Article:

Mann, RP orcid.org/0000-0003-0701-1274 and Helbing, D (2017) Optimal incentives for collective intelligence. Proceedings of the National Academy of Sciences, 114 (20). pp. 5077-5082. ISSN 0027-8424

https://doi.org/10.1073/pnas.1618722114

Optimal incentives for collective intelligence

Richard P. Mann[a,1] **and Dirk Helbing**[b]

[a]Department of Statistics, School of Mathematics, University of Leeds, Leeds LS2 9JT, UK.; [b]Computational Social Science, ETH Zurich, 8092 Zurich, Switzerland.

This manuscript was compiled on April 3, 2017

Collective intelligence is the ability of a group to perform more effectively than any individual alone. Diversity among group members is a key condition for the emergence of collective intelligence, but maintaining diversity is challenging in the face of social pressure to imitate one's peers. Through an evolutionary game-theoretic model of collective prediction we investigate the role incentives may play in maintaining useful diversity. We show that market-based incentive systems produce herding effects, reduce information available to the group and restrain collective intelligence. Therefore, we propose a new incentive scheme that rewards accurate minority predictions, and show that this produces optimal diversity and collective predictive accuracy. We conclude that real-world systems should reward those who have demonstrated accuracy when the majority opinion has been in error.

collective intelligence | game theory | diversity | markets

The financial crisis and its aftermath have reopened longstanding debates about the collective wisdom of our societal organisations [1–3]. Financial and prediction markets seem unable to foresee major economic and political upheavals such as the credit crunch or Brexit. This lack of collective foresight could be the result of insufficient diversity among decision-making individuals [4]. Diversity has been identified as a key ingredient of successful groups across many facets of collective behaviour [5–7]. It is a crucial condition for collective intelligence [6–10] that can be more important than the intelligence of individuals within a group [11]. As collective behaviour ultimately results from individual actions, incentives play a major role for diversity and collective performance [12, 13]. While most previous research has focused on explaining how collective intelligence emerges [14], less is known about how to optimise the wisdom of crowds in a quantitative sense.

Harnessing collective wisdom is important. Global systems of communication, governance, trade and transport grow rapidly in complexity every year. Many of these real world problems have a large number contributing factors. For example, predicting future economic fluctuations requires integrating knowledge about credit markets and supply chains across the world, as well as the ramifications of political developments in different countries and the shifting sentiments of individual investors and consumers. Political developments are themselves the result of many factors, both direct (e.g. political parties' strategies) and indirect (e.g. technological change). Scientific questions are also increasingly complex. For instance, building a complete model of an ecosystem requires bringing together expertise on many scales, from individual animal behaviour to complex networks of predation and codependency [15]. In each case, knowledge about the diverse contributing factors is dispersed. For these high-dimensional problems, it is becoming impossible for any single individual or agency to gather and process enough data to understand the entire system [16]. In many cases we do not even have full knowledge of what the potential causal factors are, let alone a full understanding of them.

Attention is therefore shifting towards distributed systems as a means of bringing together the local knowledge and private expertise of many individuals [12, 17]. In machine-learning, researchers have found that a pluralistic modelling approach maximises prediction accuracy [18]. In politics, the forecasts of prediction markets [19, 20] are now commonly reported alongside opinion polls during elections. Scientists are also turning to crowd-sourcing collective wisdom as a validation tool [21–23]. However, as highlighted by the failure of financial and prediction markets to foresee the results of recent elections in the UK and USA, collective wisdom is not a guaranteed property of a distributed system [2], partly due to herding effects [24, 25]. In science as well, the incentive structure undervalues diversity: low-risk projects with assured outcomes are more likely to be funded than highly novel or interdisciplinary work [26, 27]. Rewards for conformity with institutional cultures can severely limit useful diversity [28]. Previous work [29] has investigated mechanisms to elicit truthful minority views to counter herding effects in expressed opinion. This raises the question: how can minority viewpoints be fostered in the first place, so as to enhance diversity and its potential benefits for collective intelligence?

Here we analyse an evolutionary game-theoretic model of collective intelligence amongst unrelated agents motivated by individual rewards. We show that previously proposed incentive structures [13] are sub-optimal from the standpoint of collective intelligence, and in particular produce too little diversity between individuals. We propose a new incentive system that we term 'minority rewards', wherein agents are rewarded for expressing accurate minority opinions, and show that this produces stable, near-optimal collective intelligence at equilibrium. Our results demonstrate that common real-world

Significance Statement

Diversity of information and expertise amongst group members has been identified as a crucial ingredient of collective intelligence. However, many factors tend to reduce the diversity of groups, such as herding, groupthink and conformity. We show why the individual incentives in financial and prediction markets and in the scientific community reduce diversity of information, and how these incentives can be changed to improve the accuracy of collective forecasting. Our results therefore suggest ways to improve the poor performance of collective forecasting seen in recent political events, and how to change career rewards to make scientific research more successful.

RPM performed analysis. RPM and DH conceived the study and wrote the paper

Please declare any conflict of interest here.

