Bayesian markets to elicit private information

Aurélien Baillon^{a,1}

^aErasmus School of Economics, Erasmus University Rotterdam, 3000 DR Rotterdam, The Netherlands

Edited by Jose A. Scheinkman, Columbia University, New York, NY, and approved June 13, 2017 (received for review March 3, 2017)

Financial markets reveal what investors think about the future, and prediction markets are used to forecast election results. Could markets also encourage people to reveal private information, such as subjective judgments (e.g., “Are you satisfied with your life?”) or unverifiable facts? This paper shows how to design such markets, called Bayesian markets. People trade an asset whose value represents the proportion of affirmative answers to a question. Their trading position then reveals their own answer to the question. The results of this paper are based on a Bayesian setup in which people use their private information (their “type”) as a signal. Hence, beliefs about others’ types are correlated with one’s own type. Bayesian markets transform this correlation into a mechanism that rewards truth telling. These markets avoid two complications of alternative methods: they need no knowledge of prior information and no elicitation of metabeliefs regarding others’ signals.

prediction markets | economic incentives | truth telling | mechanism design | Bayesianism

When trading in a market or submitting a price to an auctioneer, people reveal the extent to which they value a good. In finance, option and future markets reveal investors’ beliefs. Although such revelations of tastes and beliefs originally were a by-product of regular markets, they have led to prediction markets whose primary goal is to reveal beliefs. In such markets, people can buy or sell a simple asset whose value is 1 unit (e.g., \$10) if a specified event occurs and is 0 otherwise. Buying a unit at a given price is a bet that the event is more likely than people think to be on average. The resulting market price reveals aggregate expectations. Prediction markets have been used by various public organizations and companies (1, 2); for example, social scientists use this type of market to predict the replicability of experiments (3, 4). Prediction markets also have been shown to outperform polls in predicting election results (5). Unfortunately, they can only be used for events whose occurrence can be objectively verified. When collecting personal data such as opinions and self-assessed measurements, objective verification often is conceptually or practically impossible.

This paper introduces Bayesian markets, which are designed to elicit private information in binary settings (yes-or-no questions) when objective verification is impossible. Bayesian markets rely on the assumption that the private information that people possess influences their belief about others. Such inference is justified by Bayesian reasoning, a widely used theory of rational reasoning (6). Answering yes (Y) provides information (a signal) that can be used to update one’s prior expectation about the proportion of Y answers. In Bayesian markets, the assets traded have a value determined by the proportion of Y answers to a given question, for example, “Are you satisfied with your life?” Using Bayesian reasoning, we can predict that people who actually answer Y will expect a higher asset value than those who answer no (N). Hence, there exists a range of prices for which Y people want to buy the asset and N people want to sell it. In other words, for any price in this range, Y people bet that the asset will be worth more than the price, and N people bet it will be worth less. Bayesian markets use the difference in betting behavior to provide incentives for people to tell the truth about unverifiable information.

In prediction markets, betting on an event reveals one’s beliefs about that event. In Bayesian markets, betting on how many people are satisfied with their lives, for instance, reveals the bettor’s beliefs about others’ life satisfaction, which in turn reveals the bettor’s own life satisfaction. Bayesian markets can also be used to elicit people’s opinions about an event far in the future, such as the very long-term consequences of climate change, for which prediction markets are not adequate.

Bayesian markets complement alternative methods proposed to elicit private information, which are the Bayesian truth serum (7), the peer prediction method (8), and their refinements (9–11). Bayesian markets share the same Bayesian setting with these methods, notably the assumption that agents have a common prior about the population. However, by using simple betting decisions instead of eliciting a probability or estimating metabeliefs, Bayesian markets are simpler and more transparent than these alternatives and are robust to certain deviations from the common prior assumption (unlike refs. 7–10). They are restricted to binary questions, however. After presenting the setup, the main result, and an extension to small samples, I discuss related literature and situations in which the required assumptions are satisfied.

