewriting, $\int_{t_l(\pi)}^{t_u(\pi)} \min\{\nu^*, \hat{\nu}\} dQ = \hat{\nu}[Q(t_u) - Q(t_l)]$, since $\hat{\nu}$ is binding. A change in π changes this term by $\hat{\nu}[t_u' q(t_u) - t_l' q(t_l)]$, and changes $1 - Q(t_u)$ by $-t_u' q(t_u)$. Hence, $\hat{\nu}[t_u' q(t_u) - t_l' q(t_l)] - t_u' q(t_u) < 0$, since $\hat{\nu} \leq 1$ and $t_u', t_l' > 0$. Thus, the probability mass investing in A is strictly decreasing in π, so there is a unique π^* fulfilling equality (5).

iv. For ω with non-binding $\nu_i^* = \frac{q_i \omega}{8\pi} + \frac{(1-q_i)\omega}{8(1-\pi)} + 1/2$, the precision has to be constrained for $\pi \to 1$ and $\pi \to 0$ nonetheless. Thus

$$\frac{\partial \alpha/8}{\partial \pi} = \begin{cases} -\frac{\omega}{8\pi^2} & \text{if } \nu_i^* < \hat{\nu} \\ 0 & \text{if } \nu_i^* \geq \hat{\nu}. \end{cases}, \quad \frac{\partial \beta/8}{\partial \pi} = \begin{cases} \frac{\omega}{8(1-\pi)^2} & \text{if } \nu_i^* < \hat{\nu} \\ 0 & \text{if } \nu_i^* \geq \hat{\nu}. \end{cases}$$

The demand for contract A (for the case $\theta = 1$) is

$$\begin{aligned} 1 - Q(t_u) + \int_{t_l}^{t_u} \frac{q_i \alpha + (1-q_i)\beta + 4}{8} dQ(q_i) &= 1 - Q(t_u) + \left[\frac{\beta}{8} + \frac{1}{2}\right](Q(t_u) - Q(t_l)) \\ + \frac{\alpha}{8}\left[t_u Q(t_u) - t_l Q(t_l) - \int Q(q_i) dq_i\right] &- \frac{\beta}{8}\left[t_u Q(t_u) - t_l Q(t_l) - \int Q(q_i) dq_i\right] \end{aligned}$$

using integration by parts. Differentiating with respect to π using Leibniz' integral rule and simplifying,

$$\begin{aligned} \frac{\partial}{\partial \pi} =& q(t_u) t_u' \left[-\frac{1}{2} + t_u \frac{\alpha}{8} + (1 - t_u)\frac{\beta}{8}\right] + q(t_l) t_l' \left[-\frac{1}{2} - \frac{\alpha}{8} t_l - (1 - t_l)\frac{\beta}{8}\right] \\ &+ \frac{\partial \beta/8}{\partial \pi}\left[Q(t_u)(1 - t_u) - Q(t_l)(1 - t_l) + \int Q(q_i) dq_i\right] \\ &+ \frac{\partial \alpha/8}{\partial \pi}\left[t_u Q(t_u) - t_l Q(t_l) - \int Q(q_i) dq_i\right]. \end{aligned} \tag{7}$$

Now, $-\frac{1}{2} + t_u \frac{\alpha}{8} + (1 - t_u)\frac{\beta}{8} \leq 0$, since $\frac{1}{2} \leq \nu^* = \frac{q_i \alpha + (1-q_i)\beta + 4}{8} \leq \hat{\nu} \leq 1 \; \forall q_i \in [0, 1]$. Therefore, the first two terms are negative ($t_u', t_l' > 0$). The remaining two terms have a positive upper bound of $-\frac{\partial \alpha/8}{\partial \pi} + 2\frac{\partial \beta/8}{\partial \pi}$, because $0 \leq \int_{t_l}^{t_u} Q(q_i) dq_i \leq 1$ and $0 \leq Q(t_u) - Q(t_l) \leq 1$. This upper bound can be decreased arbitrarily by reducing ω. For $\omega \to 0$, the entire expression converges to $-1/2(q(t_u) t_u' + q(t_l) t_l') < 0$, so demand is decreasing in response to a decreasing π for small ω, and uniqueness follows. □

Proposition 3. *Comparative statics if demand is non-increasing in the price, ω is homogeneous and $\theta = 1$.*

i. Any change in the prior distribution from Q to R such that $Q(q_i) > R(q_i) \; \forall q_i \in (0, 1)$ increases the market price π^* if $\hat{\nu}$ is binding, unless all agents in the Q-economy already acquire information.