[1]To whom correspondence should be addressed. E-mail: r.p.mann@leeds.ac.uk

reward structures are unlikely to produce optimal collectively intelligent behaviour, and we present a superior alternative that can inform the design of new reward systems.

Results

To investigate the effect of incentives on collective intelligence, we use an abstract model of collective information gathering and aggregation [13]. Complex outcomes are modelled as a result of n independent, causal factors. A large population of individual agents gather information in a decentralised fashion, each being able to pay attention to just one of these factors at any given time. Collective prediction is achieved by aggregation of individual predictions via simple voting. Agents are motivated by an incentive scheme that offers rewards for making accurate predictions. It is assumed that the accuracy of an individual's prediction can be judged after the event. We exclude cases where the ground truth is either never discoverable or where no such ground truth exists (for instance in questions regarding taste or voter preferences). Instead we consider questions such as the prediction of future events (which are known once they occur) or scientific questions (which may be resolved at some later point in time). For example, one might consider whether national GDP will rise above trend in the coming year, whether a certain party will win an election, or whether global temperatures will change by more than 1°C in the next decade. The proportion of agents attending to different sources of information evolve depending on the rewards they receive, where less successful agents tend to imitate their more successful peers.

Consider a binary outcome, Y, which is the result of many factors, $x_1, x_2, \ldots, x_n$. We model this outcome as the sign of a weighted sum of the contributing factors:

$$Y = \text{sign}\left(\sum_{i=1}^{n} \beta_i x_i\right). \quad [1]$$

For simplicity we assume that each contributing factor takes binary values, such that $Y, x_i \in \{-1, 1\}$, and that the values of these factors are uncorrelated (see SI Appendix for instances with correlated factors). Without loss of generality $\beta_i > 0$ for all factors.

An individual attending to factor i observes the value of x_i. Having observed the value of x_i, this individual then votes in line with that observation. Thus, if the proportion of individuals attending to factor i is ρ_i, the collective prediction $\hat{Y}$ is given by:

$$\hat{Y} = \text{sign}\left(\sum_{i=1}^{n} \rho_i x_i\right). \quad [2]$$

Collective accuracy, C, is the probability that the collective vote agrees with the ground truth, given the distribution, $\{\rho\}$, of agents attending to each factor:

$$C = P\left(\hat{Y} = Y \mid \{\rho\}\right). \quad [3]$$

The reward given to an agent for an accurate vote depends on the proportion of other correct votes in any given collective decision. Let z_i be the proportion of agents that will vote identically to those attending to factor i, i.e. the proportion of agents attending to factors whose value matches x_i: $z_i = \sum_{j=1}^{n} \rho_j \delta_{x_i, x_j}$, where δ is the Kronecker delta. Then the reward is determined by a function, $f(z_i)$, such that an agent receives a reward proportional to $f(z_i)$ if and only if their prediction is accurate. We will investigate three potential reward systems for deciding how each agent is rewarded for their accurate votes, the first two of which are taken from previous work by Hong *et al.* [13]. The first of these is 'binary rewards': agents receive a fixed reward if they make an accurate prediction, corresponding to the reward function $f(z_i) = 1$. The second is 'market rewards': a fixed total reward is shared equally amongst all agents who vote accurately, corresponding to the reward function $f(z_i) = 1/z_i$. This adds an incentive to be accurate when others are not, and closely mimics the reward system of actual prediction markets. Finally, we introduce 'minority rewards': agents are rewarded for an accurate prediction when fewer than half of the other agents also vote accurately, corresponding to the reward function $f(z_i) = 1 - H(z_i - 1/2)$, where $H(\cdot)$ is the Heavyside step function. This explicitly rewards agents who hold accurate *minority* opinions, and incentivises agents to be accurate on questions where the majority prediction is wrong.

The expected reward a player receives by attending to factor i is determined by the expected value of $f(z_i)$, conditioned on voting accurately (see eq. 8). Players adapt their behaviour in response to the rewards they and others receive. In alignment with previous evolutionary game theory work, we model changes in individual attention to factors as being the result of imitation; agents who are observed to be gaining greater rewards are imitated by those gaining fewer. This leads to the classic replicator equation [30], describing the evolution of the proportion of agents, ρ_i, that pay attention to factor i (see eq. 6)

We studied the behaviour of the model under the three incentive schemes described above. We initialised the model by assigning uniform proportions of agents to each factor, with values of β randomly drawn from a uniform distribution (the absolute scale of β does not affect the model). We followed the evolutionary dynamics described by the replicator equation until the population converged to equilibrium. This was repeated over a range of problem dimensionalities from $n = 3$ to $n = 10000$. Expected rewards were calculated either by exhaustive search over all possible values of $x_1, \ldots, x_n$ (for $n < 10$) or by using appropriate normal-distribution limits for large numbers of factors (see Methods).

Figure 1 shows how collective accuracy and diversity evolve towards equilibrium for the three rewards systems of binary, market and minority rewards in simulations with $n = 100$, $n = 1000$ and $n = 10,000$ independent factors. Note the logarithmic scale on the x-axis, to better illustrate the early evolution. For each reward system two initial allocations of agents' attention are used: (i) a uniform allocation to each factor; and (ii) an allocation where half of all agents attend to the single most important factor, with others allocated uniformly across the other factors. This demonstrates that the equilibrium distribution of attention is the same, no matter whether agents initially attend to arbitrary factors or initially favour the most obvious ones. The convergence time to equilibrium depends on the magnitude of rewards; in our simulations we normalise rewards such that the mean reward *per agent* is one at each time step.