The Agents and Their Information

There are n agents (referred to as “he” in the singular). Consider a question Q about the agents’ private information, with two possible answers $\{0, 1\}$ to choose from. The type $t_i \in \{0, 1\}$ of agent $i \in \{1, \dots, n\}$ corresponds to his truth, which is private information. The proportion of type 1 agents is denoted $\omega = (\sum_{i=1}^n t_i / n) \in [0, 1]$. Following the literature (7–13), I assume that it is common knowledge that all agents share a prior belief $f(\omega)$ describing how likely they would consider various proportions to be, had they not (yet) known their own type. Harsanyi (14) provided justifications of this common prior assumption.

It is also common knowledge that types are impersonally informative, as defined by Prelec (7): $f(\omega|t_i) = f(\omega|t_j)$ is equivalent to $t_i = t_j$. This property includes two aspects. First, types are impersonal. That is, all agents i with $t_i = 0$ have the same updated belief $f(\omega|t_i = 0)$, with expectations denoted $\bar{\omega}_0$, and all agents j

Significance

People’s private information can be revealed by the way in which they trade specifically designed assets in a new type of market. People trade an asset whose value is the proportion of affirmative answers to a question. Their trading position then reveals their own answer to the question. In Bayesian markets, people can be rewarded for telling the truth even when the truth is not verifiable. Bayesian markets are simpler and more transparent than alternative methods, avoiding the measurements of metabeliefs about others and prior beliefs.

Author contributions: A.B. wrote the paper.

The author declares no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹Email: baillon@ese.eur.nl.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1703486114/-DCSupplemental.

with $t_j = 1$ have the updated belief $f(\omega|t_j = 1)$, with expectations denoted $\bar{\omega}_1$. Hence, the agent type includes all of the non-common information. Second, types are informative [or “stochastically relevant” (8)]. If agent i is of type 1, then because of this signal, he will consider large proportions ω to be more likely than he did a priori, and types 0 will consider them less likely, so that $\bar{\omega}_1 > \bar{\omega}_0$.

Two convenience assumptions are made to keep the presentation of the main result simple: (i) There are infinitely many agents, and (ii) f does not assign probability mass to the degenerate distributions 0 and 1. That is, it is certain that there are strictly positive proportions of both types of agents. These two assumptions will be discussed after the main result.

The Market

A market is set up for Q , the question of interest. It is a one-shot market in which all agents can simultaneously participate. The agents are offered the possibility of trading an asset whose value v will be the proportion of agents reporting 1. It is common knowledge that agents maximize their subjective expected payoff (what they expect to earn based on their beliefs) and only participate in the market if they expect a strictly positive subjective expected payoff. Agent i first reports his answer $r_i \in \{0, 1\}$ to question Q . Then, a market price p is randomly, uniformly drawn from the unit interval. Agent i is asked whether he wants to trade at p . Specifically, he is asked whether he wants to buy if he reported 1, and whether he wants to sell if he reported 0. Implementing a random price in the market is a simple adaptation of a mechanism commonly used in experimental economics to elicit people’s valuations of goods and assets (15). It guarantees that the agents rely only on what they expect the value of the asset to be because they cannot influence the market price and there is nothing to learn from it.

Agents do not trade with each other but with a market maker (referred to as “she”). However, whether the market maker buys or sells with a given agent does depend on what the others do. If a majority (taken as strict in this paper) of the agents reporting 1 is willing to buy, then the market maker buys from all of the agents reporting 0 and willing to sell; otherwise, she does not trade with the agents reporting 0. Symmetrically, if a majority of the agents reporting 0 are willing to sell, then the market maker agrees to sell to any agent reporting 1. The market maker trades with both sides simultaneously. She ensures that all agents could trade if they wanted, even if the number of buyers and sellers differs. Although it is important for the result to come that the market maker’s actions are based on the agents’ decisions, all results of this paper and of [Supporting Information](#) would still hold if we replace majority by unanimity. The main result would also hold for any strictly positive threshold, for instance requiring one-third of agents willing to trade. However, from an implementation perspective, the majority rule is less sensitive to extreme behavior or the mistakes of a few agents than unanimity, and it is more natural than other thresholds.

