MANAGERIAL AND DECISION ECONOMICS
Manage. Decis. Econ. **33**: 323–334 (2012)
Published online 19 July 2012 in Wiley Online Library
(wileyonlinelibrary.com) DOI: 10.1002/mde.2560

Incentives, Information, and Emergent Collective Accuracy

Lu Hong[a], Scott E. Page[b,*] and Maria Riolo[c]

[a]*Department of Finance, Loyola University, Chicago, IL, USA*
[b]*Departments of Political Science and Economics, and Center for the Study of Complex Systems, University of Michigan, Ann Arbor, MI, USA*
[c]*Department of Applied Mathematics, University of Michigan, Ann Arbor, MI, USA*

In this paper, we construct a framework within which we explore how incentives and information structures influence the ability of a collection of individuals to make an accurate aggregate prediction. In our framework, individuals of bounded ability predict outcomes that depend on the values of a set of attributes. Individual construct models consider only a subset of those attributes, and those models depend on their incentives and their information environments. We consider two types of incentive structures: one in which individuals get paid on the basis of accuracy and one based on market like, for example, parimutuel payoffs. We also consider two information environments: one in which individuals learn in isolation and another in which they can copy more successful predictors. We find that market incentives and isolated learning environments produce the most accurate aggregate predictions but that these same incentives and information structures also produce the least accurate individuals. Thus, the incentives and informational structures that produce collective wisdom may hinge on their ability to produce and maintain diversity.

1. INTRODUCTION

The efficient functioning of organizations as well as markets and democracies requires that collections of individuals be capable of making accurate assessments of current and future events (Von Hayek, 1945; Ober, 2008). Often, these institutions demand that the crowd be more accurate than its constituent members (Woolley *et al*., 2010). Theoretical models demonstrate that this desired emergent collective accuracy requires a combination of individual expertise and collective diversity (Ladha, 1992; Armstrong, 2001; Page, 2007; Ray, 2006). These extant theoretical results provide sufficient conditions for collective accuracy but leave open the question of how those conditions emerge. That is the subject that we address here: how do wise crowds emerge? And, in particular, what roles do information structures and incentives play in producing them?

To answer these questions, we construct a framework within which we vary incentive and information structures. Incentives to make correct predictions should promote individual accuracy, yet if individuals are bounded or if the problem is computationally complex, individuals may not identify the best predictors (Aragones *et al*., 2005). Incentives to be correct need not, however, create diversity. Although some diversity should occur just by chance—the opinions and forecasts of individuals who possess only partial and diverse information and models of limited cognitive

*Correspondence to: Scott E. Page, Departments of Political Science and Economics, and Center for the Study of Complex Systems University of Michigan, Ann Arbor, MI, USA. Email: scottepage@gmail.com

depth need not be the same even under rational learning (Al-Najjar, 2009)—more is probably needed. To that end, we explore the ability of market like incentives, in which being correct when others are wrong yields large rewards, to produce and maintain diversity.

Our main findings are that market-based incentives improve collective accuracy and that social learning, in which individuals can share models, reduces it. Interestingly, we find that the incentive and information structures that lead to the most accurate collectives also produce the least accurate individuals. This result arises because collective accuracy requires diversity, and the mechanisms that produce diversity do so by reducing individual accuracy. Thus, organizations that estimate values, be they investment banks evaluating stocks, government agencies evaluating solvency, or firms evaluating market demands might do well to worry as much about incentivizing people to think differently as they do about improving individual accuracy.

In our framework, we assume an *outcome function* that maps a set of attribute values into either a good or bad outcome. These attributes take on binary values. Individuals choose to look at some subset of those attributes and construct *predictive models* that map the attribute values they consider into predicted outcomes. The accuracy of an interpreted signal depends on the number of attributes that the individual considers (the more the better) and the nature of the outcome function.

Our model builds from the interpretive signal framework of Hong and Page (Hong & Page, 2009), which enables us to capture diversity by the overlap in the attributes considered. If two individuals include non-overlapping subsets of attributes, their predictions will be negatively correlated. If the individuals consider identical attributes, their predictions will be perfectly correlated.

In what follows, we consider outcome functions that rely on a threshold. If a weighted average of the attribute values exceeds a threshold, the outcome is good. Otherwise, the outcome is bad. Individuals learn which attributes to include in their model as well as what weights to place on those attributes. Individuals also have a cost associated with adding more attributes. This cost prevents them from considering all of the attributes in their models. Over time, they learn to balance the informational benefits of adding more attributes against the cognitive or acquisitional costs of having a more sophisticated model.

We consider two information structures. In the first, we *isolate* the individuals. Each must construct a predictive model on its own. In the second, the individuals meet others and can copy their models. We call this *social learning*.[1] We also consider two types of incentives. Under *individual incentives*, an individual earns a payoff if it correctly predicts the outcome in that period, whereas under *market incentives*, an individual earns a payoff proportional to the inverse of the percentage of other individuals who predicted correctly. Thus, if twice as many individuals predict correctly, an individual only earns half as much.[2] In total then, we have four scenarios, and for each we measure (i) the accuracy of the individual predictive models that emerge; (ii) the diversity of those models; and (iii) the collective accuracy of all of the models.

Our emphasis on how incentives and informational structure influence the predictive models that individuals select differs from the bulk of the information aggregation literature in economics that focuses on the incentives for individuals to truthfully reveal and to gather information (Feddersen & Pesendorfer, 1997; Piketty, 1999; Hanson, 1999; Chen *et al.*, 2010). Here, we assume that individuals have incentives to reveal their predictions and to gather the relevant information to make those predictions. We do not deny the incentive issues exist with regards to revelation; we just choose to focus on an aspect of the problem that to date has been largely unexplored—the emergence of the accuracy and diversity of individual predictions based on information and incentives.

