

A market to read minds

Aurélien Baillon

Erasmus University Rotterdam

6 April 2016

Abstract: Financial markets reveal what investors think about the future. Nowadays, markets are also used to predict the results of the next US presidential election. But could markets make people reveal personal information that only they know? For instance, their possibly secret thoughts or hidden actions? This paper shows how to design such markets, called Bayesian markets. People trade an asset whose value is the proportion of yes answers to a given question. Then their trading positions reveal their own answer. Bayesian markets can transform the way social scientists, survey companies, and governments collect data.

When trading on a market or submitting a price to an auctioneer, people reveal how much they value a good. In finance, option and future markets reveal investors' beliefs. Revealing tastes or beliefs was originally a 'by-product' of regular market functioning, but it has led to the design of prediction markets whose only goal is to reveal people's beliefs. On such markets, people can buy or sell a simple asset whose value is 1 unit (e.g., \$10) if a specified event occurs and 0 otherwise. The resulting market price reveals the aggregate of people's expectations. Prediction markets have been used by various public organizations and large companies (1). Economists used them to predict the replicability of experiments (2). They have been shown to outperform polls in predicting election results (3). Yet, they are restricted to beliefs about events that can be objectively verified.

Obviously, there is no such thing as an objective verification of someone's feelings or personal thoughts. Could another form of markets be used to read minds (or hearts) in such cases?

This paper introduces such markets, called Bayesian markets, which can be used for any type of opinion or personal information, especially what people never revealed about themselves and what only they know. On these markets, assets are traded whose value is the proportion of Yes (Y) answers to a given question. For instance, "have you ever said 'I love you' to someone without meaning it?" From the psychological literature, we can predict that people who would actually answer Y will expect a higher asset value than those who would answer No (N). Such a correlation between one's own truth and beliefs about others has been robustly observed in various domains concerning behavior, feelings, and thoughts (4,5). It can be explained by a common theory for rational reasoning called Bayesian reasoning (6). Answering Y provides information (a signal) that can be used to update one's prior expectation about the proportion of Y answers.

Bayesian markets make use of the link between own truth and expectations about others' truths. In a nutshell, they work as follows. Agents (referred to as "he" in singular) first report an answer (Y or N) and then a bid price for the asset if they answer Y or an ask price if they answer N. A market maker (or bookmaker, referred to as "she") computes the average bid price and the average ask price from the submissions and uses these as her own bid and ask prices. Next a market price between 0 and 1 is randomly drawn, as typically done in experimental economics (7). If it falls between an agent's ask price and the market maker's bid price or between the market maker's ask price and an agent's bid price, the corresponding trade occurs. Finally, the settlement value of the asset is computed from the collected answers and buyers receive this value from the market maker, who receives it from the sellers. As shown below, truth-telling is best in such a market. Hence, Bayesian markets can be used to read minds.

The following paragraphs formally define the market and present the main results. Consider any Y/N question about agents' private information. The type $t_i \in \{Y, N\}$ of agent i corresponds to his own truth, his private information. The proportion of Y agents

is denoted ω . Following the literature (8-11), I assume that it is common knowledge that all agents share a prior belief $f(\omega)$ describing how likely they would consider various proportions ω to be a priori, before they knew their type.¹ They then use their type as a signal to derive their posterior belief about the population $f(\omega|t_i)$ (for agent i). If agent i is of type Y then, because of that signal, he will consider large proportions ω more likely than he did a priori, and types N will consider them less likely.² It implies that beliefs about others are correlated with own type, as commonly found in psychology (3-4).

Further, agents with the same type have the same posterior belief. Hence the type of an agent comprises all the non-common information that they have. For simplicity, this analysis assumes that there are only two types, and this is an aspect to be generalized in future work.³ Denoting by $\bar{\omega}_k$ what agents of type k expect ω to be, we obtain $\bar{\omega}_Y > \bar{\omega}_N$. Two convenience assumptions are made to keep the presentation of the main results simple. There are infinitely many agents and the extreme proportions 0 and 1 are impossible. That is, it is certain that there are nonzero proportions of both types of agents.

The agents are offered the possibility to trade an asset whose value will be the proportion of agents reporting Y . 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. I first assume that they all participate. We will later see that the expected payoffs are indeed strictly positive.

Agent i reports $r_i \in \{Y, N\}$. If he reports N , he becomes a seller of the asset and must submit an ask price a_i . If he reports Y , he becomes a buyer and submits a bid price b_i . The market maker defines her ask and bid prices as the average submitted prices \bar{a} and \bar{b} .⁴ The market price p is randomly, uniformly⁵ drawn from the unit interval. If $b_i \geq p \geq \bar{a}$, then agent i buys the asset from the market maker at p . Similarly, if $a_i \leq p \leq \bar{b}$, then agent i sells the asset to the market maker at the market price. Having collected all

¹ See (12) for a justification of the common prior assumption, should it only be for the agent to be able to solve the problem at hands.