Having collected all reports, the market maker determines the settlement (or liquidation) value v , which is defined as the proportion of 1 among the reports r_i . Because there are infinitely many agents, no one can influence the settlement value on his own. As in prediction markets, when the value of the asset is known, the market stops and payment occurs. Buyers of the asset receive the settlement value from the market maker and sellers give it to the market maker. Hence, agent i ’s net payoff is $v - p$ if he bought the asset, $p - v$ if he sold it, and 0 if he did not trade.

Main Result

The theorem below establishes that truthful reporting, defined as $r_i = t_i$ for all agents i , provides a Bayesian Nash equilibrium. This means that, given that everyone else is telling the truth, no agent

wants to lie. The next paragraph clarifies this statement by providing the proof of the theorem.

I first assume that all agents participate in the market. We will later see that all expected payoffs are strictly positive, ensuring that the agents indeed will participate. Consider agent i and assume that all other agents report the truth, which implies $v = \omega$. Type 1 agents become buyers and will agree to buy if the market price p is less than their expectation of the asset value $\bar{\omega}_1$. Similarly, type 0 agents become sellers and will sell if the market price exceeds $\bar{\omega}_0$. By assumption, it is certain that there are both types of agents and the market maker will trade with all agents who want to if a majority on each side wants to trade. Hence, trades will occur if and only if $\bar{\omega}_0 < p < \bar{\omega}_1$. What should agent i do? If his type is 1, he expects the asset to be worth $\bar{\omega}_1$. To make a profit as a seller, he would only sell if the market price exceeds $\bar{\omega}_1$, but no trades will occur at such a high price. However, he is willing to buy up to a market price of $\bar{\omega}_1$. He therefore expects to trade and make a profit if the market price is between $\bar{\omega}_0$ and $\bar{\omega}_1$. To reap this payoff, he must first report 1 and hence be truthful. He can then expect to receive $\int_{\bar{\omega}_0}^{\bar{\omega}_1} (\bar{\omega}_1 - p) dp = ((\bar{\omega}_1 - \bar{\omega}_0)^2 / 2) > 0$, that is, a payoff of $\bar{\omega}_1 - p$ for all market price p between $\bar{\omega}_0$ and $\bar{\omega}_1$. So-called mixed strategies (defined as reporting 1 with some probability and 0 otherwise) make no sense because they only decrease the probability of reaping the payoff. The proof that truth telling is best for a type 0 agent is symmetric, and the expected payoff is the same, thus proving the *Theorem*.

Theorem. In a Bayesian market, truthful reporting is a Bayesian Nash equilibrium.

Interestingly, Bayesian markets also prevent (degenerate) equilibria in which everyone reports the same answer. If all agents were to decide to report 0, then no trades would occur; there would be no buyers, and therefore, no strict majority of buyers willing to trade. All agents would end up with 0 and therefore would not participate. The same holds if all agents were to decide to report 1. Hence, there is no Bayesian Nash equilibrium in which agents participate and all report the same. The equilibrium in which no agents participate (and therefore all get 0) is dominated by participation and truthful reporting because the latter gives a strictly positive expected payoff.

As we have seen, all agents expect a strictly positive payoff. This expectation is only possible because the market maker subsidizes the market. At the equilibrium and if p falls between $\bar{\omega}_0$ and $\bar{\omega}_1$, she will pay $(\omega - p)$ to ωn buyers and receive it from $(1 - \omega)n$ sellers. Hence, the ex post cost will be either 0 or $n(2\omega - 1)(\omega - p)$, which is bounded below by $-(n/8)$ and above by n ([Supporting Information](#)). The maximum relative cost of 1 per agent can become problematic in an absolute sense for very large samples.

Until now, we assumed that both types were present with certainty, but we can relax this assumption and only exclude some degenerate cases. It is enough to require that there is no $x \in (0, 1)$ such that $f(0) + f(x) + f(1) = 1$. [Supporting Information](#) proves the theorem under this assumption. In the present setup, the market price is randomly determined from a uniform distribution. Any other distribution whose density is strictly positive for all values in $(0, 1)$ would work as well. A distribution ensuring that agents are often rewarded can help them to quickly learn the best strategy (16).