ii. *With priors* $q_i \sim \mathcal{N}(\mu, \sigma^2)$ *and* $\omega < (1-\pi^*)4$*, an upward shift of* μ *increases the equilibrium price* π^* *for each* μ, σ *if* $\pi^* \leq 1/2$*.*

iii. *With priors* $q_i \sim \mathcal{N}(\mu, \sigma^2)$ *and endowment* $\omega < (1-\pi^*)4$ *sufficiently small, an upward shift of* μ *increases the equilibrium price* π^* *for each* μ, σ *and* π^**.*

iv. *With priors from the triangular distribution, an increase in the mode value* d *leads to an increase of the market prediction* π^* *if* $d < t_l$ *or* $d > t_u$*.*

v. *Higher endowment increases information in the economy as measured by*

$$D_A = \int_{t_l}^{t_u} \min\{\nu^*(q_i), \hat{\nu}\} dQ(q_i),$$

i.e., more traders obtain information of higher precision.

vi. *For any distribution of priors with non-increasing density* q *the market prediction* π^* *increases in response to a non-binding endowment increase if* $\pi^* \geq 1/2$*.*

vii. *With priors* $q_i \sim \mathcal{N}(\mu, \sigma^2)$ *the market prediction* π^* *increases in response to an endowment increase for* $\pi^* \geq 1/2$*,* $\hat{\nu}$ *non-binding and* $|\mu - t_l| < |\mu - t_u|$*.*

Proof. i. Not all agents acquire information, which implies $t_l \in (0,1)$ or $t_u \in (0,1)$. The demand for the A-contract is $1 - Q(t_u) + \int_{t_l}^{t_u} \hat{\nu} q(q_i) dq_i = 1 - Q(t_u) + \hat{\nu}[Q(t_u) - Q(t_l)]$. For any change $Q(q_i) > R(q_i)\ \forall q_i \in (0,1)$, the demand changes by $Q(t_u) - R(t_u) + \hat{\nu}[R(t_u) - Q(t_u) - R(t_l) + Q(t_l)] > 0$, since $[1-\hat{\nu}]Q(t_u) \geq [1-\hat{\nu}]R(t_u)$ and $Q(t_l) \geq R(t_l)$ with at least one strict inequality, as t_l or t_u is in the interior. The increase of the equilibrium price is implied by the implicit function theorem and non-increasing demand.

ii. Defining $E(q_i) := \int_{t_l}^{t_u} q_i \phi((q_i - \mu)/\sigma)/\sigma dq_i$, the definition of the equilibrium price (5) (assuming non-binding $\hat{\nu}$) with normally distributed priors is

$$1 + \Phi\left(\frac{t_u - \mu}{\sigma}\right)\left[\frac{\beta}{8} - \frac{1}{2}\right] + E(q_i)\left[\frac{\alpha}{8} - \frac{\beta}{8}\right] + \Phi\left(\frac{t_l - \mu}{\sigma}\right)\left[-\frac{\beta}{8} - \frac{1}{2}\right] - \pi^* = 0.$$

By the implicit function theorem,

$$\frac{d\pi^*}{d\mu} = -\frac{\partial(5)/\partial\mu}{\partial(5)/\partial\pi^*} = -\frac{\overbrace{-\frac{1}{\sigma}\phi\left(\frac{t_u - \mu}{\sigma}\right)\frac{\beta - 4}{8}}^{>0 \text{ since } \omega<(1-\pi^*)4} \overbrace{+E(q_i)'\frac{\alpha - \beta}{8}}^{\leq 0 \text{ if } \pi^* \geq 1/2} \overbrace{-\frac{1}{\sigma}\phi\left(\frac{t_l - \mu}{\sigma}\right)\frac{-\beta - 4}{8}}^{>0}}{\underbrace{\partial(5)/\partial\pi^*}_{<0}} > 0. \tag{8}$$

The weight of the $E(q_i)'$ term can be scaled down arbitrarily by lowering ω, while the weight of the ϕ-terms in the numerator converge to $1/(2\sigma)$, showing iii. The $E(q_i)'$ term is always non-positive for $\pi^* \leq 1/2$, thus showing ii.

iii. Shown in ii.

iv. For case $d < t_l$

$$C_A + D_A = \frac{(1-t_u)^2}{1-d} + \int_{t_l}^{t_u} \min\{\nu^*, \hat{\nu}\} \frac{2(1-q_i)}{(1-d)} dq_i > 0,$$
$$\frac{\partial(C_A + D_A)}{\partial d} = \frac{(1-t_u)^2}{(1-d)^2} + \int_{t_l}^{t_u} \min\{\nu^*, \hat{\nu}\} \frac{2(1-q_i)}{(1-d)^2} dq_i > 0.$$