Figure 2 shows how the resulting collective accuracy varies across problem dimensionalities from $n = 3$ to $n = 10000$ for the three different reward systems and for a uniform allocation of agents to factors. For simple problems ($n < 10$), all reward schemes produce high collective accuracy (over 90%). In these cases the strong predictive power of only one or two meaningful independent factors means that individual accuracy is high, and collective aggregation only leads to relatively small increases in collective accuracy. However, even for these 'small n' problems we observe that minority rewards outperform other schemes. The differences in collective accuracy become more substantial as n increases. As Figure 1 shows, these differences become apparent after only a few iterations, well before equilibrium is reached. Consistent with [13], we find that market rewards increase diversity and collective accuracy relative to binary rewards. However, collective accuracy under market rewards declines rapidly with increasing n, falling to $\sim 65\%$ for $n = 10000$. For comparison we also show the accuracy achieved under a uniform allocation of agents, which reaches a stable value of approximately 80% for large n. Market rewards therefore produce lower accuracy than a uniform allocation for all but the lowest values of n. In contrast, minority rewards lead to a far higher accuracy than any of the investigated alternative reward systems, regardless of system complexity, and achieve close to 100% accuracy up to $n = 10000$. Our mathematical analysis shows that minority rewards will continue to produce near-perfect accuracy for any problem size, if the population of agents is large enough (see SI Appendix). Our analysis of finite group sizes shows that minority rewards outperform other reward schemes for problem dimensions up to ten times bigger than the population size, assuming best-response dynamics (see SI Appendix, Fig. S1).

The different levels of collective accuracy across reward systems are a reflection of the differing equilibrium distributions of the proportion of agents attending to each factor. Minority rewards outperform both market rewards and unweighted approaches, as attention is automatically redirected if the collective prediction would otherwise be wrong; only those outcomes where the majority opinion is wrong contribute to agents' rewards. Under minority rewards the system converges towards a state where the number of agents paying attention to any factor is proportional to factor importance. This optimal distribution is both a stationary and stable state of the minority rewards system (see our mathematical analysis in the SI Appendix). Further analysis (see SI Appendix, Fig. S2) shows that varying the cutoff value for minority rewards (for example by rewarding those voting with less than 40% of the group, or 60%), invariably reduces collective accuracy. In Figure 3 we plot the equilibrium distribution for each reward system for a high-dimensional problem ($n = 10000$). Using binary rewards, almost all agents attend to the single most important factor. Under market rewards agents distribute themselves in proportion to the predictive value of the factors, but only among the top 10% of factors; 90% of factors receive essentially no attention at all (this proportion decreases as n increases, and is therefore larger for smaller values of n). By comparison, under minority rewards the proportion of agents paying attention to a factor is also proportional to its importance, but agents cover the full range of factors down to the least important ones, thereby providing more information to the group and improving predictions.The evolution of this distribution towards equilibrium is shown in detail in SI Appendix, Fig. S3.

Discussion

We proposed a novel reward system, minority rewards, that incentivises individual agents in their choice of which informational factors to pay attention to when operating as part of a group. This new system rewards agents both for making accurate predictions and for being in the minority of their peers or conspecifics. As such it encourages a balance between seeking useful information that has substantive predictive value for the ground truth, and seeking information that is currently under-utilised by the group. Conversely, where the collective opinion is already correct, no rewards are offered and therefore no agent is motivated to change their strategy. Over time, therefore, agents are motivated to change their behaviour only in ways that benefit collective accuracy.

The poor performance of market rewards relative to a uniform unweighted allocation for $n > 10$ shows that a market reward system incentivises herding behaviour and suppresses useful diversity, as illustrated by the equilibrium distribution in Figure 3b. This suggests that stock markets and prediction markets tend to systematically underweight a large pool of informational factors that are of limited predictive power individually, but which can contribute powerfully to aggregate predictions if agents can be persuaded to pay attention to them. This sheds doubt on the accuracy of existing markets as a tool for aggregating dispersed knowledge to predict future profits or events, and motivates further work on how to design collectively more accurate market mechanisms. The relatively high performance of uniform allocations of attention supports work showing that models with equally weighted predictors can match or even improve on more closely fitted prediction models [31, 32]. The inclusion of all relevant predictors is often more important than determining their appropriate weights in making predictions; too much diversity is less harmful than too little, especially for complex problems.