The remainder of this paper consists of three parts. In the first, we describe the basic framework and present some basic mathematical insights that follow from the model's construction. In the second part, we describe the results from a computational model for four classes of functions (Miller & Page, 2008). In the final part, we discuss the implications of our findings for organizations and societies that wish to produce emergent collective wisdom.

2. FRAMEWORK

In our framework, there exists a *binary outcome*, that is, either good or bad, that depends on the *state of the world*. The state of the world can be written as a vector of M binary attributes.

The **state of the world** $x = (x_1, x_2, \ldots x_M)$ *where* $x_i \in \{-1, 1\}$
The **outcome function** $F : (x_1, x_2, \ldots x_M) \rightarrow \{0, 1\}$.

DOI: 10.1002/mde

We consider the ability of a population of N individuals to predict that *outcome* using majority rule voting. Each of these individuals relies on a *predictive model* that considers a subset of the attributes. Let S_j denote the subset of the attributes $\{1, 2, \ldots M\}$ considered by individual j.

The **predictive model** *of individual j* $M_j\left(\{x_j\}_{j \in S_j}\right) \rightarrow \{0, 1\}$.

We say that the predictive model of an individual is *correct* if it agrees with the outcome function, and the *accuracy* of a predictive model equals the percentage of the time that it is correct. The *collective accuracy* of the population of N predictive models equals the percentage of the time that a majority of the individuals make the correct prediction.

We will say that a predictive model is *optimal* if it maximizes accuracy conditional on the attributes included in the model. As a general rule, if two individuals consider many of the same attributes, and arrive at optimal predictive models, then their predictions will be positively correlated. If they consider different attributes, their predictions will be negatively correlated. We can therefore use the overlap in attributes used in models as a proxy for the diversity of predictions from those models.

Incentives and Information. We initially endow the individuals with random predictive models in a way we make formal below. Over time they learn based on their available information and their incentives. We consider two different incentive structures. Our first incentive structure considers only an individual's accuracy. Individuals who predict correctly get paid. Those that predict incorrectly do not. We call these *individual incentives*. Our second incentive structure creates a pool of money that gets divided among those who pick correctly. Thus, we call these *market incentives* because they capture an essential feature of financial markets: predicting correctly when everyone else does as well yields a much smaller payoff than predicting correctly when everyone else is wrong.

individual incentives: *the payoff to an individual equals 1 if his or her model predicts correctly and 0 otherwise.*
market incentives: *the payoff to an individual who predicts correctly equals* N/N_c *where* N_C *equals the total number of individuals who predict correctly and 0 otherwise.*

We also consider two learning environments. In the first, no information is shared among the individuals. In the second, individuals can see the performance of other agents and copy their models if the other agents are earning higher payoffs.

isolated learning: *an individual compares a model based on a randomly chosen set of attributes against his or her current model. He or she switches only if the new model produces a higher payoff.*
social learning: *with probability* θ *an individual relies on isolated learning and with probability* $(1 - \theta)$ *he or she chooses a random individual (including possibly himself or herself) and compares that individual's model against his or her own. In each case, he or she switches only if the new model produces a higher payoff.*

Combining the incentive structures and learning environments produces four scenarios shown in Figure 1.

2.1. Learning Rules for Individuals

We now describe the specific learning rules that we encode in our computational implementations of these four scenarios. The computer programs consist of four parts: *initialization*, *model refinement*, *model switching*, and *evaluation*. The particulars of the various outcome functions that we use will be postponed until the results section.

Initialization. We construct N individual agents as follows. Each agent is assigned a set of K randomly chosen attributes from the M attributes that describe the state of the world. These K attributes combine to partition the 2^M possible states of the world into 2^K *categories*. For example, if $K = 1$, then the set of possible worlds is partitioned into one category in which the attribute has value 1 and one category in which it has value -1. We then construct a *look-up table* for each individual that keeps track of past states of the world as well as payoffs, by assigning each state of the world to its corresponding category and

	Individual Incentives	Market Incentives
Isolated Learning	I	II
Social Learning	III	IV

Figure 1. Our four scenarios.

DOI: 10.1002/mde

denoting the associated payoff for each prediction. This look-up table also contains the individual's prediction. As an individual has no prior knowledge about the relationship between the attributes considered in his or her model and the outcome, the initial predictions are random. Each agent also has a *cumulative payoff*. This equals the sum of all payoffs to date minus the costs of looking at attributes. We assume that each attribute considered has a cost c per period. This cost corresponds to the informational or attentional cost of paying attention to an attribute's value.

Model Refinement. We now describe how the agents update this look up table to refine their models. Under individual incentives, the payoff associated with a prediction will be 1 if the prediction is correct and 0 if the prediction is incorrect. Therefore, a lookup table for an individual who considers two attributes after 20 periods might look like Figure 2. Notice that this individual would predict good outcomes when his or her two attributes take opposite values and bad outcomes when they take the same value.

Under market incentives, the payoffs depend on the number of people who predict correctly, so the individual may prefer to predict the less likely outcome because it has a higher expected payoff. Under market incentives, rather than add one to the correct column after each period, we add an amount equal to N/N_c, where N_C is the number who predicted correctly. In either scenario, after each period, an individual records the reward he or she received and updates his or her look-up table by incrementing the incentives for making positive and negative predictions given the state of the world.