² We assume stochastic relevance (11), which can be written as $f(\omega|t_i) = f(\omega|t_j) \Rightarrow t_i = t_j$.

³ $t_i = t_j \Rightarrow f(\omega|t_i) = f(\omega|t_j)$.

⁴ With a default ask price 1 if there are no sellers and a default bid price 0 if there are no buyers.

⁵ Any other distribution whose density is strictly positive for all values in (0,1) would work as well.

reports, the market maker determines what is called the settlement value v , being the proportion of Y among the reports r_i . Buyers of the assets receive the settlement value from the sellers. Hence, agent i 's payoff is $v - p$ if he bought the asset, $p - v$ if he sold it, and 0 otherwise.

The first result demonstrated in the supplementary online material (SOM) states that if an agent submits an ask price or a bid price on a Bayesian market, then this price is the agent's subjective expectation of the settlement value, given his type. Such a result is not surprising. Implementing a random price on the market is a simple adaptation of a mechanism commonly used in experimental economics to elicit people's valuation of good and assets (7).

The second result of the SOM establishes that truthful reporting, defined as $r_i = t_i$ for all agents i , is a Bayesian Nash equilibrium. In other words, if an agent thinks everyone else is reporting the truth, then so should he. Consider agent i and assume that all other agents are reporting the truth. It implies $v = \omega$. Y agents become buyers and submit the bid price $\bar{\omega}_Y$, leading to $\bar{b} = \bar{\omega}_Y$. Similarly, N agents become sellers and submit the ask price $\bar{\omega}_N$, yielding $\bar{a} = \bar{\omega}_N$.

What should agent i do? If he is a Y agent, he expects the asset to be worth $\bar{\omega}_Y$, which is also the market maker's bid price \bar{b} . This does not leave room to make a profit when selling. However, he is willing to pay $\bar{\omega}_Y$, which is more than the market maker's ask price $\bar{\omega}_N$. So the agent expects a positive payoff as a buyer. To reap this payoff, he must report Y, hence being truthful. He can then expect to receive $\frac{(\bar{\omega}_Y - \bar{\omega}_N)^2}{2} > 0$. The proof for an N agent is symmetric and the expected payoff is the same.⁶

THEOREM 1: On a Bayesian market, truthful reporting is a Bayesian Nash equilibrium.

⁶ Only pure strategies (reporting either 0 or 1) have been explored here. Mixed strategies (reporting 1 with a probability that he chooses and 0 otherwise) cannot be the best response to truthful reporting of all other agents because they only decrease the probability of reaping the payoffs.

The next theorem is a strengthening of Theorem 1. Assume that Y agents do not always tell the truth but are simply more likely than N agents to report Y.⁷ Then the best response in such a case is still to tell the truth.

THEOREM 2: On a Bayesian market, truthful reporting is optimal as soon as Y agents are more likely to report Y than N agents.

Theorem 2 implies a remarkable property of Bayesian markets: equilibria are very *stable*. If agents expect that others sometimes make mistakes or that some agents might not understand or not be convinced by the mechanism, truthful reporting remains optimal. Moreover, Theorem 2 has a useful practical consequence: the market maker may ensure anonymity by implementing a random mechanism sometimes changing the reports. For instance, with 50% chance, a random report will be recorded instead of the agent's actual report. Truthful reporting remains optimal and no one can know whether a report was actually given by the agent or randomly given by an independent device, thereby protecting anonymity.

An infinite number of agents was assumed for mathematical elegance but a simple modification makes Bayesian markets work for even small samples. Only a minimum of 4 agents is required. To prevent an agent's own report from influencing the asset value he is betting on, the market maker may simply buy and sell 'individualized' assets, whose value will be the proportion of Y reports among three randomly selected 'other' agents (hence, excluding the agent's own report).⁸ Bayesian markets are also robust to deviations from the common prior assumption, so long as agents' expectations remain closer to the average expectations of the other agents sharing their type than to the expectations of agents with opposite type.⁹ Bayesian markets rely on beliefs being strongly influenced by agents' types. This is why they are particularly well-suited for

⁷ By symmetry, N agents are also more likely than Y agents to report N.

⁸ This adaptation is inspired from peer prediction mechanisms (11,13,14). All previous results are proven for $n \geq 4$ in the SOM.

⁹ See formal definition and proof in the SOM.

secrete or taboo information. In such cases, agents' own type is the only relevant piece of information they can use.

The market maker ensures that all agents can trade, even though the number of buyers and sellers may differ.¹⁰ By doing so, she subsidizes the market and makes sure that all agents will expect a strictly positive profit and will be willing to participate. In practice, her role is played by the researcher, social scientist, or investigator, who is interested in the answer to the question and wants to reward respondents for telling the truth.