Adaptation to Small Samples

Although an infinite number of agents was assumed for mathematical simplicity, a simple modification makes Bayesian markets work even for small samples. A minimum of four agents is required. To prevent an agent’s own report from influencing the asset value he is betting on, the market maker simply buys and sells individualized assets whose value will be the proportion of 1 answers reported by three randomly selected other agents

(hence, excluding the agent's own report). This adaptation is inspired from peer prediction mechanisms (8–10).

The market maker provides the additional guarantee that the trade is canceled if the three other agents, whose answers an agent is betting on, are of the same type, that is, she will neither buy nor sell if $v=0$ or $v=1$. The agent then knows that if he trades, v is either $1/3$ or $2/3$. If he assumes everyone else is telling the truth, he expects $E(v|t_i) = (1 + \bar{\omega}_{t_i})/3$ (*Supporting Information*). Type 1 agents will therefore expect a higher asset value than type 0 agents, and the optimal strategy remains truthful reporting.

Related Literature

Prediction markets are popular for eliciting people's beliefs (1–5) because they only rely on simple betting decisions. However, they are restricted to beliefs about objectively verifiable events. Bayesian markets are as close as possible to prediction markets in terms of simplicity but extend their domain to unverifiable subjective data and beliefs about unverifiable events.

Since the 1980s, various Bayesian revelation mechanisms have been proposed (7–13). Some mechanisms were especially designed for data collection methods such as surveys. The peer prediction mechanism (8) assumes that the investigator actually knows the common prior (estimated, for instance, from previously collected data). An agent's answer is then used to update the prior and to reward the agent based on the performance of the obtained posterior in predicting what another agent answered. Other peer prediction mechanisms also require knowing (in practice: estimating) some forms of metabeliefs, for instance about type correlations between tasks (17). In practical implementations, the payment rule may not be transparent to the respondents because it is based on additional assumptions or estimations. Bayesian markets avoid this issue at the cost of asking for a betting decision. Moreover, in Bayesian markets, degenerate equilibria are dominated by truth telling, unlike in many peer prediction mechanisms (17).

The Bayesian truth serum (7) and its refinements (9–11), like Bayesian markets, do not require the investigator to know the prior beliefs of the agents. This feature makes the Bayesian truth serum implementable in previously unasked questions. It has been used, for instance, in recognition tasks (18) or in evaluating the prevalence of scientifically questionable practices (19). The Bayesian truth serum and its refinements work as follows. Agents must first answer a question but then must also report a prediction of the rate of each possible answer. Predicting the rate of an answer can be a difficult task for the general population, because it involved probabilistic reasoning. The agents' score is then based on two components: whether their own answer occurred more often than predicted on average and whether their prediction was accurate. This accuracy is typically measured with a logarithmic or quadratic distance measure (scoring rule). The payment with the Bayesian truth serum is difficult for respondents to understand, to the point that it was not even explained in most applications. Bayesian markets have three advantages over the Bayesian truth serum and its refinements: (i) simple betting decisions are more natural to laypeople than scoring rules; (ii) people do not have to report a prediction (a value) in Bayesian markets, but only have to make a binary decision (trade or not); (iii) the final payment is more transparent, and anyone can compute it. Point ii makes clear that Bayesian markets are not a special case of the Bayesian truth serum and its refinements because they require less information.

Unlike the Bayesian truth serum, Bayesian markets are as yet restricted to binary settings. This restriction can be problematic even for binary questions if, for instance, the population consists of experts and nonexperts. Then there are four types (expert and yes, expert and no, nonexpert and yes, nonexpert and no), violating the assumptions of Bayesian markets. The Bayesian truth

serum solves this problem by asking a question with the four types as possible answers (7). Because markets only have buyers and sellers, they cannot accommodate questions with more than two answers. Although a question with more than two answers can be reformulated as a series of binary questions, having multiple markets related to the same initial question can create dependencies that affect the agents' optimal strategy. There is a solution if a researcher is only interested in the proportion of each answer and has access to a sufficiently large sample. For a question with k possible answers, she can split the sample into $k-1$ groups. Each group participates in only one market related to the question, "Is your answer h ?" for h from 1 to $k-1$. Each agent only considers one market, and the results of the present paper may apply. The researcher can then recombine the full distribution of answers.