Model Switching. In all of our scenarios, agents can also switch their models. In the isolated learning environment, model switching involves adding, dropping, or switching an attribute. In the social learning environment, model switching can also include abandoning one's own model for that of someone who has earned a higher payoff. First, for isolated learning, we assume that with some fixed probability, p_d, an individual decides to assess the usefulness of each the attributes he or she considers. With probability$(1-p_d)$ he or she does nothing to his or her model.

When an individual evaluates his or her attributes, he or she does so by comparing its added value in accuracy to the costs of looking at that additional attribute. He or she does this by rerunning history for 300 periods and determining his or her payoff had he or she not included the attribute. If he or she does not have 300 periods of data in his or her memory, he or she uses what data he or she has available. Dropping an attribute will generally lower accuracy, but it could increase the payoff by reducing costs. Note also that in the market incentive scenarios, lower accuracy need not imply a lower payoff even without considering costs. In all cases, we assume that if dropping the attribute would have led to a higher payoff, the individual drops it. Note that if an individual drops an attribute, then rows of the look-up table can be combined because he or she now relies on fewer categories.

If, on the other hand, all of his or her attributes are worth keeping, then the agent chooses a random attribute to add. If an individual adds a new random attribute to his or her model, he or she must add columns to his or her look-up table. We assume therefore that he or she starts with an empty look-up table. This asymmetry could be avoided if all individuals kept track of all of the attributes of all states of the world, but this would be empirically implausible. Thus, even though this assumption introduces an asymmetry related to the adding and dropping of attributes, we believe it to be the most sensible choice.

Under social learning, individuals apply the procedure just described with probability θ, and with probability $1-\theta$, they randomly pick some other

Attributes	Good Outcomes	Bad Outcomes	Prediction
$x_i = -1, x_j = -1$	1	4	Bad
$x_i = -1, x_j = 1$	4	0	Good
$x_i = 1, x_j = -1$	4	1	Good
$x_i = 1, x_j = 1$	1	5	Bad

Figure 2. A look-up table after 20 periods given individual incentives.

DOI: 10.1002/mde

individual with whom to compare. If that other individual has a higher cumulative payoff, then the first individual copies the look-up table of the second. We assume that he or she also copies the values within those columns. In other words, under our social learning assumption, we assume that individuals share the details of their models with one another. With some small probability (we used 0.05), the first individual makes a mistake in copying, in which case he or she substitutes a random attribute for one of the attributes to be copied and starts with an empty look-up table. This captures an individual getting an idea for a new model upon hearing someone else's model.

Evaluation. To evaluate the individuals and the collective, we keep track of three time series of data. First, we keep track of the average *individual accuracy*: the average percentage of periods in which each individual predicts correctly. Note that this equals the average payoff in the individual incentives scenario. We also keep track of the *collective accuracy*, which we calculate to be the average accuracy of the majority decision. If a majority predicts that the outcome is good and it is good, then the collective is correct. Finally, as mentioned in our construction of the framework, diversity over the attributes considered will be a proxy for negatively correlated individual predictions and therefore a more accurate collective. Therefore, we also keep track of the diversity of the predictive models.

There are several ways that we could capture model diversity. One way would be to consider all possible models that an individual might use and then to compute a diversity measure over that distribution. This approach creates computational difficulties in that the set of possible models is so large. An alternative approach is to capture model diversity by considering the distribution over the attributes considered in the models. This second approach gives a cruder measure of the actual model diversity, but it is more easily calculated and as our findings reveal, proves sufficient to capture model diversity. Note also that when each individual considers only a single attribute, these two approaches will be the same.

To make this formal, we first compute R, the total number of attributes considered by all individuals. We then calculate the number of individuals that consider attribute i. Call this r_i. The probability vector $\left(\frac{r_1}{R}, \frac{r_2}{R}, \ldots, \frac{r_M}{R}\right)$ denotes the percentage of people who look at each attribute. To calculate the diversity of this vector, we use the *diversity index*.[3]

Let p_i be the fraction of individuals who consider attribute i. We define the **diversity index** *of the population as a function of the fraction of individuals considering each of the M attributes*, $\text{DIV}\ (p_1, p_2, \ldots, p_M) = \frac{1}{\sum_{i=1}^{M} p_i^2}$

The diversity index is constructed so that if each individual considers the same two attributes, the index will equal two ($\text{DIV}\left(\frac{1}{2}, \frac{1}{2}, 0, \ldots, 0\right) = 2$), and if the individuals spread attention evenly across all M attributes, the diversity will equal M.

2.2. Classes of Functions

We analyze this model for four classes of functions: *equally weighted linear*, *linear with a dominant attribute*, *general linear functions*, and *linear with nonlinear interactions*. By analyzing the first class, we can compare the abilities of the various incentive and informational structures to maintain diversity against drift. The second class enables us to see how the structures create a balance between individual accuracy (focusing on the dominant attribute) and collective diversity. The third class subsumes the first two classes and enables us to see if the results found for the more special cases hold generally. The final class provides a context to see some of the limits of collective accuracy (Page & Toole, 2011). Unless the individuals can absorb all the interaction terms within their models, they will not be able to make accurate collective predictions.

2.3. The Dynamics over Models

Prior to presenting our findings, we first present some mathematical formalism that provides a foundation for the results that follow. In the first case that we consider, we assume that the outcome function takes the following form:

$$F_\beta(x) = \delta\left[\sum_{i=1}^{M} \beta x_i\right],$$

where $\delta(y) = 1$ if $y \geq 0$ and 0 otherwise, and $\beta > 0$ We furthermore assume that M is odd so that all outcomes are either good or bad. Note that under individual learning and no model switching, in the long run an individual will construct an optimal predictive model based on the attributes he or she considers. This follows directly from the law of large numbers.