From the 1980s on, various Bayesian revelation mechanisms have been proposed (8-11). Only one of them, called Bayesian truth-serum (10,13,14,16,17), does not require for the investigator to fully know the prior beliefs of the agents. Bayesian markets have the same property. The market maker only knows *how* the agents think (how they update the prior), not *what* they think (what the prior itself is). This feature made the Bayesian truth-serum implementable in marketing studies (17) or to evaluate the prevalence of scientifically questionable practices (16), but it is difficult to understand for respondents (to the point that it was even not explained in most applications so far). By contrast, participants in Bayesian markets can understand how the payment relates to what they do, as on prediction markets. Yet, unlike prediction markets, Bayesian markets can be applied to all sorts of personal information.

One may expect that markets make agents act less moral (18) and thus less truth telling. In Bayesian markets, to the contrary, the optimal selfish behavior and optimal moral behavior, understood as truth telling, align thanks to a new version of Adam Smith's invisible hand: the invisible hand of the Bayesian market maker.

¹⁰ A technical point: the market maker avoids that agents learn from each other, which would lead to no trade at all according to the no-trade theorem (15).

References

1. K.J. Arrow, R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J.O. Ledyard, S. Levmore, R. Litan, P. Milgrom, F.D. Nelson, G.R. Neumann, M. Ottaviani, T.C. Schelling, R.J. Shiller, V.L. Smith, E. Snowberg, C.R. Sunstein, P.C. Tetlock, P.E. Tetlock, H.R. Varian, J. Wolfers, E. Zitzewitz. The Promise of Prediction Markets. *Science* **320**, 877-878 (2008).
2. C. Camerer, A. Dreber, E. Forsell, T.H. Ho, J. Huber, M. Johannesson, M. Kirchler, J. Almenberg, A. Altmeld, T. Chan, E. Heikensten, F. Holzmeister, T. Imai, S. Isaksson, G. Nave, T. Pfeiffer, M. Razen, H. Wu. Evaluating replicability of laboratory experiments in economics. *Science* 10.1126/science.aaf0918 (2016).
3. J.E. Berg, , F.D. Nelson, T.A. Rietz. Prediction market accuracy in the long run. *International Journal of Forecasting* **24**, 285-300 (2008).
4. L. Ross, D. Greene, P. House. The “false consensus effect”: An egocentric bias in social perception and attribution processes. *Journal of experimental social psychology* **13**, 279-301 (1977).
5. M. Gary, N. Miller. Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychological Bulletin* **102**, 72 (1987).
6. R.M. Dawes. Statistical criteria for establishing a truly false consensus effect. *Journal of Experimental Social Psychology* 25, 1-17 (1989).
7. G.M. Becker, M.H. de Groot, J. Marschak. Measuring Utility by a Single-Response Sequential Method. *Behavioral Science* **9**, 226–232 (1964).
8. J. Crémer, R.P. McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 1247-1257 (1988).
9. S. Johnson, J.W. Pratt, R.J. Zeckhauser. Efficiency despite mutually payoff-relevant private information: The finite case. *Econometrica*, 873-900 (1990).
10. D. Prelec. A Bayesian Truth Serum for Subjective Data. *Science* **306**, 462-466 (2004).
11. N. Miller, P. Resnick, R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science* **51**, 1359-1373 (2005).

12. J.C. Harsanyi. Games with Incomplete Information Played by "Bayesian" Players, Part III. The Basic Probability Distribution of the Game. *Management Science* **14**, 486–502 (1968).
13. J. Witkowski, D.C. Parkes. A Robust Bayesian Truth Serum for Small Populations. *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, 1492-1498 (2012).
14. G. Radanovic, B. Faltings. A robust bayesian truth serum for non-binary signals. *Proceedings of the 27th AAAI Conference on Artificial Intelligence*, 833-839 (2013).
15. P. Milgrom, N. Stokey. Information, trade and common knowledge. *Journal of Economic Theory* **26**, 17-27 (1982).
16. L.K. John, G. Loewenstein, D. Prelec. Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological science* **23**, 524-532 (2012).
17. R. Weaver, D. Prelec. Creating Truth-telling Incentives with the Bayesian Truth Serum. *Journal of Marketing Research* **50**, 289-302 (2013).
18. A. Falk, N. Szech. Morals and markets. *Science* **340**, 707-711 (2009).

Supplementary online material of

“A market to read minds”

Aurélien Baillon

Erasmus University Rotterdam

6 April 2016

1 Main results from the paper

For mathematical simplicity, types Y and N are replaced by 1 and 0.

Consider a question Q with two possible answers $\{0,1\}$. An agent, typically denoted $i \in \{0, \dots, n\}$, has a *type* $t_i \in \{0,1\}$, which corresponds to his own truth. The proportion of agents with type 1 is $\omega \equiv \frac{\sum_{i=1}^n t_i}{n} \in [0,1]$.

