Avoiding Moral Hazards in Organizational Forecasting

Tad Hogg and Bernardo A. Huberman
HP Labs
Palo Alto, CA 94304

June 13, 2002

Abstract

We describe a new mechanism that induces accurate forecasts within an organization while reducing moral hazards and the stigma associated with negative opinions. It is based on the notion of identity escrow, whereby the identity of a forecaster is kept anonymous and only revealed when a number of his subordinates detect an attitude that is contrary to the interests of the organization. An analysis of the relative payoffs between forecasting and production shows that through the use of identity escrows one can adjust the size of the prediction group so as to ensure both production and accurate forecasts.

1 Introduction

Predicting the future of uncertain situations is an extremely important problem for most organizations. From problems of planning, to evaluating technology scenarios and assessing the state of the market for given products, forecasting these uncertain events is a prerequisite to making any strategic decision. As a result, large resources are spent trying to forecast sales and revenues, the success or failure of given technologies and the fate of certain trends determined by competitors and outsiders. And while there is a plethora of methods to choose from, ranging from committees of experts to consultants and the use of statistical inference techniques, the results tend to suffer in terms of accuracy and ease of implementation.

A possible alternative for the prediction of future outcomes is the construction of markets within organizations, where the asset is information rather than a physical good. Such markets involve the trading of state-contingent securities. If these markets are large enough and properly designed, they can be more accurate than other techniques for extracting diffuse information, such as surveys and opinions polls. There are problems, however, with information markets as they tend to suffer from information traps [3], illiquidity [11], manipulation [8], and lack of equilibrium [2]. These problems are exacerbated when the groups involved are small.

Recently, a market based method for forecasting the future by small groups was developed and tested in the laboratory [4]. The method works by treating individual ideas about possible events as assets, and incorporating the risk attitudes of the participants into the overall forecast much in the spirit of portfolio theory. The mechanism successfully aggregates scattered information and produces reliable forecasts of uncertain events. Equally important it is easy to implement and to deploy inside an organization.

In spite of the advantages this method offers, it still does not solve a pervasive problem associated with any forecasting methodology, posed by the fact that those directly involved in producing forecasts often have conflicting interests

at stake. While their accuracy in predicting future sales or production might be highly rewarded and even compensated accordingly, managers participating in these prediction processes also have control over production resources and human capital, and can therefore adjust the output of the units under their control to fit the forecast. If high output forecasts result in increased pressure to produce, the organization benefits as a whole. But there are also situations when knowledge of a pessimistic forecast will make managers slack in their supervisory roles, since compensation which takes into account the accuracy of a forecast also gives an incentive to reduce production.

Thus the desire to obtain accurate forecasts can create a moral hazard problem for the organization. This hazard occurs when individuals expecting lowered production are then tempted to work against the interests of their group to ensure a high payoff from their forecasts. Even if group members don't actually follow such a strategy, the appearance of a conflict of interest caused by an information market could inhibit its widespread adoption, and thus prevent the organization from realizing its potential for improved forecasts.

This problem is an example of the more general agency problem in which one group (e.g., managers) is unable to fully monitor the efforts of another (e.g., workers) and so needs to align, as much as possible, the incentives of both groups. Agency theory explains how to best organize relationships in which one party (the principal) determines the work, which another party (the agent) undertakes . The theory argues that under conditions of incomplete information and uncertainty, which characterize most business settings, two agency problems arise: adverse selection and moral hazard. Adverse selection is the condition under which the principal cannot ascertain if the agent accurately represents his ability to do the work for which he is being paid. Moral hazard is the condition under which the principal cannot be sure if the agent has put forth maximal effort.

While appealing in principle, a forecast mechanism's usefulness in practice will depend on suitable choices for the mechanism design. In particular, it must encourage sufficient participation in the forecasting to achieve accuracy while

avoiding negative side-effects. Moreover, people might be reluctant to publicize negative expectations for their own group, leading to unwillingness to forecast unanticipated downturns unless the forecasts are made private, which could in turn inadvertently affect actual group production adversely.

To solve these pervasive problems, we introduce a mechanism that avoids both the moral hazard associated with forecasting and the social stigma associated with negative forecasts. The mechanism involves the notion of identity escrow, whereby one's identity is hidden in a crowd in such a way that can only be revealed if sufficient members of that community decide to do so. This can balance the requirements of privacy, to encourage truthful reporting of negative forecasts, with the ability to identify conflicts of interest.