Also note that given the symmetry, any two predictive models that consider the same number of attributes will be equally accurate in the long run. Thus,

DOI: 10.1002/mde

our focus in this case will be on how the various incentive and information structures produce and maintain diversity. The production of diversity will not be a problem given our assumption that individuals choose models randomly. In fact, if the individuals could maintain their initial diversity over attributes and optimize those models, then they would always make accurate predictions for this class of functions. This result follows from the fact that the distribution over models will be a Bernoulli distribution. However, given that an individual can also change the attributes that they consider, both through model refinement and model switching, the distribution over models will drift over time. This drift away from an equal distribution will result in less accuracy, at least in this case.

To give some insight into the dynamics of model refinement and model switching, we first consider a restricted case in which each individual's model may only consider one attribute at any time. In this special case, the dynamics of the system become analytically tractable. The results we derive will allow us to form some hypotheses about the diversity that will emerge under each of the four learning and reward scenarios.

For simplicity, we assume that in each period, only one individual can switch his or her model. Let $\left(k_1^t, k_2^t, \ldots, k_M^t\right)$ denote the distribution of individuals across attributes at time t. In what follows, we assume that the number of states of the world for which the two models are compared during the learning phase, is sufficiently large that if one model has a higher expected payoff, then the individual will switch to that model. Without loss of generality, we assume that individual who learns in period t considers attribute 1. Let $J = \{j : k_j < k_1\}$ denote the attributes that are chosen with less frequency than 1 and $L = \{\ell : k_\ell = k_1\}$ denote the set of attributes chosen with identical frequency. Let $\delta_J(i) = 1$ if $i \in J$ and 0 otherwise. Define $\delta_L(i)$ similarly.

Given these assumptions, we can write the distribution over models in the next period for each of our four scenarios as follows:

$$\begin{aligned}
&\text{Isolated, Individual}: && \left(\frac{1}{2}+\frac{1}{2M}, \frac{1}{2M}, \frac{1}{2M}, \ldots, \frac{1}{2M}\right)\\
&\text{Social, Individual}: && \left(\frac{1}{2}+\frac{\theta+(1-\theta)k_1^t}{2M}, \frac{\theta+(1-\theta)k_2^t}{2M}, \ldots, \frac{\theta+(1-\theta)k_M^t}{2M}\right)\\
&\text{Isolated, Market}: && \left(\frac{M-\left|J+\dfrac{L}{2}\right|}{M}, \frac{2\delta_J(2)+\delta_L(2)}{2M}, \ldots, \frac{2\delta_J(M)+\delta_L(M)}{2M}\right)\\
&\text{Social, Market}: && \left(\frac{N-\left|J+\dfrac{L}{2}\right|}{M}, \frac{(\theta+(1-\theta)k_2)[2\delta_J(2)+\delta_L(2)]}{2M}, \ldots, \frac{(\theta+(1-\theta)k_M)[2\delta_J(M)+\delta_L(M)]}{2M}\right)
\end{aligned}.$$

Comparing the expressions shows that isolated learning with individual incentives will result in an unbiased random walk across the M models. Maximal collective accuracy does not require equal numbers of each model. For example, with five models, a distribution of (3, 5, 5, 6, 6) would be just as accurate as an equal distribution. Overall then, we should expect isolated learning with individual incentives to perform well but at times it may drift toward distributions in which a few models dominate and accuracy will fall.

Under social learning with individual incentives, the process will be more likely to drift toward models based on attributes that have greater representation in the population even though those models are not more accurate. Thus, we should expect social learning to be much less accurate than under isolated learning.

Under isolated learning with market incentives, those models that are less represented will receive higher payoffs. As a result, the drift will be toward those models, thereby creating a drift that is biased toward equal numbers of the models. We should expect greater collective accuracy because a more even distribution across models will be maintained. Finally, under social learning with market incentives, we see a bias toward models that exist in equal or lower proportion but that bias is attenuated by the fact that less represented attributes will be chosen less frequently during social learning. Comparing across these distributions, we expect that social learning should lead to less collective accuracy than individual learning, and market incentives should lead to higher collective accuracy.

So far, we have assumed that each model is equally accurate. If we suppose that some models are more accurate than others, then we may no longer want an equal distribution across the models. Instead, the more accurate models should obtain more weight. Under individual incentives, we should expect a drift toward those more accurate models leading to less diversity, although more accurate individual predictions. That drift will be even more pronounced under social learning. Under market based incentives, the larger payoff from being in the minority should attenuate that drift toward individual accuracy and maintain some diversity. However, the market incentives may not prove sufficient in all cases. Consider functions of the following form:

$$F_\beta(x) = \delta\left[\sum_{i=1}^{4} \beta_i x_i\right],$$

where $\beta_1 > \beta_2 > \beta_3 > \beta_4 > 0$ and $\beta_2 + \beta_3 + \beta_4 > \beta_1 > \beta_2 + \beta_3$. The outcome will be good if F has a positive value and bad otherwise. Given these assumptions, a model that considers only attribute one will be correct unless all of the other three attributes take the opposite sign. Thus, it will be correct with probability 7/8. The models that consider the other attributes will be correct if either attribute one takes the same sign, (probability 1/2) or if attribute one takes the opposite sign and the other two attributes take the same sign (probability 1/8). The optimal allocation of individuals across attributes would be to have fewer than half of the individuals consider attribute one and to also have fewer than half of the individuals consider any subset of the other attributes. It can be shown that any distribution that satisfies those two assumptions will always yield an accurate collective prediction. In other words, the optimal distribution across attributes will not be maximally diverse.