Consider next the case $d \geq t_u$

$$C_A + D_A = 1 - \frac{t_u^2}{d} + \int_{t_l}^{t_u} \min\{\nu^*, \hat{\nu}\} \frac{2q_i}{d} dq_i,$$
$$\frac{\partial(C_A + D_A)}{\partial d} = \frac{t_u^2}{d^2} - \int_{t_l}^{t_u} \min\{\nu^*, \hat{\nu}\} \frac{2q_i}{d^2} dq_i \geq \frac{t_u^2}{d^2} - \hat{\nu}\left[\frac{t_u^2}{d^2} - \frac{t_l^2}{d^2}\right] > 0,$$

because $\hat{\nu} < 1$. The comparative static is implied by non-increasing demand and the implicit function theorem.

v. Using Leibniz' integral rule,

$$\frac{\partial D_A}{\partial \omega} = \int_{t_l}^{t_u} \left(\frac{q_i}{8\pi^*} + \frac{(1-q_i)}{8(1-\pi^*)}\right) q(q_i) dq_i$$
$$+ \frac{t_u\alpha + (1-t_u)\beta + 4}{8} q(t_u)t_u' - \frac{t_l\alpha + (1-t_l)\beta + 4}{8} q(t_l)t_l' > 0,$$

because $\partial t_l/\partial \omega < 0$ and $\partial t_u/\partial \omega > 0$.

vi. Using Leibniz' integral rule,

$$\frac{\partial(5)}{\partial \omega} = \overbrace{\int_{t_l}^{t_u} \left(\frac{q_i}{8\pi^*} + \frac{(1-q_i)}{8(1-\pi^*)}\right) q(q_i) dq_i}^{>0}$$
$$+ \underbrace{\frac{t_u\alpha + (1-t_u)\beta - 4}{8}}_{<0, >-1/2} q(t_u)t_u' - \underbrace{\frac{t_l\alpha + (1-t_l)\beta + 4}{8}}_{<1, >1/2} q(t_l)t_l'.$$

Since $t_u' > 0, t_l' < 0$, but $t_u' < -t_l' \iff \pi^* > 1/2$ when $\hat{\nu}$ is not binding, and since density q is non-increasing, the expression is positive.

vii. Whenever μ is closer to t_l than to t_u, the former point has higher density by symmetry. The rest follows from the proof of the previous result. □

Lemma 4. *If the economy consists of two groups with* $\pi_1^* \neq \pi_2^*$*, and demand is non-increasing in the price for both groups, then the unique market price* π^* *fulfills* $\min\{\pi_1^*, \pi_2^*\} < \pi^* < \max\{\pi_1^*, \pi_2^*\}$.

Proof. If $\theta = 1$, the equilibrium price is

$$\begin{aligned} &\gamma\omega_1\left[1 - Q_1(t_u(\pi^*, \omega_1)) + \int_{t_l(\omega_1)}^{t_u(\omega_1)} \min\{\nu^*(\pi^*, \omega_1, q_i), \hat{\nu}\} dQ_1(q_i) - \pi^*\right] \\ &+(1-\gamma)\omega_2\left[1 - Q_2(t_u(\pi^*, \omega_2)) + \int_{t_l(\omega_2)}^{t_u(\omega_2)} \min\{\nu^*(\pi^*, \omega_2, q_i), \hat{\nu}\} dQ_2(q_i) - \pi^*\right] = 0. \end{aligned} \tag{9}$$

Suppose without loss of generality $\pi_1^* > \pi_2^*$. By monotonicity of demand, the LHS when evaluated at $\pi = \pi_2^*$ is positive and therefore cannot be a market clearing price. Since demand for the A-contract is weakly decreasing in π in both groups, the LHS is strictly decreasing in π. Thus, the unique market clearing price must fulfill $\pi^* > \pi_2^*$. Similarly, the LHS evaluated at $\pi = \pi_1^*$ is positive, which implies $\pi^* < \pi_1^*$. □

Proposition 6. *Suppose the economy consists of two endowment groups* $\omega_1 > \omega_2$*, and demand is non-increasing in the price for both groups. If* $\pi_1^* > \pi_2^*$*, then* $\pi^* > I^*$*. Similarly,* $\pi_1^* < \pi_2^*$ *implies* $\pi^* < I^*$.

Proof. For $\theta = 1$, I^* is defined as

$$\begin{aligned} &\gamma\left[1 - Q_1(t_u(I^*, \omega_1)) + \int_{t_l(I^*,\omega_1)}^{t_u(I^*,\omega_1)} \min\{\nu^*(I^*, \omega_1, q_i), \hat{\nu}\} dQ_1(q_i) - I^*\right] \\ &+(1-\gamma)\left[1 - Q_2(t_u(I^*, \omega_2)) + \int_{t_l(I^*,\omega_2)}^{t_u(I^*,\omega_2)} \min\{\nu^*(I^*, \omega_2, q_i), \hat{\nu}\} dQ_2(q_i) - I^*\right] = 0. \end{aligned}$$

Setting the price equal to π^*, which fulfills $\pi_1^* > \pi^* > \pi_2^*$ (Lemma 4), the first term within brackets is positive, while the second is negative. Since the relative weight of the first term is reduced compared to (9), because $\omega_1 > \omega_2$, while the relative weight of the second is larger, the LHS is negative. Consequently, the unique solution I^* must be smaller than π^*. □

References

Ali, M. (1977): "Probability and utility estimates for racetrack bettors," *Journal of Political Economy*, 85, 803–815.