On a technical level, the presence of the market maker prevents agents from relating the likelihood of trading to the actual ω and thus avoids the no-trade theorem (20). Indeed, in a pure Bayesian setup, if an agent knows that he trades with an agent from the opposite type, then they both learn the type of the opposite party and end up with the same posterior. The market maker prevents this learning from occurring and therefore has an important role in keeping the agents' type private until the settlement value is determined. This is why Bayesian markets, as proposed in this paper, are one-shot markets in which all agents act simultaneously. One may wonder whether sequential or dynamic versions are possible. Such versions would require solving the challenge of preserving some degree of privacy, a similar problem as the one highlighted by Cummings et al. (21). If agents' types are constant, simply revealing the number of buyers and sellers immediately reveals v and therefore ω (at the equilibrium). Once ω is known, there is no profit to expect from trading because everyone agrees and there is no longer a role for Bayesian markets.

Discussion of the Assumptions

Bayesian markets perfectly reveal beliefs whenever the assumptions are perfectly satisfied. In applications, the assumptions will be violated to some extent. Then, the mechanism of Bayesian markets induces people to move closer to speaking the truth to the extent that its assumptions hold approximately, where the common prior assumption and the sensitivity of information are most important. Noise and biases remain to the extent that assumptions are violated. I next discuss violations of the various assumptions in detail.

The theoretical literature on Bayesian mechanisms discussed above relies heavily on the existence of a common prior. Relaxing the common prior assumption is not an easy task and requires specifying how much agents know about others' beliefs about types. A refinement of the Bayesian truth serum (11) allows priors to vary; however, agents should still expect their posterior to be closer to the posterior of the other agents sharing their type than to any opposite-type agents. In *Supporting Information*, it is shown that Bayesian markets are robust to deviations from the common prior under a similar assumption as in ref. 11. Such a robustness result is important because the common prior assumption is unlikely to perfectly hold in practice. Still, Bayesian markets rely on beliefs being strongly influenced by agent types. Questions that the agents never asked themselves (and their relatives) are more likely to satisfy this requirement. For practical implementation, it is desirable that respondents have little reason to believe that others have different priors.

There is ample evidence that people deviate from Bayesian updating in the psychology literature (22). The crucial implication of Bayesian updating for us is that posteriors are correlated with types ($\bar{\omega}_1 > \bar{\omega}_0$). First, deviations from Bayesian updating such as the confirmation bias (23) will not reverse this correlation. Second, and more importantly, correlations between one's

own truth and beliefs about others have been robustly observed in various domains of behavior, feelings, and thoughts (24, 25). Such correlations, initially seen as a bias and called the false-consensus bias, were later reconciled with Bayesianism (6). People can still improve their predictions of others' behavior by weighting their own behavior more heavily (26). Even if a bias exists, it points in the same direction as Bayesianism and therefore supports the mechanism on which Bayesian markets rely.

An example of applications in which Bayesian markets will not work well concerns questions about shameful behavior; in this regard, people may exhibit a different bias known as pluralistic ignorance. They then neglect the extent to which people deviate from social norms when predicting others' answers (27). Pluralistic ignorance concerns average beliefs, whereas the false consensus bias concerns beliefs of one group relative to another. The two phenomena may occur simultaneously (27), but it remains safer not to implement Bayesian markets in domains in which pluralistic ignorance has been identified. The belief assumptions of Bayesian markets would also be violated if some agents view themselves as fundamentally different from the others. Because of idiosyncratic tastes or specific political views, they may expect their answers to be negatively correlated with the answers of the others.

For the Bayesian market to work, agents should use Bayesian Nash equilibria. Empirical studies have found that people may deviate (28), although deviations are reduced by learning (16). Hence, it will be useful to explain the desirability of equilibrium to agents. For instance, an agent can be told that being the only one to deviate from truth telling will reduce his expected payoff.