In the long run, individual incentives will lead to everyone choosing the same model, which will not be optimal. And, as we now show, market incentives will also fail to produce an optimal distribution in equilibrium. Let p_i denote the proportion of individuals that choose a model that considers attribute i. By symmetry, it follows that $p_2 = p_3 = p_4$. Let this proportion equal p so that $p_1 = 1 - 3p$. Given market incentives, the payoff to the model that considers attribute 1 as a function of p equals

$$\pi_1(p) = \frac{1}{8}\left(1 + \frac{3}{1-p} + \frac{3}{1-2p}\right).$$

The payoff to an individual whose model considers attribute $j \neq 1$ equals

$$\pi_j(p) = \frac{1}{8}\left(1 + \frac{2}{1-p} + \frac{1}{1-2p} + \frac{1}{3p}.\right)$$

Setting these two expressions equal gives $\frac{1}{1-p} + \frac{2}{1-2p} = \frac{1}{3p}$, and rearranging terms gives $3p - 6p^2 + 6p - 6p^2 = 1 - 3p + 2p^2$. This reduces to $14p^2 - 12p + 1 = 0$, which has a solution $p = 0.0935$. This formal argument shows that even though market incentives produce diversity in equilibrium, that diversity will be insufficient to produce collective accuracy. In this particular case, the diversity index in equilibrium will equal 1.836. One can show that in order for the collective to predict accurately, the diversity index must lie in the interval $[0, 2]$.[4]

3. COMPUTATIONAL RESULTS

We now present results from a computational model. We present results from experiments with 101 agents although the qualitative results hold for a wide range of parameters.[5] In what follows, we take average values over individual runs of 30,000 periods. If we treat each period as a data point, then all of the differences that we present are significant at the $p = 0.01$ level.

Linear Functions with Five Equal Attributes. We first consider linear functions with five attributes. In this first set of results, we let costs per attribute equal 0.15. These costs are set so that an optimal model considers only a single attribute. As will be clear from the results, in any given time period, some percentage of the individuals will be using a two-attribute model. This occurs because the individuals continually experiment with new models. Given our assumptions, the expected accuracy of each individual equals 68.75%.[6]

Figure 3 shows the average accuracy of the group (the top line) as well as the moving average accuracy for each of two individuals under social learning with individual incentives. The up and down movements in this average are consistent with the random walk posited in the previous section. The theoretical results implied that market incentives should produce higher group accuracy than individual incentives and that isolated learning should produce higher group accuracy than social learning. We find both to be true as shown in Table 1.

DOI: 10.1002/mde

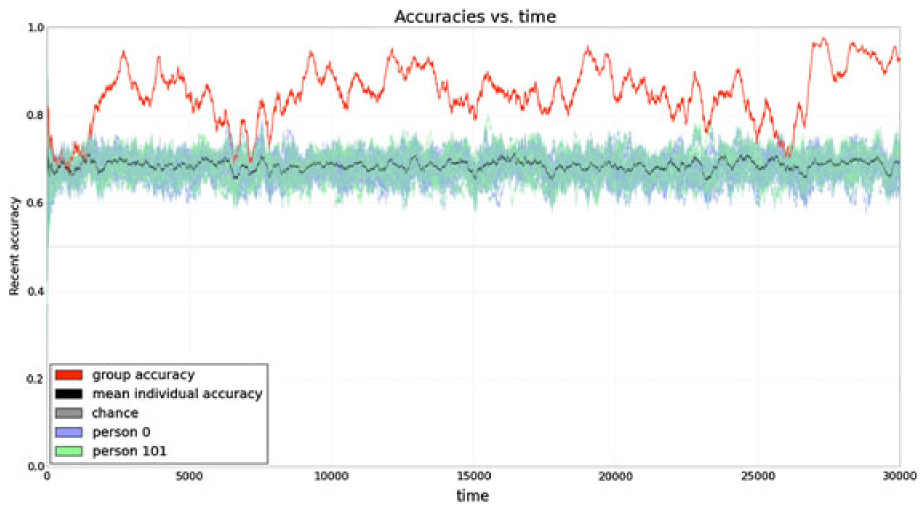

Figure 3. Social learning, individual incentives with five equal attributes.

Table 1. Group Accuracy for Five Equally Weighted Attributes 101 Individuals, 30,000 Periods

Isolated, individual	Social, individual	Isolated market	Social market
87.39	84.83	96.35	95.44

We also find that market incentives result in more diversity than do individual incentives. The advantage of market incentives stems from their ability to maintain diversity in the population. Recall that in this example, maximal diversity equals five. Figures 4 and 5 show the time series for the diversity index for isolated learning and market incentives (the most diverse population) and social learning and individual incentives (the least diverse). Isolated learning and market incentives produce almost maximal diversity (five). Not only does social learning produce less diversity, it is also more volatile. Notice too that the periods of low accuracy seen in Figure 3 corresponds to periods of low diversity in Figure 3.