We show that under general conditions, three possible regimes exist. In the first one, members contribute to production but do not participate in the forecasting. For different values of the payoffs members contribute to both production and forecast, while in a third regime members forecast but work against the group's production. We compute the conditions for the appearance of each of these regimes and show how one can adjust the payoffs so as to be in the desirable state, i.e. one in which forecasters contribute to both accurate predictions and production.

2 Model

Consider an organization of N individuals organized hierarchically in groups of various sizes. Setting the parameters of a forecast mechanism can be viewed as a choice made by the organization managers, in the hope of improving forecasts through information aggregation. These parameters include the types of forecasts made, e.g., the group sizes over which individuals make forecasts, and the relative payoffs for the forecasters vs. group production.

Once such a mechanism is created, individuals must then decide whether to participate, based on anticipated payoffs and costs, as well as how much to contribute to group production.

In the remainder of this section, we describe the payoffs of various choices to the organization as a whole and the individuals, thereby producing two coupled games, operating at different time scales.

2.1 Organizational Choices

Consider a situation whereby a company institute a forecasting system with the following design choices:

- size of payouts for correct predictions

  These payouts are a cost for running the forecast. Higher payouts encourage more participation and better predictions.

- the group size n over which people forecast

  Smaller groups will allow more accurate individual predictions but also make it more difficult to aggregate forecasts for the groups into an overall forecast. (The smaller groups are also more likely to have moral hazards since individuals have more influence over them.)

  The company could also change size of production groups, even if just in deciding how large a group to use for profit sharing, e.g., the extremes are everyone gets a fraction of company profit or a bonus based only on your own group's production. I.e., production group size for profit sharing and forecast group size need not be the same.

- privacy of forecasts

  More privacy can encourage truthful reporting of bad news, but is also more likely to create moral hazards.

The payoff for the company is a combination of production, the forecast error and the cost involved in the forecast, which we take to contribute with equal weights:

$$U = f_{\text{coop}} bN - E - C_{\text{organization}} \tag{1}$$

with the first term proportional to the production and E denoting the detriment of forecast errors. The last term $C_{\text{organization}}$ represents the cost of operating the forecast operation (which we take to be dominated by the payments made to individual participants). Here the production payoff is the sum of individual productions, hence proportional to the fraction cooperating with respect to production. If the forecast is used for planning of sales, inventory, etc., it is important to avoid errors either too large or too small in which case E will be large for either sort of error.

The forecast error involves the fraction participating in the forecast mechanism, f_{forecast}, and the group size over which individuals make predictions. In principle, the global accuracy of a forecast should improve when more people participate. This accuracy also depends on the size of the groups over which individual predictions are being made. In particular, individuals have more accurate information about smaller groups they participate in than about larger groups. On the other hand, predicting behavior for the organization as a whole by aggregating forecasts about small group behavior is more difficult, e.g., since such information may not include correlations among the groups and each small group will have fewer participants for each group. We can thus expect overall accuracy to be low when the group size is either too small or too large. Finally, when no one chooses to participate in the forecast, the organization has some default estimation method giving an error that is relatively larger than that possible when people participate.

While the precise functional form will depend on the particular situation, a simple functional form with these properties is

$$E = \frac{n + g^2/n}{1 + N f_{\text{forecast}}} \tag{2}$$

where n is the size of the groups over which forecasts are made, and N is the total number of individuals in the organization and g is the optimal group size for prediction. While this functional form is somewhat arbitrary, it does contain the qualitative behaviors, in particular is nonnegative over the relevant range $0 < n < N$ and has a minimum, as a function of n, at $n = g$.

2.2 Individual Choices

In our model, the individuals make two choices, which we take to be binary:

- production: cooperate or defect
- forcast: cooperate (i.e., participate accurately) or defect

The overall utility for the individual is the sum of the reward for production and the payoff from forecasts.