As in the paper, I first assume that it is common knowledge that all agents share a prior $f(\omega)$ (called *common prior*) satisfying:

- *Stochastic relevance*: $f(\omega|t_i) = f(\omega|t_j) \Rightarrow t_i = t_j$.
- *Impersonal signaling*: $t_i = t_j \Rightarrow f(\omega|t_i) = f(\omega|t_j)$.

An agent uses his type as a signal to derive his posterior belief about the population $f(\omega|t_i)$. With these assumptions, types are “impersonally informative” (11), i.e. $f(\omega|t_i) = f(\omega|t_j)$ is equivalent to $t_i = t_j$.

Bayes rule implies $E(\omega|t_i = 1) = \frac{\int_0^1 \omega^2 f(\omega) d\omega}{\int_0^1 \omega f(\omega) d\omega} = \frac{E(\omega^2)}{E(\omega)} > E(\omega)$ (the strict inequality arises from stochastic relevance) and similarly $E(1 - \omega|t_i = 0) = \frac{E((1-\omega)^2)}{E(1-\omega)} = \frac{\int_0^1 (1-\omega)^2 f(\omega) d\omega}{\int_0^1 (1-\omega) f(\omega) d\omega} > 1 - E(\omega)$. Hence, with $\bar{\omega} \equiv E(\omega)$ and $\bar{\omega}_k \equiv E(\omega|t_i = k)$, we obtain $\bar{\omega}_1 > \bar{\omega} > \bar{\omega}_0$.

Two convenience assumptions are made to derive the main results of the paper. The number of agents, n , is infinite and *non-degeneracy* holds, defined as f not assigning probability mass on 0 or 1. These assumptions are not crucial and will be relaxed in Section 2. Together, they imply that agents know for sure that there will be nonzero proportions of each type. Moreover, agents cannot influence the value of the asset on their own.

We consider a Bayesian market for the question of interest, Q . On this market, the agents can trade an asset whose value will be the proportion of agents who participate in the market and report an answer 1 to Q . Agents maximize their *subjective expected payoff* (what they expect to earn based on their beliefs) and only participate in the market if they can expect a strictly positive subjective expected payoff (*strict participation constraint*). We first assume that they all participate and will verify in due course that the equilibrium of the market satisfies the strict participation constraint.

The *report* of agent i is denoted $r_i \in \{0,1\}$. If agent i reports $r_i = 0$, he becomes a seller of the asset and must submit an *ask price* a_i . If he reports $r_i = 1$, he becomes a buyer and submits a *bid price* b_i . The market maker defines her ask and bid prices as the average submitted prices \bar{a} and \bar{b} , with a default ask price 1 if there are no sellers and a default bid price 0 if there are no buyers. The *market price* p is randomly, uniformly drawn from the unit interval. If $b_i \geq p \geq \bar{a}$, then agent i buys the asset from the market maker at p . Similarly, if $a_i \leq p \leq \bar{b}$, then agent i sells the asset to the market maker at the market price. Having collected all reports, the market maker determines the *settlement value* v , with $v = \frac{\sum_{i=1}^n r_i}{n}$. Buyers of the assets receive the settlement value

from the sellers. Hence, agent i 's *payoff* is $v - p$ if he bought the asset, $p - v$ if he sold it, and 0 otherwise.

PROPOSITION 1: If an agent submits an ask price or a bid price on a Bayesian markets, then this price is the agent's subjective expectation of the settlement value given his type.

PROOF OF PROPOSITION 1:

Agents i expects the settlement value of the asset to be $E(v|t_i)$. As a buyer, his expected payoff would be $\int_{E(\bar{a}|t_i)}^{b_i} E(v|t_i) - p dp$. If $E(v|t_i) > \bar{a}$, maximizing $\int_{E(\bar{a}|t_i)}^{b_i} E(v|t_i) - p dp$ in b_i gives the first order condition $b_i = E(v|t_i)$. If $E(v|t_i) \leq \bar{a}$, the agent's expected payoff is zero and he will either not participate in the market or be a seller. In both cases, he would not submit a bid price. Hence, all submitted bid prices satisfy $b_i = E(v|t_i)$.

The expected payoff of a seller being $\int_{a_i}^{E(\bar{b}|t_i)} p - E(v|t_i) dp$. If $E(v|t_i) < E(\bar{b}|t_i)$, maximizing $\int_{a_i}^{E(\bar{b}|t_i)} p - E(v|t_i) dp$ in a_i gives the first order condition $a_i = E(v|t_i)$. Hence, the optimal ask price is $a_i = E(v|t_i)$ if the agent expects the value of the asset to be less than the market maker's bid price. If he expects the value of the asset to be equal to or more than the average bid price, then he will either not participate or be a buyer. Again, all submitted bid prices satisfy the condition $a_i = E(v|t_i)$. \square

THEOREM 1: On a Bayesian markets, truthful reporting is a BNE.