Linear Functions with a Dominant Attribute. For our second class of functions, we consider a function defined over five attributes in which one of the attributes is dominant. We consider the specific class of functions described previously where

$$F_\beta(x) = \delta\left[\sum_{i=1}^{4} \beta_i x_i\right],$$

and $\beta_1 > \beta_2 > \beta_3 > \beta_4 > 0$ and $\beta_2 + \beta_3 + \beta_4 > \beta_1 > \beta_2 + \beta_3$. Recall that the most accurate model considers only attribute one and is accurate with probability 87.5%. Here, we find that for all of the cases except for market

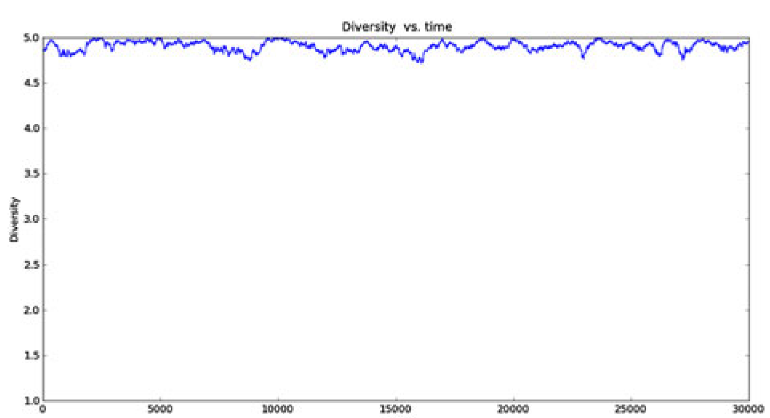

Figure 4. Diversity: isolated learning, market incentives with five equal attributes.

DOI: 10.1002/mde

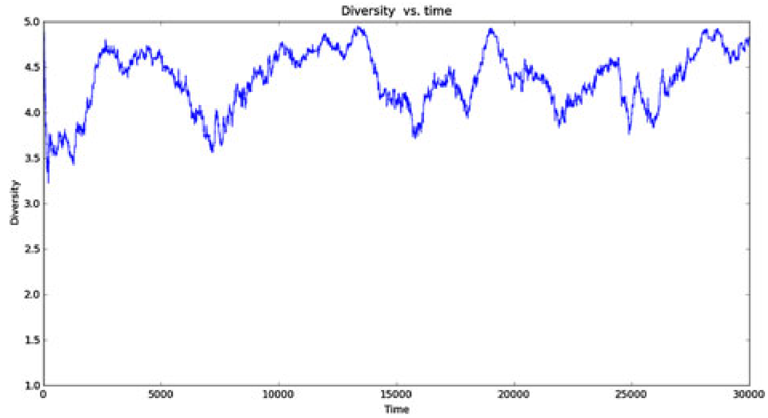

Figure 5. Diversity: social learning, individual incentives with five equal attributes.

incentives and isolated learning, the group's accuracy is not significantly different from this percentage. The data from the runs show that a strict majority of the individuals consider only attribute one. However, under market incentives, with isolated learning, the average individual accuracy equals only 78.94%. Yet, the group's accuracy equals 88.16%, only a slight improvement over the other scenarios. The market incentives maintain some diversity as shown in Figure 6. The amount of diversity exceeds that found in the theoretical model because some portion of the agents consider more than one attribute. In fact, the diversity levels lie in the range that would make it possible for the collective to be 100% accurate. Yet, they are not. Thus, it seems that market incentives may produce sufficient diversity, but they do not always produce an accurate collective prediction.

General Linear Functions. We next consider a more general class of linear functions. Here, we allow the coefficients to be drawn uniformly at random from $[-1, 1]$, and we increase the number of attributes to nine. Formally, we let

$$F_\beta(x) = \delta\left[\sum_{i=1}^{9} \beta_i x_i\right],$$

where x_i is, with equal probability, -1 or 1. In each run of the model, we drew new coefficients. Thus, some of the model runs may be like our equal attributes model, whereas others may have one or two dominant attributes. Each run of the model consisted of 100,000 periods. We present average data over 200 runs.

Figure 7 shows results on average individual and group level accuracy as well as diversity. The dark bars show average individual accuracy, and the light bars show collective accuracy across the 200 runs. The graph also includes error bars that mark out two standard deviations in each direction. The data show that markets create less accurate individuals, more diversity, and more accurate groups. Isolated learning

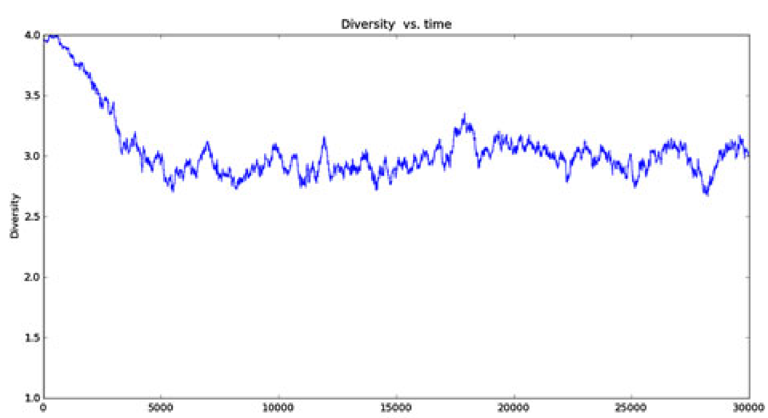

Figure 6. Diversity: isolated learning, market incentives, with one dominant attribute.