Incentives are a fundamental part of any effort to harness the potential of collective intelligence. In this paper we have presented evidence that rewarding accurate minority opinions can induce near-optimal collective accuracy within a model of collective prediction. Therefore, to maximise the collective wisdom of a group, we suggest that individuals should not be rewarded simply for having made successful predictions or findings, nor should a total reward be equally distributed amongst those who have been successful or accurate. Instead, rewards should be primarily directed towards those who have made successful predictions in the face of majority opposition from their peers. This can be intuitively understood as rewarding those who contribute information that has the potential to change collective opinion, since it contradicts the current mainstream view. In our model groups rapidly converge to an equilibrium with very high collective accuracy, after which the rewards for each agents become less frequent. We anticipate that once this occurs, agents would move on to new, unsolved problems. This would produce a dynamic system in which agents are incentivised not only to solve problems collectively, but also to address issues where collective wisdom is currently weakest. Future work should investigate how our proposed reward system can be best implemented in practice, from scientific career schemes, to funding and reputation systems [33], to prediction markets, and democratic procedures [34]. We suggest experiments to determine how humans respond to minority rewards, and further theoretical work to determine the effects of stochastic rewards, agent learning and finite group dynamics. In conclusion, how best to foster collective intelligence is an important problem we need to solve collectively.

Materials and Methods

Terminology. Throughout this paper we use the following conventions for describing probability distributions:

- $\mathbb{E}(x)$ denotes the expectation of x
- $\mathcal{N}(x; \mu, \sigma^2)$ denotes the normal probability density function with mean μ and variance σ^2, evaluated at x
- $\mathcal{N}(x; \mu, \Sigma)$ for vector-valued x and mu, and matrix Σ denotes the multi-variate normal probability density function with mean μ and covariance matrix Σ, evaluated at x
- $\Phi(x)$ denotes the standard normal cumulative probability distribution function with mean 0 and standard deviation 1.

Ground truth and voting. We consider a binary outcome, Y that is the result of many independent factors, $x_1, x_2, \ldots, x_n$ (for correlated factors see SI Appendix). We model this outcome as being determined by the sign of ψ: a weighted sum of the contributing factors.

$$Y = \text{sign}(\psi), \; \psi = \sum_{i=1}^{n} \beta_i x_i. \quad [4]$$

In computational implementation of this model we sample values of $\{\beta\}$ independently from a uniform distribution (the scale of which is arbitrary and does not influence the analysis). We assume without loss of generality that factors are ordered such that $\beta_i \geq \beta_{i+1}$, and further we normalise the values of the coefficients such that $\sum_{i=1}^{n} \beta_i = 1$, without affecting the value of Y. Our analytical results (see SI Appendix) do not depend on the exact distributon of $\{\beta\}$. Any sampling distribution for $\{\beta\}$ that has a finite moment of order $m, \; m > 2$ will obey the Ljapunov and Lindeberg conditions [35], guaranteeing convergence in distribution of ψ to a normal distribution, from which our results are obtained.

Each individual attends to one factor at a given time; an individual attending to factor i therefore observes the value of x_i. Having observed the value of x_i this individual then votes in line with that observation. The collective prediction, $\hat{Y}$ is given by the sign of the collective vote V, which is a sum over the contributing factors, weighted by the proportion of individuals attending to each factor:

$$\hat{Y} = \text{sign}(V), \; V = \sum_{i=1}^{n} \rho_i x_i. \quad [5]$$

Evolutionary dynamics. We model changes in individual attention to factors as being motivated by imitation; agents who are observed to be gaining greater rewards are imitated by those gaining fewer [30], leading to the classic replicator equation [36–38] describing the evolution of p_i, the proportion of agents attending to factor i:

$$\dot{\rho_i} = \rho_i \left(\mathbb{E}(R_i) - \sum_{j=1}^{n} \rho_j \mathbb{E}(R_j) \right), \quad [6]$$

where $\sum_{i=1}^{n} \rho_i = 1$ by definition. The expected reward ($\mathbb{E}(R_i) = 1$) is the mean reward an agent attending factor i will receive, averaging over all possible values of both x_i and the other factors x_j. It is thus determined by both the proportion of times that the agent will vote correctly (when $x_i = Y$) and the magnitude of the reward received on those occasions (determined by the reward system). To calculate this expectation we either exhaustively enumerate all possibilities (for $n < 10$) or numerically evaluate an approximation considering the normally distributed limiting behaviour (see below). When solving these n equations (one for each factor) numerically, we normalise the rewards given to all agents such that $\sum_{i=1}^{n} \rho_i \mathbb{E}(R_i) = 1$. This is equivalent to adaptive variation of the time step and does not change the relative rewards between options, nor the final steady state, but ensures smoother convergence to that state. This also mimics a real constraint on any practical reward system where the total reward available may be fixed. In our model we assume that agents reliably receive the expected reward for the factor that they attend to. Similar models with stochastic rewards (e.g. [13]) may show slower convergence to equilibrium. In our simulation of the collective dynamics of the system we used the Runge-Kutta order 2(3) algorithm, as implemented in R by Soetaert *et al.* [39].

The three reward schemes. We present three possible systems for rewarding agents for making accurate predictions. Each reward scheme corresponds to a choice of reward function, $f(z)$, which determines the magnitude of the reward when an agent makes an accurate prediction, as a function of the proportion, z, of other agents that also do so. These are:

1. Binary rewards: $f(z) = 1$
2. Market rewards: $f(z) = 1/z$
3. Minority rewards: $f(z) = 1 - H(z - 1/2)$, where H is the Heavyside step-function.