As is common in games using language to coordinate, there is another Bayesian Nash equilibrium in which 1 can be interpreted as meaning 0 and conversely. Expected payoffs would remain unchanged. The problem with this approach is that agents can only adopt it if all others also do so in a coordinated action; however, such coordination is difficult to imagine. The natural first answer coming to mind in this game is the true answer. Agents are aware that this also occurs for other agents. Hence, truthful answering is the most natural procedure for coordination, making it a focal equilibrium (ref. 29; see also ref. 30). The exclusion of degenerate equilibria reduces the set of candidates for focal points, making truth telling the most convincing one. Hence, the *Theorem* implies that Bayesian markets induce truth telling.

By rewarding truth telling, Bayesian markets motivate respondents to think deeply about their answers. Bayesian markets provide incentives to tell the truth when agents do not have any other incentives. They reduce the effect of opposite incentives (interests to lie) if there are other incentives, but the outcome then is uncertain. In such cases, Bayesian markets are less effective (i) the more the respondents are suspicious that their answers may be used against them, (ii) when the information is so sensitive that the payoff does not compensate the cost of admitting the truth, and (iii) when the respondents expect others to report untruthfully for strategic reasons even if they themselves would not.

Truth-telling incentives show that a carefully considered and honest answer is important to the researcher. If the researcher fears that providing monetary incentives crowds out respondents' intrinsic motivation, the latter can be reinforced by transferring incentives to charity giving as in ref. 19.

Application Domains

In summary, Bayesian markets are especially useful for private information that is not sensitive but that does require effort from the respondents. In health studies and in social and human sciences, it is common to ask respondents to evaluate various aspects of their lives (health, well-being, life satisfaction, etc.) or to remember and report certain past actions. For such questions, respondents may not have any reason to deliberately lie but also have little reason to think deeply about each answer. Truth-telling incentives help motivate them to provide careful answers.

Bayesian markets can solve a central problem in experimental economics. Although providing monetary incentives is a cornerstone of experimental economics, experimenters often also collect nonincentivized survey data to complement the main choice data of their experiments. It is therefore salient that this survey part is not incentivized, thereby jeopardizing the quality of the data. Bayesian markets can help incentivize what is not incentivized so far, such as simple personality questionnaires.

Bayesian markets offer an alternative to prediction markets when the event to predict cannot be verified in the lifetime of the forecaster (for instance, whether extraterrestrial life will be discovered before the year 2200 or whether climate change will become a serious problem for future generations). Other events that cannot be observed and that therefore are unsuited for prediction markets are counterfactual events. In Bayesian markets, experts can be incentivized to predict, or rather conjecture, what could have occurred, such as whether an alternative action would have yielded a better outcome than the action that was implemented. Answering such a question requires effort, and if not too sensitive or strategic, is suited for Bayesian markets. Here, again, experts can be rewarded for reporting what they truly believe even when we can never check whether they are right.

Conclusion

Through Bayesian markets, researchers interested in the answer to a question can become market makers to incentivize truth telling. The market maker subsidizes the market and ensures that all agents can expect a strictly positive profit. Currently, social scientists and companies often simply reward agents for completing surveys. Using Bayesian markets, they can reallocate resources to rewarding truth telling in a simpler and more transparent way than would be possible with alternative solutions.

ACKNOWLEDGMENTS. I am very grateful to Drazen Prelec, Benjamin Tereick, and Peter Wakker for their remarks on previous drafts. This research was made possible by European Research Council Starting Grant 638408 BayesianMarkets.