PROOF OF THEOREM 1:

Assume that all agents $j \neq i$ are reporting the truths ($r_j = t_j$), which leads to $v = \omega$. Hence, all agents with type 1, and only them, become buyers and they submit the bid price \bar{w}_1 , leading to $\bar{b} = \bar{w}_1$. Similarly, type 0 agents become sellers and submit an ask price equal to \bar{w}_0 , yielding $\bar{a} = \bar{w}_0$.

What should agent i do? First consider the case $t_i = 1$. He expects the asset to be worth $\bar{\omega}_1$, which is also the market maker's bid price \bar{b} . Consequently, he expects zero payoffs should he decide to be a seller. However, the previously established condition $\bar{\omega}_1 > \bar{\omega}_0$ guarantees $\bar{\omega}_1 > \bar{a}$. The agent expects a positive payoff by submitting his optimal bid price $b_i = \bar{\omega}_1$. To reap this payoff, he must submit $r_i = 1$, hence he must be truthful. By doing so, he can expect to receive $\int_{\bar{\omega}_0}^{\bar{\omega}_1} (\bar{\omega}_1 - p) dp = \frac{(\bar{\omega}_1 - \bar{\omega}_0)^2}{2} > 0$. The strict participation constraint is thus satisfied.

The agent may want to consider mixed strategies, i.e. reporting 1 with a probability that he chooses and 0 otherwise. Such a mixed strategy cannot be the best response to truthful reporting of all other agents because it would only decrease the probability of reaping the payoffs (he expects zero payoffs as a seller). The case $t_i = 0$ is symmetric. The expected payoff is also $\frac{(\bar{\omega}_1 - \bar{\omega}_0)^2}{2}$. \square

As in all games purely based on language, there is another BNE in which Y can be interpreted as meaning N and conversely. Expected payoffs would remain unchanged. Yet, such a situation is unlikely to happen as it contradicts natural language and therefore is no focal equilibrium. Note that it cannot serve as a way to avoid telling the truth: if everyone tells the opposite of the truth, then types are still revealed. Finally, if all agents were to decide to sell, then $v = 0$, which is the default bid price of the market maker and no one could expect any profit from this situation. The same holds if all agents were to decide to buy.

The second theorem considers the case in which each agent chooses the probability to report 1 (hence, implementing a mixed strategy) and when this probability is positively correlated with types. For the sake of simplicity, we assume that agents are agnostic whether this probability could be correlated with other individual characteristics.

THEOREM 2: On a Bayesian markets, truthful reporting is optimal as soon as type 1 agents are more likely to report 1 than type 0 agents.

PROOF OF THEOREM 2:

Assume that agents ($\neq i$) with $t_j = 1$ decide to report $r_j = 1$ with probability q_j while those with $t_j = 0$ report it with probability p_j (hence, telling the truth with probability $1 - p_j$). Denote q the average q_j and p the average p_j and assume $q > p$ (Y agents are more likely to answer 1 than N agents). The expected value of the asset becomes

$E(v|t_j) = q\bar{\omega}_{t_j} + p(1 - \bar{\omega}_{t_j})$. The condition $\bar{\omega}_1 > \bar{\omega}_0$ implies $E(v|t_j = 1) > E(v|t_j = 0)$. Hence, type 1 agents still expect a higher asset value than type 0 agents do.

Let \bar{a}_k and \bar{b}_k be the market maker's ask and bid prices that agents of type k expect.

From Proposition 1, we obtain

$$\bar{a}_k = \frac{[(1-q)\bar{\omega}_k[q\bar{\omega}_1+p(1-\bar{\omega}_1)]+(1-p)(1-\bar{\omega}_k)[q\bar{\omega}_0+p(1-\bar{\omega}_0)]}{(1-q)\bar{\omega}_k+(1-p)(1-\bar{\omega}_k)} \text{ and}$$

$$\bar{b}_k = \frac{[q\bar{\omega}_k[q\bar{\omega}_1+p(1-\bar{\omega}_1)]+p(1-\bar{\omega}_k)[q\bar{\omega}_0+p(1-\bar{\omega}_0)]}{q\bar{\omega}_k+p(1-\bar{\omega}_k)}.$$

These prices are weighted averages of type 1 agents' expected value of the asset and type 0 agents' expected value. Note that $q > p$ implies $\frac{q}{p} > \frac{1-q}{1-p}$. Multiplying both terms by

$\frac{\bar{\omega}_k}{1-\bar{\omega}_k}$ gives $\frac{q\bar{\omega}_k}{p(1-\bar{\omega}_k)} > \frac{(1-q)\bar{\omega}_k}{(1-p)(1-\bar{\omega}_k)}$, which means that \bar{b}_k gives a higher coefficient to the

high expectations $q\bar{\omega}_1 + p(1 - \bar{\omega}_1)$ than \bar{a}_k does. Hence, the market maker's bid price is always higher than her ask price.