The expected reward an agent receives for attending to factor i is therefore the expected value of $f(z_i)$, conditional on their vote being accurate:

$$\mathbb{E}(R_i) = \int_{\epsilon}^{1} f(z_i) P(Y = x_i \mid z_i) p(z_i) dz. \quad [7]$$

where z_i is the proportion of agents voting identically to those attending to factor i: $z_i = \sum_{j=1}^{n} \rho_j \delta_{x_i, x_j}$, where δ is the Kronecker delta. The lower limit of the integral above is $\epsilon > 0$ to account for the limiting case of a single individual attending to the factor. As the population size N tends to infinity, ϵ tends to zero. For our implementation we take $\epsilon = 10^{-6}$.

Normal approximation for expected rewards. For $n \geq 10$ an exhaustive search over all 2^n combinations of $x_1, \ldots x_n$ is computationally infeasible. Instead we use the Central Limit Theorem to approximate the expected reward received for attending to any given factor. Focusing on a single individual who attends to factor i, we can calculate the expected reward received by the individual as follows. Firstly, we assume without loss of generality by symmetry that the focal individual observes $x_i = 1$. The expected reward, $\mathbb{E}(R_i)$ is then:

$$\mathbb{E}(R_i) = \int_{\epsilon}^{1} f(z_i) P(\psi > 0 \mid x_i = 1, z_i) p(z_i) dz \quad [8]$$

Given the independence of the individual values of x_i, the mean and variance of ψ can be determined by the linearity of expectations and by the sum rule for variances of independent variables:

$$\mathbb{E}(\psi \mid x_i = 1) = \beta_i \sum_{j \neq i}^{n} \beta_j \mathbb{E}(x_j) = \beta_i$$
$$\text{VAR}(\psi \mid x_i = 1) = \sum_{j \neq i}^{n} \beta_j^2 \mathbb{E}(x_j^2) = \sum_{j \neg i}^{n} \beta_j^2 \quad [9]$$
$$\Rightarrow p(\psi \mid x_i = 1) \simeq \mathcal{N}\left(\psi; \beta_i, \sum_{j \neq i} \beta_j^2\right)$$

In the case of binary rewards, where $f(z) = 1$, the value of z_i does not impact on the reward for attending to any factor. In this case the expected reward is calculated directly from the distribution of ψ:

$$\mathbb{E}_{\text{binary}}(R_i) = P(\psi > 0 \mid x_i = 1) = \Phi\left(\frac{\beta_i}{\sum_{j \neq i} \beta_j^2}\right) \quad [10]$$

For other reward schemes where the value of z_i affects the reward, we also require an approximation for $p(z_i)$. Again we calculate the mean and variance of z_i:

$$\mathbb{E}(z_i \mid x_i = 1) = \rho_i + \sum_{j \neq i} \rho_j (\mathbb{E}(x_j) + 1)/2 = (1 + \rho_i)/2$$

$$\begin{aligned} \text{VAR}(z_i \mid x_i = 1) &= \sum_{j \neq i}^{n} (\rho_j/2)^2 \mathbb{E}(x_j^2) = \frac{1}{4} \sum_{j \neq i} \rho_j^2 \\ \Rightarrow p(z_i \mid x_i = 1) &\simeq \mathcal{N}\left(z; \frac{1+\rho_i}{2}, \frac{1}{4} \sum_{j \neq i} \rho_j^2\right) \end{aligned} \quad [11]$$

The convergence of z_i in distribution to a normal distribution depends on the values of $\{\rho\}$ meeting the Lindeberg condition [35]. In practice this means that all elements of $\{\rho\}$ should tend to zero as the number of dimensions, n tends to infinity, i.e. the distribution should not be dominated by a small subset of elements. As illustrated in Figure 1, when the system is initialised in a state conforming to these requirements it will remain so for market and minority reward systems, but not for the binary reward system. Since the binary reward system does not depend on the value of z_i the failure of this approximation in this case does not have any repercussions for our results.

ψ and z_i are correlated due to the shared dependence on the values of $x_1, \ldots, x_n$, with a covariance of:

$$\text{COV}(z_i, \psi \mid x_i = 1) = \frac{1}{2} \sum_{j \neq i} \sum_{k \neq i} \beta_j \rho_k \mathbb{E}(x_j x_k) = \frac{1}{2} \sum_{j \neq i} \beta_j \rho_j \quad [12]$$

In the normal distribution limit,the joint distribution may be approximated as

$$p(\psi, z_i \mid x_i = 1) = \mathcal{N}\left(\begin{bmatrix}\psi \\ z_i\end{bmatrix}; \begin{bmatrix}\mu_\psi \\ \mu_z\end{bmatrix}, \begin{bmatrix}K_{\psi,\psi} & K_{\psi,z} \\ K_{\psi,z} & K_{z,z}\end{bmatrix}\right) \quad [13]$$

with,

$$\begin{aligned} \mu_\psi &= \mathbb{E}(\psi \mid x_i = 1) \\ \mu_z &= \mathbb{E}(z_i \mid x_i = 1) \\ K_{\psi,\psi} &= \text{VAR}(\psi \mid x_i = 1) \\ K_{z,z} &= \text{VAR}(z_i \mid x_i = 1) \\ K\psi, z &= \text{COV}(\psi, z_i \mid x_i = 1) \end{aligned}$$