Arrow, K., R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J. Ledyard, S. Levmore, R. Litan, P. Milgrom, F. Nelson, et al. (2008): "The promise of prediction markets," *Science*, 320, 877–878.

Aumann, R. J. (1976): "An elementary proof that integration preserves uppersemicontinuity," *Journal of Mathematical Economics*, 3, 15–18.

Barlevy, G. and P. Veronesi (2000): "Information acquisition in financial markets," *Review of Economic Studies*, 67, 79–90.

Berg, J., F. Nelson, and T. Rietz (2008): "Prediction market accuracy in the long run," *International Journal of Forecasting*, 24, 285–300.

Berg, J. E. and T. A. Rietz (2006): "The Iowa Electronic Markets: Stylized facts and open issues," in *Information Markets: A New Way of Making Decisions*, ed. by R. W. Hahn and P. C. Tetlock, AEI-Brookings Joint Center for Regulatory Studies, 142–169.

Blume, L. and D. Easley (2006): "If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets," *Econometrica*, 74, 929–966.

Forsythe, R., F. Nelson, G. R. Neumann, and J. Wright (1992): "Anatomy of an experimental political stock market," *American Economic Review*, 82, 1142–1161.

Forsythe, R., T. Rietz, and T. Ross (1999): "Wishes, expectations and actions: a survey on price formation in election stock markets," *Journal of Economic Behavior & Organization*, 39, 83–110.

Gjerstad, S. (2005): "Risk aversion, beliefs, and prediction market equilibrium," Working Paper, University of Arizona.

Grossman, S. J. and J. E. Stiglitz (1980): "On the impossibility of informationally efficient markets," *American Economic Review*, 70, 393–408.

Hanson, R. and R. Oprea (2009): "A manipulator can aid prediction market accuracy," *Economica*, 76, 304–314.

Jullien, B. and B. Salanié (2000): "Estimating preferences under risk: The case of racetrack bettors," *Journal of Political Economy*, 108, 503–530.

Manski, C. (2006): "Interpreting the predictions of prediction markets," *Economics Letters*, 91, 425–429.

Ottaviani, M. and P. N. Sørensen (2009): "Surprised by the parimutuel odds?" *American Economic Review*, 99, 2129–2134.

——— (2010): "Noise, information, and the favorite-longshot bias in parimutuel predictions," *American Economic Journal: Microeconomics*, 2, 58–85.

——— (2013): "Price Reaction to Information with Heterogeneous Beliefs and Wealth Effects: Underreaction, Momentum, and Reversal," Working paper.

Page, L. and R. Clemen (2013): "Do prediction markets produce well calibrated probability forecasts?" *Economic Journal*, 123, 491–513.

Peress, J. (2004): "Wealth, information acquisition, and portfolio choice," *Review of Financial Studies*, 17, 879–914.

Quah, J. (2003): "The law of demand and risk aversion," *Econometrica*, 71, 713–721.

Radner, R. (1979): "Rational expectations equilibrium: Generic existence and the information revealed by prices," *Econometrica*, 655–678.

Rothschild, D. M. and R. Sethi (2013): "Trading Strategies and Market Microstructure: Evidence from a Prediction Market," Working paper.

Sauer, R. (1998): "The economics of wagering markets," *Journal of Economic Literature*, 36, 2021–2064.

Servan-Schreiber, E., J. Wolfers, D. M. Pennock, and B. Galebach (2004): "Prediction markets: Does money matter?" *Electronic Markets*, 14, 243–251.

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

Spann, M. and B. Skiera (2009): "Sports forecasting: a comparison of the forecast accuracy of prediction markets, betting odds and tipsters," *Journal of Forecasting*, 28, 55–72.

Thaler, R. and W. Ziemba (1988): "Parimutuel betting markets: Racetracks and lotteries," *Journal of Economic Perspectives*, 2, 161–174.

Tziralis, G. and I. Tatsiopoulos (2007): "Prediction markets: An extended literature review," *Journal of Prediction Markets*, 1, 75–91.

Verrecchia, R. E. (1982): "Information acquisition in a noisy rational expectations economy," *Econometrica*, 50, 1415–1430.

Wolfers, J. and E. Zitzewitz (2004): "Prediction Markets," *Journal of Economic Perspectives*, 18, 107–126.

——— (2006): "Interpreting prediction market prices as probabilities," *NBER Working Paper*, 12200.