There is a cost c to cooperate, but we suppose the organization payoffs and structure is such as to induce cooperation when considering production costs alone [1]. This assumption allows us to examine the conditions under which introducing a forecasting method to an already functioning organization can produce useful forecasts while maintaining the prior high productivity. In terms of the payoffs used here, we need $b/n > c$ (where n is the size of the group). The individual payoff is

$$U = U_{\text{production}} + U_{\text{forecast}} \tag{3}$$

We focus on the portion of the payoff for production spread among the group (i.e., a form of profit sharing). For the production, each person receives a payoff of bf_{coop} where f_{coop} is the fraction of the production group that chooses to cooperate. For the individual, there is a cost c for cooperating. Thus the choice is between the cost c for cooperating and the loss of shared benefit, b/n, for defecting [7].

In the case of forecasting the organization gives a payoff for accurate predictions (if the person participates, i.e., "cooperates"). The accuracy of an individual's predictions decreases with the size of the group, reflecting the decreasing information about the entire group available to its members.

There is a social cost to participate in the forecasting exercise, c_p when revealing negative bets about the future. This cost applies only if trades are public. In addition, if a negative forecaster encourages a slowdown in production (same as defecting in the production side of the game) and his/her identity is revealed he/she incurs an extra cost, C.

We characterize the likelihood a forecaster's identity is revealed by two probabilities: P_1 and P_2 as the probability the identity is revealed when the individual is not engaged in a moral hazard or is, respectively. The expected cost for an individual participating in the forecasting but not working against the organization is thus $c_p P_1$ while that for one who is working against the organization is $(c_p + C)P_2$.

Offsetting these costs is the utility accrued for accuracy. One component of this accuracy is the ability of the individual to directly influence the production of the group by influencing what other do. That is, we suppose the accuracy of forecasts improves if the individual's production choice matches the forecast position. Thus, if the individual expects production to be low and makes such a forecast, then reducing production and encouraging others in the group to do likewise, will result in higher accuracy than if the individual cooperates to increase the group's production.

As a specific form, we take the individual accuracy, and hence individual payoff for participating in the market to be

$$M_{\pm} = \frac{A}{n}(1 \pm \alpha) \tag{4}$$

where the plus sign is taken when individual's production choice matches his/her forecast, and the minus sign otherwise. Here $0 \leq \alpha < 1$ characterizes the difference in accuracy between these cases. The parameter A characterizes the payoff given to individuals for accurate forecasting.

The utilities for the choices available to an individual are summarized in Table 1.

3 Avoiding Moral Hazards

In what follows we present a mechanism based on notions of group detection and anonymous reports that makes it unlikely for a forecaster to depress production so as to meet a pessimistic forecast while still retaining privacy. This privacy enables the pessimistic forecasts without incurring the social cost of announcing

	forecast	
production	cooperate	defect
cooperate	$M_- - c_p P_1 + b f_{\text{coop}} - c$	$b f_{\text{coop}} - c$
defect	$M_+ - (c_p + C) P_2 + b(f_{\text{coop}} - 1/n)$	$b(f_{\text{coop}} - 1/n)$

Table 1: Payoffs and choices for individuals in the situation in which predictions are for lower performance (the situation relevant for moral hazards). Without the forecasts, we assume payoffs are such as to encourage cooperation, i.e., $b/n > c$.

bad news.

One possibility would be for the company to keep track of the actions of all forecasters so as to make sure that they don't act against the interest of the company. Unfortunately, just as it is difficult for companies to monitor individual efforts, i.e. the agency problem [5] the same is true for monitoring forecasters. Instead, members of the group are more likely to notice any actions that are detrimental to the common good. But on the other hand, this would eliminate privacy and therefore make people reluctant reveal pessimistic forecasts.

What is needed instead is a mechanism which addresses all these problems, and which we describe below.

3.1 Identity Escrow

One problem with having a forecasting mechanism within organizations has to do with the degree to which the identities of the forecasters are public, both to the group in charge of making the predictions, and to the organization as a whole. To encourage people to reveal accurate information, even if it reflects a negative outlook on the group, one may want to insure a certain level of privacy to the forecasters, so that they can use their anonymity to reveal perceived truths that might not be pleasant for others to learn about.

But the very anonymity of the forecaster also creates the a moral hazard,

since he or she are now in a position to influence the production of the units under their control so as to fulfill the prediction. If that would mean increased production to meet an ambitious quota it would be acceptable, but the alternative, inducing slower production levels to come close to prediction is not.