Using standard relations for conditional normal distributions we therefore have:

$$\begin{aligned} p(\psi \mid x_i = 1, z_i) &= \mathcal{N}\left(\psi; \mu_\psi + (z_i - \mu_z)\frac{K_{\psi,x}}{K_{z,z}}, K_{\psi,\psi} - \frac{K_{\psi,x}^2}{K_{z,z}}\right) \\ \Rightarrow P(\psi > 0 \mid x_i = 1, z) &= \Phi\left(\frac{\mu_\psi + (z_i - \mu_z)\frac{K_{\psi,x}}{K_{z,z}}}{K_{\psi,\psi} - \frac{K_{\psi,x}^2}{K_{z,z}}}\right) \end{aligned} \quad [14]$$

Combining the above expressions gives the complete equation for the expected reward of attending to factor i, conditioned on the values of β, the current distribution of attention, ρ, and the reward function $f(z)$

$$\mathbb{E}(R_i) = \int_\epsilon^1 f(z_i) \mathcal{N}(z_i; \mu_z, K_{z,z}) \, \Phi\left(\frac{\mu_\psi + (z_i - \mu_z)\frac{K_{\psi,x}}{K_{z,z}}}{K_{\psi,\psi} - \frac{K_{\psi,x}^2}{K_{z,z}}}\right) dz_i \quad [15]$$

This integral may be evaluated numerically to give the expected reward for any general reward modulation function $f(z)$.

Calculating collective accuracy. The collective accuracy, C, is the probability that the collective vote will correctly predict the ground truth, conditioned on the current distribution of attention to different factors. For small numbers of factors (we use $n < 10$) this can be determined exactly by exhaustive search over all 2^n possible combinations of the values of $x_1, \ldots x_n$. For larger values of n we use the following normal approximation (similarly defined as above) for the joint distribution of the latent ground truth function ψ and the collective vote V.

$$p(\psi, V) \simeq \mathcal{N}\left(\begin{bmatrix}\psi \\ V\end{bmatrix}; \begin{bmatrix}0 \\ 0\end{bmatrix}, \begin{bmatrix}S_{\psi,\psi} & S_{\psi,V} \\ S_{\psi,V} & S_{V,V}\end{bmatrix}\right) \quad [16]$$

where

$$S_{\psi,\psi} = \sum_{i=1}^{n} \beta_i^2, \; S_{V,V} = \frac{1}{4} \sum_{i=1}^{n} \rho_i^2, \; S_{\psi,V} = \frac{1}{2} \sum_{i=1}^{n} \beta_i \rho_i, \quad [17]$$

implying the following conditional probability distribution for V given ψ:

$$p(V \mid \psi) \simeq \mathcal{N}\left(V; \psi \frac{S_{\psi,V}}{S_{\psi,\psi}}, S_{V,V} - \frac{S_{\psi,V}^2}{S_{\psi,\psi}}\right). \quad [18]$$

Considering without loss of generality the case where $Y = 1$,

$$\begin{aligned} C &= P(\hat{Y} = 1 \mid Y = 1) \\ &= P(V > 0 \mid \psi > 0) \\ &= 2 \int_0^\infty \int_0^\infty \mathcal{N}\left(V; \psi \frac{S_{\psi,V}}{S_{\psi,\psi}}, S_{V,V} - \frac{S_{\psi,V}^2}{S_{\psi,\psi}}\right) dV \mathcal{N}\left(\psi; 0, S_{\psi,\psi}\right) d\psi \\ &= 2 \int_0^\infty \Phi\left(\frac{\psi \frac{S_{\psi,V}}{S_{\psi,\psi}}}{S_{V,V} - \frac{S_{\psi,V}^2}{S_{\psi,\psi}}}\right) dV \mathcal{N}\left(\psi; 0, S_{\psi,\psi}\right) d\psi, \end{aligned} \quad [19]$$

which can be evaluated numerically. The normal approximation limit becomes invalid when the distribution of $\{\rho\}$ is concentrated on very few elements; in these cases (which we identify as 99% of the distribution mass being concentrated on fewer than 10 elements) we use exhaustive search over the values of $\{x\}$ corresponding to the remaining factors with non-negligible values of ρ.

ACKNOWLEDGMENTS. Cédric Beaume, Viktoria Spaiser and Jochen Voss provided valuable feedback on the manuscript. This work was supported by ERC Advanced Investigator Grant 'Momentum', reference number 324247.