DOI: 10.1002/mde

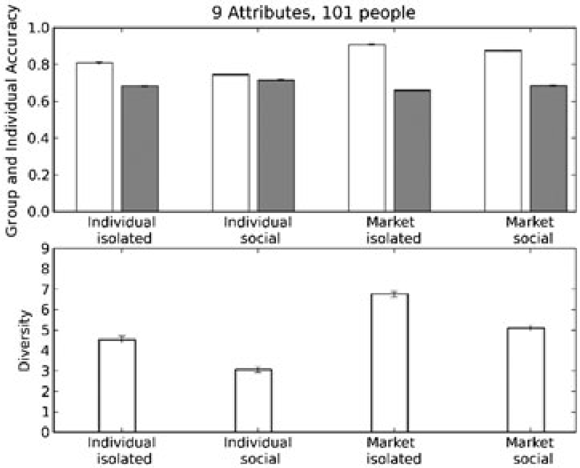

Figure 7. Individual and collective accuracy, diversity (220 runs).

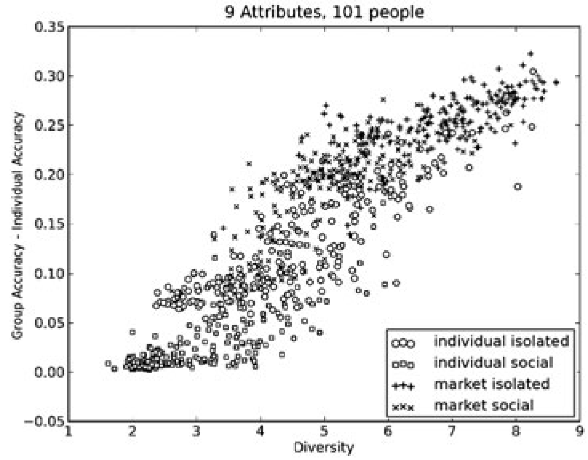

Figure 8. Correlation between diversity bonus and diversity index (220 runs).

has the same effects, although of less magnitude. Overall, the results from the computational model paint a consistent picture. The ordering of the four scenarios by collective accuracy is identical to their ordering by the amount of diversity that they maintain.

Market, **Isolated** $\succ$ **Market**, **Social**
$\succ$ **Individual**, **Isolated**
$\succ$ **Individual Social**

Somewhat surprisingly, the ordering by individual accuracy turns out to be the exact reverse. At least in our model, incentives that produce accurate individuals prove less capable of producing accurate collectives. Overall, although market incentives and isolated learning both produce more accuracy markets have a larger effect because they produce far more diversity (as measured by the diversity index of the attributes considered). In fact, market incentives with isolated learning produce nearly twice the diversity as do individual incentives and social learning.

This relationship between diversity and group accuracy can also be seen by plotting the difference between group accuracy and average individual accuracy against the diversity index. Figure 8 shows data from 200 runs for each scenario. This plot shows that model diversity, as measured by the diversity index of the distribution of attributes considered by the individuals, correlates strongly with what we call the *diversity bonus*, the difference between group accuracy and individual accuracy. This correlation holds both within each scenario and across the four scenarios. This plot shows how markets create much more model diversity and a corresponding larger diversity bonus.

4.1. Linear with Nonlinear Interaction Terms

For our final set of experiments, we consider linear functions with pairwise nonlinear terms. Our functions take the following form:

$$F(x) = \sum_{i=1}^{5} \beta_i x_i + \gamma_{12} x_1 x_2 + \gamma_{34} x_3 x_4,$$

where the β_i are uniformly drawn from $[-1, 1]$ and the γ_{ij} are uniformly drawn from $[-4, 4]$. Here again, we find that isolated information environments maintain more diversity and produce greater collective accuracy. However, the effects of market incentives do not have as large of an effect on collective accuracy. In cases with nonlinear terms, collective accuracy requires that some individual captures each nonlinear interaction. Markets create incentives to locate those nonlinear interactions and include those attributes. However, for an individual learner in isolation, once a nonlinear effect is located, it will not be abandoned for the simple reason that dropping an attribute drops both the linear term and the interaction term. Therefore, the fact that the market does not improve on individual learning speaks less to the inability of markets than to the power of individual, isolated learning in the presence of nonlinear effects.

4. DISCUSSION

In this paper, we have constructed a model to explore how collective accuracy emerges in a model in which individuals choose models in various informational

DOI: 10.1002/mde

environments under different incentives. We present two main findings: First, we find that isolated learning produces greater model diversity and therefore more accurate collectives. Second, we find that market incentives, in which individuals payoffs are inversely proportional to the percentage of correct predictors produce more diversity but less individual accuracy. This increase in diversity more than offsets the loss in individual accuracy implying that market incentives with isolated learning provide the best environment for collective accuracy to emerge. These results broadly agree with those derived by Golub and Jackson (Golub & Jackson, 2010), who study the wisdom of crowds in social networks. They take the distribution of initial signals as given and examine how much weight people should place on the opinions of others. They find that the wisdom of crowds requires vanishing weight on the most connected people, in other words, more equal weighting of signals.

We also find that although markets incentives create greater diversity, they do not produce sufficient diversity for the crowd to be 100% accurate, which theoretically, they should be able to achieve. In the case of five equally weighted attributes, the aggregate comes close to perfect accuracy, but when we increase the number of attributes to nine and average across a variety of functional forms for the outcome function, accuracy falls to around 90%. Individual accuracy lies far below that level, implying that collective accuracy does in fact emerge, but interestingly, it emerges partly through a lack of interaction—the isolated learning—and partly through linked payoffs—the market incentives.