Assume $t_i = 1$. Then $a_i = b_i = q\bar{\omega}_1 + p(1 - \bar{\omega}_1) > \bar{b}_1 > \bar{a}_1$. Consequently, the optimal strategy for agent i remains being a buyer, and therefore truthful reporting with probability 1 (because being a seller yields zero payoffs).

Assume now $t_i = 0$. Then $a_i = b_i = q\bar{\omega}_0 + p(1 - \bar{\omega}_0) < \bar{a}_0 < \bar{b}_0$. Consequently, the optimal strategy for agent i is selling, i.e., truthful reporting with probability 1 (because being a buyer yields zero payoffs). \square

2 Robustness

So far, for convenience, non-degeneracy was assumed but it can easily be relaxed. If it is possible that all agents are of the same type, the market maker may not be able to determine both her ask and her bid price. In such a case, she sets her default ask or bid prices to 1 and 0 respectively and no trade occurs. Consequently, agents will make their decisions about a_i , b_i , and r_i conditional on trades occurring, i.e. conditional on both types being present. Yet, the theorems still hold if types are stochastically relevant once the agents have ruled out the possibility that everyone is of the same type.¹

Formally, define $g(\omega|t_i)$ over $[0,1]$ by $g(0|t_i) = g(1|t_i) = 0$, and $g(\omega|t_i) = \frac{f(\omega|t_i)}{1-f(1|t_i)-f(0|t_i)}$ otherwise. It is the posterior of agent i conditional on both types being present. Instead of non-degeneracy of f and the equivalence between $f(\omega|t_i) = f(\omega|t_j)$ and $t_i = t_j$, simply assume the equivalence between $g(\omega|t_i) = g(\omega|t_j)$ and $t_i = t_j$. This assumption is weaker than the previous set of assumptions. Under non-degeneracy, $g = f$. Hence, non-degeneracy and the equivalence of $f(\omega|t_i) = f(\omega|t_j)$ and $t_i = t_j$ imply the equivalence between $g(\omega|t_i) = g(\omega|t_j)$ and $t_i = t_j$.

The number of agents, n , need not be infinite for Bayesian markets to work. From now on, we only assume $n \geq 4$. To prevent an agent's own report from influencing the asset value he is betting on, the market maker may simply buy and sell "individualized" assets, whose value will be the proportion of Y reports among three randomly selected "other" agents (hence, excluding the agent's own report). An agent can never influence the market maker's ask or bid price applied to him. For instance, as a buyer, he submits a bid price and therefore cannot influence the market maker's ask price, which is the only price relevant to him. The market maker provides the additional guarantee that the trade is cancelled if the three other agents are of the same type, i.e. she will neither buy nor sell

¹ Here is a counterexample, in which types are not stochastically relevant once the agents have ruled out the possibility that everyone is of the same type: If $\omega \in \{0.5,1\}$ and $f(1) = f(0.5) = 0.5$. Someone whose type is 0 knows for sure that $\omega = 0.5$. An agent of type 1 then faces two cases. Either there will only be buyers ($\omega = 1$) and his payoff will be zero, or there will be sellers as well ($\omega = 0.5$) but then everyone knows the true value of the asset. In such cases, no one would participate because there is no profit to be expected.

to agent i if $v = 0$ or $v = 1$. The agent then knows that if he trades, v is either $\frac{1}{3}$ or $\frac{2}{3}$. If he assumes everyone else is telling the truth, he expects $E(v|t_i) = \frac{1+\bar{\omega}_{t_i}}{3}$ (see below). Type 1 agents will therefore expect a higher asset value than type 0 agents and the optimal strategy remains truthful reporting.

Proposition 1 and Theorems 1 and 2 are now demonstrated for $n \geq 4$ and for the weaker assumptions discussed in the previous paragraphs.

PROOF OF PROPOSITION 1 FOR $n \geq 4$ AND ASSUMING THE EQUIVALENCE OF $g(\omega|t_i) = g(\omega|t_j)$ AND $t_i = t_j$:

The new setting does not change any of the arguments proving Proposition 1 in the main text. \square

PROOF OF THEOREM 1 FOR $n \geq 4$ AND ASSUMING THE EQUIVALENCE OF $g(\omega|t_i) = g(\omega|t_j)$ AND $t_i = t_j$:

In what follows, we denote $\bar{\omega}_k = E_g(\omega|t_i = k)$, the expectations of the agents based on g . We still have $\bar{\omega}_1 > \bar{\omega}_0$. Assume that all agents $j \neq i$ are reporting the truth ($r_j = t_j$), which leads to

$$\begin{aligned}
 E_g(v|t_j) &= \frac{\bar{\omega}_{t_j} (1 - \bar{\omega}_{t_j})^2 + 2\bar{\omega}_{t_j}^2 (1 - \bar{\omega}_{t_j})}{3\bar{\omega}_{t_j} (1 - \bar{\omega}_{t_j})^2 + 3\bar{\omega}_{t_j}^2 (1 - \bar{\omega}_{t_j})} \\
 &= \frac{(1 - \bar{\omega}_{t_j}) + 2\bar{\omega}_{t_j}}{3((1 - \bar{\omega}_{t_j}) + \bar{\omega}_{t_j})} \\
 &= \frac{1 + \bar{\omega}_{t_j}}{3}
 \end{aligned}$$

All agents with type $t_j = 1$, and only they, become buyers and they submit the bid price $\frac{1+\bar{\omega}_1}{3}$, leading to $\bar{b} = \frac{1+\bar{\omega}_1}{3}$. Similarly, type 0 agents become sellers and submit an ask price equal to $\frac{1+\bar{\omega}_0}{3}$, yielding $\bar{a} = \frac{1+\bar{\omega}_0}{3}$.

First consider the case $t_i = 1$. Agent i expects the asset to be worth $\frac{1+\bar{\omega}_1}{3}$, which is also the market maker's bid price \bar{b} . Consequently, he expects zero payoffs should he decide to be a seller. However, the previously established condition $\bar{\omega}_1 > \bar{\omega}_0$ guarantees $\frac{1+\bar{\omega}_1}{3} > \bar{a}$. The agent expect a positive payoff by submitting his optimal bid price $b_i = \frac{1+\bar{\omega}_1}{3}$. To reap this payoff, he must submit $r_i = 1$, hence he must be truthful. By doing so, he can expect to receive $\frac{(\bar{b}-\bar{a})^2}{2}$. As before, mixed strategies would only reduce the agent's chance to get the payoff. The case $t_i = 0$ is symmetric. \square

PROOF OF THEOREM 2 FOR $n \geq 4$ AND ASSUMING THE EQUIVALENCE OF $g(\omega|t_i) = g(\omega|t_j)$ AND $t_i = t_j$:

Assume that agents ($\neq i$) with $t_j = 1$ decide to report $r_j = 1$ with probability q_j while those with $t_j = 0$ report it with probability p_j (hence, telling the truth with probability $1 - p_j$). Denote q the average q_j and p the average p_j and assume $q > p$ (Y agents are more likely to answer 1 than N agents). The expected value of the asset becomes

$$E(v|t_j) = \frac{1+q\bar{\omega}_{t_j}+p(1-\bar{\omega}_{t_j})}{3}. \text{ The condition } \bar{\omega}_1 > \bar{\omega}_0 \text{ implies } E(v|t_j = 1) > E(v|t_j = 0).$$

Hence, type 1 agents still expect a higher asset value than type 0 agents do. From Proposition 1, we obtain

$$\bar{a}_k = \frac{1}{(1-q)\omega_k+(1-p)(1-\omega_k)} \left[(1-q)\omega_k \frac{1+q\bar{\omega}_1+p(1-\bar{\omega}_1)}{3} + (1-p)(1-\omega_k) \frac{1+q\bar{\omega}_0+p(1-\bar{\omega}_0)}{3} \right]$$

and

$$\bar{b}_k = \frac{1}{q\omega_k+p(1-\omega_k)} \left[q\omega_k \frac{1+q\bar{\omega}_1+p(1-\bar{\omega}_1)}{3} + p(1-\omega_k) \frac{1+q\bar{\omega}_0+p(1-\bar{\omega}_0)}{3} \right].$$

Note that $q > p$ implies $\frac{q}{p} > \frac{1-q}{1-p}$. Multiplying both terms by $\frac{\bar{\omega}_k}{1-\bar{\omega}_k}$ gives $\frac{q\bar{\omega}_k}{p(1-\bar{\omega}_k)} > \frac{(1-q)\bar{\omega}_k}{(1-p)(1-\bar{\omega}_k)}$, which means that \bar{b}_k gives a higher coefficient to the high expectations $\frac{1+q\bar{\omega}_1+p(1-\bar{\omega}_1)}{3}$ than \bar{a}_k does. Hence, all agents expect that the market maker's bid price is always higher than her ask price.

Assume $t_i = 1$. Then $a_i = b_i = \frac{1+q\bar{\omega}_1+p(1-\bar{\omega}_1)}{3} > \bar{b}_1 > \bar{a}_1$. Consequently, the optimal strategy for agent i remains being a buyer, and therefore truthful reporting with probability 1 (because being a seller yields zero payoffs).