- Arrow KJ, et al. (2008) Economics. The promise of prediction markets. *Science* 320: 877–878.
- Majumder SR, Diermeier D, Rietz TA, Amaral LAN (2009) Price dynamics in political prediction markets. *Proc Natl Acad Sci USA* 106:679–684.
- Dreber A, et al. (2015) Using prediction markets to estimate the reproducibility of scientific research. *Proc Natl Acad Sci USA* 112:15343–15347.
- Camerer CF, et al. (2016) Evaluating replicability of laboratory experiments in economics. *Science* 351:1433–1436.
- Berg JE, Nelson FD, Rietz TA (2008) Prediction market accuracy in the long run. *Int J Forecast* 24:285–300.
- Dawes RM (1989) Statistical criteria for establishing a truly false consensus effect. *J Exp Soc Psychol* 25:1–17.
- Prelec D (2004) A Bayesian truth serum for subjective data. *Science* 306:462–466.
- Miller N, Resnick P, Zeckhauser R (2005) Eliciting informative feedback: The peer-prediction method. *Manage Sci* 51:1359–1373.
- Parkes DC, Witkowski J (2012) A robust Bayesian truth serum for small populations. *Proceedings of the 26th AAAI Conference on Artificial Intelligence* (Association for the Advancement of Artificial Intelligence, Palo Alto, CA), pp 1492–1498.
- Radanovic G, Faltings B (2013) A robust Bayesian truth serum for non-binary signals. *Proceedings of the 27th AAAI Conference on Artificial Intelligence* (Association for the Advancement of Artificial Intelligence, Palo Alto, CA), pp 833–839.
- Radanovic G, Faltings B (2014) Incentives for truthful information elicitation of continuous signals. *Proceedings of the 28th AAAI Conference on Artificial Intelligence* (Association for the Advancement of Artificial Intelligence, Palo Alto, CA), pp 770–776.
- Crmer J, McLean RP (1988) Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56:1247–1257.
- Johnson S, Pratt JW, Zeckhauser RJ (1990) Efficiency despite mutually payoff-relevant private information: The finite case. *Econometrica* 58:873–900.
- Harsanyi JC (1968) Games with incomplete information played by Bayesian players, part III. The basic probability distribution of the game. *Manage Sci* 14:486–502.

15. Becker GM, DeGroot MH, Marschak J (1964) Measuring utility by a single-response sequential method. *Behav Sci* 9:226–232.
16. Erev I, Roth AE (2014) Maximization, learning, and economic behavior. *Proc Natl Acad Sci USA* 111:10818–10825.
17. Shnayder V, Agarwal A, Frongillo R, Parkes D (2016) Informed truthfulness for multi-task peer prediction. *Proceedings of the HCOMP Workshop on Mathematical Foundations of Human Computation* (Proceedings of the 2016 ACM Conference on Economics and Computation, ACM, New York), pp 179–196.
18. Weaver R, Prelec D (2013) Creating truth-telling incentives with the bayesian truth serum. *J Mark Res* 50:289–302.
19. John LK, Loewenstein G, Prelec D (2012) Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychol Sci* 23:524–532.
20. Milgrom P, Stokey N (1982) Information, trade and common knowledge. *J Econ Theory* 26:17–27.
21. Cummings R, Pennock DM, Vaughan JW (2016) The possibilities and limitations of private prediction markets. *Proceedings of the 2016 ACM Conference on Economics and Computation* (Association for Computing Machinery, New York), pp 143–260.
22. Pohl R, ed (2004) *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory* (Psychology Press, New York).
23. Oswald ME, Grosjean S (2004) Confirmation bias. *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*, ed Pohl R (Psychology Press, New York), Chap 4.
24. Ross L, Greene D, House P (1977) The “false consensus effect”: An egocentric bias in social perception and attribution processes. *J Exp Soc Psychol* 13:279–301.
25. Gary M, Miller N (1987) Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychol Bull* 102:72–90.
26. Hoch SJ (1987) Perceived consensus and predictive accuracy: The pros and cons of projection. *J Pers Soc Psychol* 53:221–234.
27. Prentice DA, Miller DT (1996) Pluralistic ignorance and the perpetuation of social norms by unwitting actors. *Adv Exp Soc Psychol* 28:161–209.
28. Goeree JK, Holt CA (2001) Ten little treasures of game theory and ten intuitive contradictions. *Am Econ Rev* 91:1402–1422.
29. Schelling TC (1960) *The Strategy of Conflict* (Harvard Univ Press, Cambridge, MA).
30. Myerson RB (2013) *Game Theory* (Harvard Univ Press, Cambridge, MA), Chap 3.