1. Galton F (1907) Vox populi (the wisdom of crowds). *Nature* 75(7):450–451.
2. Mackay C (2012) *Extraordinary popular delusions and the madness of crowds.* (Start Publishing LLC).
3. Hertwig R (2012) Tapping into the wisdom of the crowd – with confidence. *Science* 336(6079):303–304.
4. Shefrin H (2009) How psychological pitfalls generated the global financial crisis. *Insights into the Global Financial Crisis* p. 224.
5. Santos FC, Pinheiro FL, Lenaerts T, Pacheco JM (2012) The role of diversity in the evolution of cooperation. *Journal of theoretical biology* 299:88–96.
6. Zafeiris A, Vicsek T (2013) Group performance is maximized by hierarchical competence distribution. *Nature Communications* 4.
7. Aplin LM, Farine DR, Mann RP, Sheldon BC (2014) Individual-level personality influences social foraging and collective behaviour in wild birds. *Proc. Roy. Soc. B* 281(1789):20141016.
8. Surowiecki J (2005) *The Wisdom of Crowds.* (Random House LLC).
9. Page SE (2008) *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies.* (Princeton University Press).
10. Page SE (2014) Where diversity comes from and why it matters? *European Journal of Social Psychology* 44(4):267–279.
11. Woolley AW, Chabris CF, Pentland A, Hashmi N, Malone TW (2010) Evidence for a collective intelligence factor in the performance of human groups. *Science* 330(6004):686–688.
12. Pickard G et al. (2011) Time-critical social mobilization. *Science* 334(6055):509–512.
13. Hong L, Page SE, Riolo M (2012) Incentives, information, and emergent collective accuracy. *Managerial and Decision Economics* 33(5-6):323–334.
14. Couzin ID (2009) Collective cognition in animal groups. *Trends in Cognitive Sciences* 13(1):36–43.
15. Purves D et al. (2013) Ecosystems: time to model all life on earth. *Nature* 493(7432):295–297.
16. Helbing D et al. (2015) Saving human lives: what complexity science and information systems can contribute. *Journal of Statistical Physics* 158(3):735–781.
17. Lämmer S, Helbing D (2008) Self-control of traffic lights and vehicle flows in urban road networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008(04):P04019.
18. Bell RM, Koren Y (2007) Lessons from the Netflix prize challenge. *ACM SIGKDD Explorations Newsletter* 9(2):75–79.
19. Wolfers J, Zitzewitz E (2006) Prediction markets in theory and practice, (National Bureau of Economic Research), Working Paper 12083.
20. Arrow KJ et al. (2008) The promise of prediction markets. *Science* 320(5878):877–878.

21. Oprea TI et al. (2009) A crowdsourcing evaluation of the NIH chemical probes. *Nature Chemical Biology* 5(7):441–447.
22. Morgan MG (2014) Use (and abuse) of expert elicitation in support of decision making for public policy. *Proceedings of the National Academy of Sciences* 111(20):7176–7184.
23. Herbert-Read JE, Romenskyy M, Sumpter DJ (2015) A turing test for collective motion. *Biology Letters* 11(12):20150674.
24. Lorenz J, Rauhut H, Schweitzer F, Helbing D (2011) How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences* 108(22):9020–9025.
25. Moussaïd M (2013) Opinion formation and the collective dynamics of risk perception. *PLoS ONE* 8(12):e84592.
26. Young NS, Ioannidis JP, Al-Ubaydli O (2008) Why current publication practices may distort science. *PLoS Med* 5(10):e201.
27. Stephan PE (2012) Research efficiency: Perverse incentives. *Nature* 484(7392):29–31.
28. Duarte JL et al. (2015) Political diversity will improve social psychological science. *Behavioral and Brain Sciences* 38:e130.
29. Prelec D (2004) A Bayesian truth serum for subjective data. *Science* 306(5695):462–466.
30. Helbing D (1996) A stochastic behavioral model and a 'microscopic' foundation of evolutionary game theory. *Theory and Decision* 40(2):149–179.
31. Dawes RM (1979) The robust beauty of improper linear models in decision making. *American Psychologist* 34(7):571.
32. Graefe A (2015) Improving forecasts using equally weighted predictors. *Journal of Business Research* 68(8):1792–1799.
33. Conte R et al. (2012) Manifesto of computational social science. *The European Physical Journal Special Topics* 214(1):325–346.
34. Helbing D (2016) Why we need democracy 2.0 and capitalism 2.0 to survive. *Jusletter IT.*
35. Feller W (1970) *An Introduction to Probability Theory and Its Applications: Volume 1.* (John Wiley & Sons) Vol. 1, 3 edition.
36. Schuster P, Sigmund K (1983) Replicator dynamics. *Journal of theoretical biology* 100(3):533–538.
37. Hofbauer J, Sigmund K (2003) Evolutionary game dynamics. *Bulletin of the American Mathematical Society* 40(4):479–519.
38. Nowak MA (2006) *Evolutionary dynamics.* (The Belknap Press of Harvard University Press, Cambridge, MA and London, UK).
39. Soetaert K, Petzoldt T, Setzer RW (2010) Solving differential equations in R: package deSolve. *Journal of Statistical Software* 33(9):1–25.