One mechanism uses cryptographic keys to represent forecasters which can then be used to provide an effective identity escrow mechanism. This is accomplished by associating each forecast and the identity of who generates it to a public key that can be read only with its associated private key. These keys can keep the forecast private. However, the mechanism should allow members of the group to identify forecasters if there is sufficient suspicion that they are acting against the interest of the group. Threshold cryptography allows to do this by breaking each private key into several pieces each of which is given to members of the production group. These pieces have the property that at least k members of the group are required to reconstruct the key, and reveal the person engaging in negative behavior. Any subgroup smaller than k individuals obtains no information at all about the key.

Another variation gives part of the keys to the company itself, so that identities are revealed only if both k group members and the management agree to do so. For instance the company might choose to act only if it noticed some unusual trades. This variation solves the problem that while group members observe only production directives but have no clue as to the position taken by the forecaster, the company is aware of the forecasts but not the details of the production directives. Thus the moral hazard is most prominent when there is a large negative forecast and a consequent decrease in production by the group. Either one by itself is not necessarily a problem.

In either case detection of a moral hazard by a subset of those holding the pieces of the key triggers an action that reconstitutes the public key corresponding to the suspicious individual and therefore identifies him/her publicly. With this key one can then read messages associated with a particular forecast and discover whether the particular person is indeed acting against the interest of the company the public key which was used.

A simple way to accomplish this key splitting without requiring all parts to be assembled together in order to get it back is to use a threshold scheme that works in the following fashion [9].

1. The public key identifying an individual forecaster is expressed as a secret integer $I > 0$ and is distributed among the n members of a group.

2. A prime p is chosen such that $p > I$ and a coefficient a_0 is defined as $a_0 = I$.

3. A server S selects $t-1$ random, independent coefficients $a_1, ...a_{t-1}$, such that $0 \leq a_j \leq (p-1)$. This defines a random polynomial $f(x) = \sum a_j x^j$.

4. The server S computes $I_i = f(i) \bmod p$, $1 \leq i \ \leq n$ (or for any n distinct points i, $1 \leq i \leq (p-1)$, and securely transfers the piece I_i to user P_i, along with the public index i.

5. Any group of t or more users can pool their pieces, which provide t distinct points $(x, y) = (i, I_i)$. This allows the computation of the coefficients a_j, $1 \leq j \leq (t-1)$ of $f(x)$, via the Lagrange interpolation scheme, which we explain below. Once this is done the secret identity is recovered by noting that $f(0) = a_0 = I$.

This technique is based on the fact that the coefficients of an unknown polynomial $f(x)$ of degree at most t, defined by the set of points (x_i, y_i), with $1 \leq i \leq t$ are given by the Lagrange interpolation formula:

$$f(x) = \sum_{i=1}^{n} \prod_{1 \leq j} \frac{x - x_j}{x_i - x_j} \tag{5}$$

since $f(0) = a_0 = I$, the secret identity can be then expressed as

$$I = \sum_{i=1}^{n} c_i y_i \tag{6}$$

where

$$c_j = \prod_{1 \leq j \leq t} \frac{x_j}{x_j - x_i} \tag{7}$$

Thus each member of the group can compute I as a linear combination of t pieces y_i, since the coefficients c_i are non-secret constants. Actually, for a fixed group of t members they can be pre-computed.

This threshold scheme, due to Shamir [10], has very useful properties for privacy, use and scalability. Equally important, even with infinite computational power, no members of the group can learn anything more than the length of the message, which each of them knows already. This is as secure as a one-time pad, for an attempt at exhaustive search will reveal that any conceivable message could be the secret key.

While threshold cryptography can in principle solve the moral hazard problem involved in organizational forecasting, one still needs to determine the minimum number of individuals that are needed in order to decode a given key without compromising the forecaster's sense of privacy. This is addressed in the next section.

3.2 Choosing a Threshold

A possible mechanism for achieving this relies on Condorcet's theorem, which states that if individuals have a given probability p of accurately detecting a detrimental action for the group, then it is possible to increase that probability by collectively aggregating them in what is analogous to a majority vote, but with a winning threshold not necessarily 50%.