To see if stronger incentives for diversity would produce more accurate collectives, we also experimented with what we call *hyper market incentives*, in which the payoff to an individual was proportional to one over the square of the number of individuals who predict correctly. These incentives might not be feasible in practice as payouts could be large if only a single individual predicted correctly. In such cases, these incentives will fail to satisfy the balance criterion from mechanism design (Myerson, 2008). Violations of the balance criterion are most relevance in market settings in which resources paid out must equal resources received. In the context of an organization, a person may be concerned with reputation enhancement. In such cases, no balance criterion need exist. And, a lone person who foresees a crash might indeed earn enormous benefits. Our analysis of hyper market incentives reveals that they perform only a little better than our market incentives.

We see our findings as suggestive of several directions for future inquiry, particularly in expanding the set of possible informational structures and incentives systems and introducing individual level cost heterogeneity. Our results suggest that there probably will not be a single incentive system that works for all outcome functions. Thus, it may be optimal for the incentive system to itself adapt. In addition, although diversity produces better outcomes, more diversity need not always be better. The scenarios that we considered produced insufficient diversity but that does not mean that all incentive systems and informational structures would do so. It would be possible as well, to produce too much diversity, at the expense of a lack of individual accuracy. An ideal mechanism would balance the incentives for individual accuracy and for collective diversity so that wise crowds can emerge.

Acknowledgements

This material is based upon work supported by, or in part by, the US Army Research Laboratory and the US Army Research Office under contract/grant number W911NF1010379. This research is also supported by the National Science Foundation's IGERT grant, an NSF research grant, and an AOR research grant.

NOTES

1. Henrich (Henrich, 2004) calls this cultural learning.
2. These are equivalent to parimutuel payoffs with the house taking no payoff. The market incentives scenarios share several features with the El Farol Problem (Arthur, 1994) and the Minority Game (Challet & Zhang, 1998). In each of these models, individuals have incentives to construct diverse models and can earn high payoffs by doing so. In those models, the outcome being predicted and the number of people at El Farol depend on the actions of the individuals. Here, the outcome occurs independently of the predictions of the individuals.
3. In Political Science, this measure also captures the effective number of parties, and in biology, the identical measure called Simpson's Index is used to capture species diversity. See Page (Page, 2010) for a survey of diversity measures.
4. The upper bound can be attained by setting $p_1 = \frac{1}{4} + 3\varepsilon$ and $p_j = \frac{1}{4} - \varepsilon$ for $j = 2, 3, 4$. The lower bound can be attained by setting $p_1 = \frac{1}{2} - \varepsilon_1$ and $p_2 = \frac{1}{2} + \varepsilon_1 - \varepsilon_2 - \varepsilon_3$, $p_3 = \varepsilon_3$ and $p_4 = \varepsilon_4$.
5. These were run in both Python and Netlogo. Code is available from the authors.
6. For each of the four scenarios, the average accuracy lied within 0.2% of that value.

REFERENCES

Al-Najjar N. 2009. Decision makers as statisticians: diversity, ambiguity, and learning. *Econometrica* **77**(5): 1339–1369.

DOI: 10.1002/mde

Aragones E, Gilboa I, Postlewaite A, Schmeidler D. 2005. Fact-free learning. *Am. Econ. Rev.* **95** (5): 1355–1368.

Armstrong JS. 2001. Combining forecasts. In *Principles of Forecasting: A Handbook for Researchers and Practitioners* Kluwer Academic Publishers: Norwell, MA.

Arthur WB. 1994. Inductive reasoning and bounded rationality. *American Economic Review Papers and Proceedings* **84**: 406–411.

Challet D, Zhang YC. 1998. On the minority game: analytical and numerical studies. *Physica A: Statistical and Theoretical Physics* **256**(3): 514–532.

Chen Y, Dimitrov S, Sami R, Reeves D, Pennock DM, Hanson RD, Fortnow L, Gonen R. 2010. Gaming prediction markets: equilibrium strategies with a market maker. *Algorithmica* **58**(4): 930–969.

Feddersen T, Pesendorfer W. 1997. Voting behavior and information aggregation in elections with private information. *Econometrica* **65** (5): 1029–1058.

Golub B, Jackson M. 2010. Naive learning in social networks: convergence, influence and the wisdom of crowds. *American Economic Journal: Microeconomics* **2**(1): 112–49.

Hanson R. 1999. Decision Markets. *IEEE Intelligent Systems* (3): 16–19.

Henrich J. 2004. Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior & Organization* **53**(1): 3–35.

Hong L, Page SE. 2009. Generated and interpreted signal. *Journal of Economic Theory* **144**(5): 2174–2196.

Ladha K. 1992. The Condorcet jury theorem, free speech, and correlated votes. *American Journal of Political Science* **36**(3): 617–634.

Miller J, Page SE. 2008. *Complex Adaptive Systems*. Princeton University Press: Princeton.

Myerson R. 2008. Perspectives on mechanism design in economic theory. *Am. Econ. Rev.* **98**(3): 586–603.

Ober J. 2008. *Democracy and Knowledge: Innovation and Learning in Classical Athens*. Princeton University Press: Princeton.

Page SE. 2010. *Diversity and Complexity*. Princeton University Press: Princeton.

Page SE. 2007 *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press: Princeton.

Page SE, Toole J. 2011. Predicting cellular automata. *Complex Systems* **19**(4): 343–362.

Piketty T. 1999. The information-aggregation approach to political institutions. *European Economic Review* **43**: 791–800.

Ray R. 2006. Prediction markets and the financial wisdom of crowds. *Journal of Behavioral Finance* nd; **7**(1): 2–4.

Von Hayek F. 1945. The use of knowledge in society. *Am. Econ. Rev.* **4**: 519–530.

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.