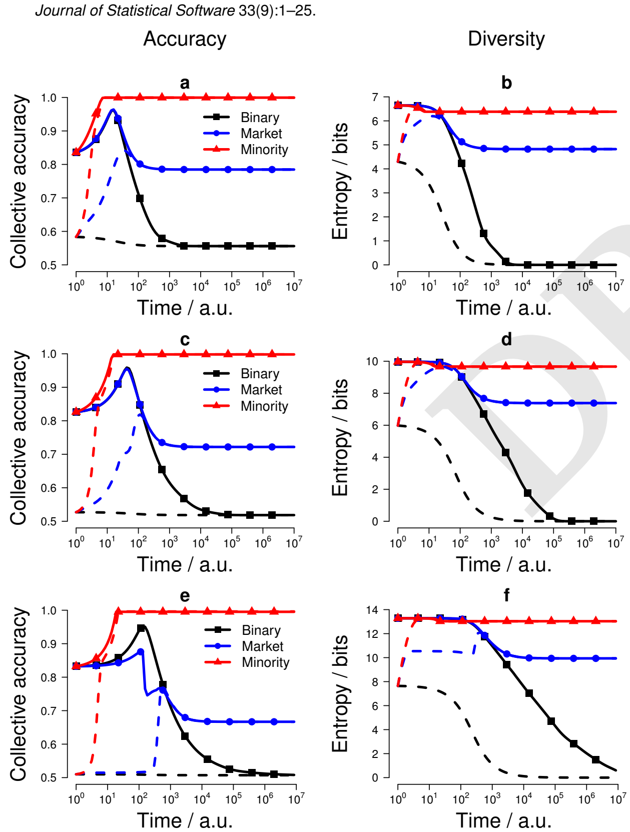

Fig. 1. Evolution of collective accuracy (left) and diversity (right) for binary rewards (black line), market rewards (blue line) and minority rewards (red line) in simulations with $n = 100$ (a,b), $n = 1000$ (c,d) and $n = 10,000$ (e,f) independent factors. Solid lines indicate results from a uniform initial allocation of agents over factors, while dashed lines indicate an initial allocation of 50% of agents to the single most important factor, with the remainder allocated uniformly over the remaining factors. Note that the number of time steps is plotted on a logarithmic scale.

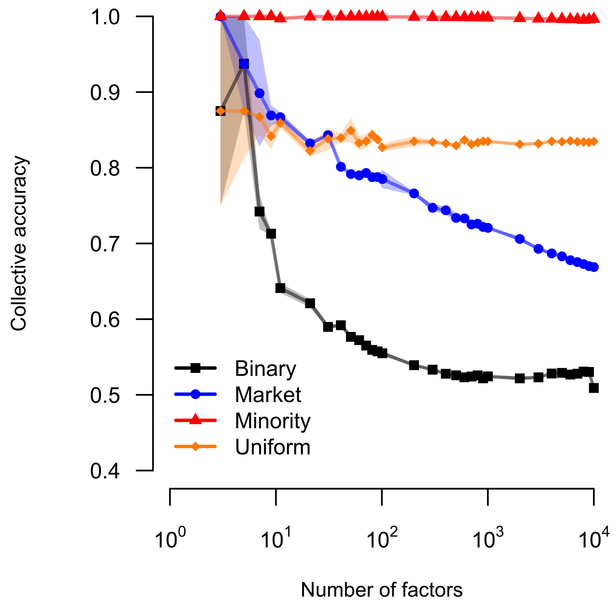

Fig. 2. Collective accuracy at equilibrium as a function of the number of independent factors across different reward systems. Solid lines and shaded regions show the mean and standard deviation of 10 independent simulations with different randomly generated values for the factor coefficients. Points on each curve show the precise values of n for which simulations were carried out, equally spaced within each multiple of 10.

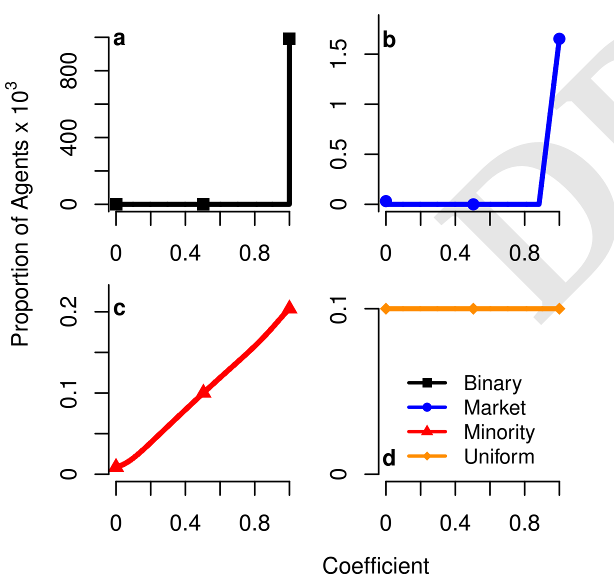

Fig. 3. Equilibrium proportions of agents paying attention to each factor, as a function of the coefficient associated with that factor. Results are shown for simulations with $n = 10000$ factors, and for the three reward systems of binary rewards (a), market rewards (b), and minority rewards (c), as well as the uniform allocation (d). Binary rewards drive almost all agents to the single most important factor (the greatest coefficient). Market rewards create a distribution proportional to coefficient size across the most important 10% of factors, while minority rewards distribute agents almost perfectly in proportion to the magnitude of the coefficient.