Assume now $t_i = 0$. Then $a_i = b_i = \frac{1+q\bar{\omega}_0+p(1-\bar{\omega}_0)}{3} < \bar{a}_0 < \bar{b}_0$. Consequently, the optimal strategy for agent i is selling, and therefore truthful reporting with probability 1 (because being a buyer yields zero payoffs). \square

Finally, the assumption of a common prior can be relaxed as well. Let us denote by f_i the prior of agent i . Define $g_i(\omega|t_i)$ over $[0,1]$ by $g_i(0|t_i) = g_i(1|t_i) = 0$, and $g_i(\omega|t_i) = \frac{f_i(\omega|t_i)}{1-f_i(1|t_i)-f_i(0|t_i)}$ otherwise. It is the posterior of agent i conditional on both types being present. The functions g^1 and g^0 are defined the same way relative to f^1 and f^0 .

We now assume

- *Stochastic relevance*: $g_i(\omega|t_i = 0) \neq g_i(\omega|t_i = 1)$.
- *Impersonal signaling*: $t_i = t_j \Rightarrow g_i(\omega|t_i) = g_i(\omega|t_j)$.

Agent i 's expectation $E_{g_i}(\omega|t_i)$ is denoted $\bar{\omega}_i$. Since we relax the assumption that it is common knowledge that all agents share the same prior, we have to specify what each agent i knows about other agents' beliefs. I assume the following: he knows $\bar{\omega}^1$ and $\bar{\omega}^0$, where $\bar{\omega}^k$ is the average $\bar{\omega}_j$ of all agents j such that $t_j = k$. As can be seen in the proofs, it is the only piece of information that is relevant in choosing the optimal strategy.

We further assume:

- *Weak deviations from common prior:* $\bar{\omega}^1 > \bar{\omega}^0$. Further, if $t_i = 1$, $\bar{\omega}^1 - \bar{\omega}_i < \bar{\omega}_i - \bar{\omega}^0$ and if $t_i = 0$, $\bar{\omega}^1 - \bar{\omega}_i > \bar{\omega}_i - \bar{\omega}^0$.
- *Common knowledge* of stochastic relevance, impersonal signaling, and weak deviations from common prior.

Proposition 1 and Theorem 1 still hold with this new set of assumptions.

PROOF OF PROPOSITION 1 WITHOUT A COMMON PRIOR:

The new setting does not change any of the arguments proving Proposition 1. \square

PROOF OF THEOREM 1 WITHOUT A COMMON PRIOR:

As before, from Bayes rule, we obtain $E_{g_i}(\omega|t_i = 1) > E_{g_i}(\omega|t_i = 0)$. Further, from the assumption of weak deviations from common prior, we also have $\bar{\omega}^1 > \bar{\omega}^0$.

Consequently, weak deviations from common prior also imply $E_{g_i}(\omega|t_i = 1) - \bar{\omega}^0 > 0$ (otherwise, $E_{g_i}(\omega|t_i = 1)$ cannot be closer to $\bar{\omega}^1$ than to the lower $\bar{\omega}^0$) and $\bar{\omega}^1 - E_{g_i}(\omega|t_i = 0) > 0$.

The asset value will be the proportion of Y reports among three randomly selected “other” agents. The market maker provides the additional guarantee that the trade is cancelled if the three other agents are of the same type, i.e. she will neither buy nor sell to agent i if $v = 0$ or $v = 1$. If u is the proportion of type 1 reports, then, for any group of 3 randomly chosen agents, there is a $u(1 - u)^2$ chance that there is one type 1 report in the group and a $u^2(1 - u)$ chance that there are two type 1 reports in the group. Hence, conditional on the group having both types of reports, the value of the asset is

$$\frac{u(1-u)^2 + 2u^2(1-u)}{3u(1-u)^2 + 3u^2(1-u)} = \frac{1+u}{3}.$$

Assume that all agents $j \neq i$ are reporting the truth ($r_j = t_j$). Agent i expects the market maker’s bid price to be $\bar{b} = \frac{1+\bar{\omega}^1}{3}$ and the ask price to be $\bar{a} = \frac{1+\bar{\omega}^0}{3}$.

First consider the case $t_i = 1$. From Proposition 1, we know he would submit $a_i = b_i = \frac{1+\bar{\omega}_i}{3}$, which must be higher than \bar{a} and may be lower than \bar{b} . As a buyer, the agent can expect to make a profit of $\int_{\bar{a}}^{b_i} b_i - p dp = \frac{(b_i - \bar{a})^2}{2} > 0$. As a seller, the profit is either 0 if $a_i \geq \bar{b}$ or $\int_{a_i}^{\bar{b}} p - a_i dp = \frac{(\bar{b} - a_i)^2}{2}$ otherwise. The assumption of weak deviations from a common prior implies $\bar{b} - a_i < b_i - \bar{a}$. Consequently, it is optimal to report $r_i = 1$, which means being truthful. As before, mixed strategies would only reduce the agent's chance to get the payoff. The case $t_i = 0$ is symmetric. \square