To see how this works, suppose we pick a threshold such that t out of n need to detect the problem to reveal the identity. Assuming independent observations, the probability $P(t, p)$ that this threshold is achieved is given by the upper tail of the binomial distribution

$$P(t,p) = \sum_{i=t}^{n} \binom{n}{i} (1-p)^{n-i} p^i \tag{8}$$

Eq. (8) relates the probability a forecaster's identity is revealed to the choice of threshold and the probability members of the group notice a potential problem. Ideally, we would like to find a threshold large enough so people feel com-

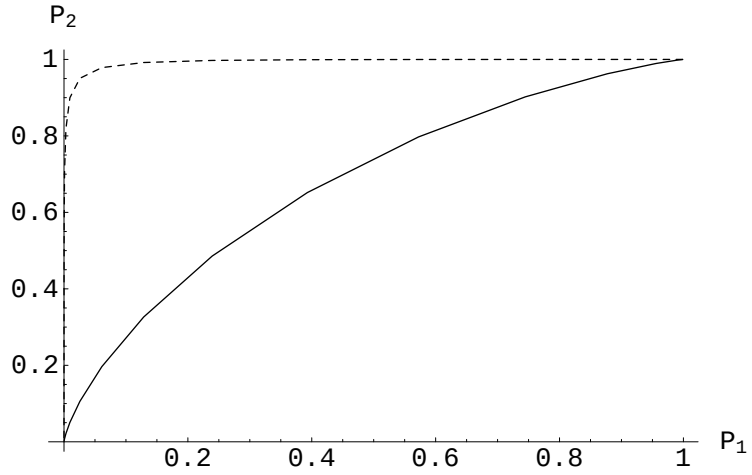

Figure 1: Behavior of P_2 vs. P_1 as the threshold varies from 0 to n. The curves show cases with high and low discrimination on the part of the group members.

fortable their privacy is unlikely to be violated when they do not act against the group, but otherwise they are likely to be caught. Whether this is possible depends on the ability of group members to discriminate among these situations. Let p_1 be the probability a group member detects a problem with a particular individual, or otherwise chooses to act to reveal that person's forecast, when in fact the individual did not work against the group. Conversely, when the individual does act against the group, we suppose the members have a somewhat higher probability p_2 to notice the problem.

For a given choice of threshold t, we then have $P_1 = P(t, p_1)$ and $P_2 = P(t, p_2)$. If group members have a large ability to discriminate between members contributing and defecting, then $p_1 \ll p_2 \approx 1$. On the other hand, if members lack this ability, p_1 will be only slightly smaller than p_2.

The choice of threshold is a tradeoff between privacy and catching bad behavior. We need the threshold larger than a typical clique size, so that people feel that there are multiple independent decisions before the identity is revealed. But the threshold should also be small enough to likely catch adverse actions made to match negative forecasts.

4 Behavior

For simplicity, our model takes the members of the group to have identical preferences. In this case, if there is a simple equilibrium either all or none of the members will participate in the forecasting and in the production. Since we assume the organization has resolved the social dilemma when there is no forecast (hence our restriction to $b/n > c$ in the payoffs described above), this leaves three possible equilibria:

- A: *no forecasts*, members contribute to production but do not participate in the forecasting, $f_{\text{coop}} = 1$, $f_{\text{forecast}} = 0$.
- B: *productive forecasts*, members contribute to production and forecast, $f_{\text{coop}} = f_{\text{forecast}} = 1$.
- C: *moral hazard*, members forecast but work against the group's production, $f_{\text{coop}} = 0$, $f_{\text{forecast}} = 1$.

Ideally, the organization would like to select a mechanism so the group equilibrium is case B. Within this model, this mechanism choice amounts to choices for the threshold t, payoff for forecasts A and penalty when identifying individuals working against the organization C. To encourage participation, the payoff may need to be sufficiently large. But then to have any possibility of preventing moral hazard, the penalty must also be correspondingly large. We take these values to be proportional, i.e., $C = \gamma A$. The choice among these equilibria is then made according to which has the largest utility to the individuals in the group, using the expressions of Table 1.

When considering forecast issues an organization has to balance the goal of accurate predictions with the forecaster's incentive to participate but not to work against the organization.

From the organizational point of view, there is a tradeoff between the size of the group's output over which the forecast is made (the smaller the more accurate) and the ability to aggregate the forecast (the bigger the better), as modelled in Eq. (2). Also, there is the additional tradeoff associated with the

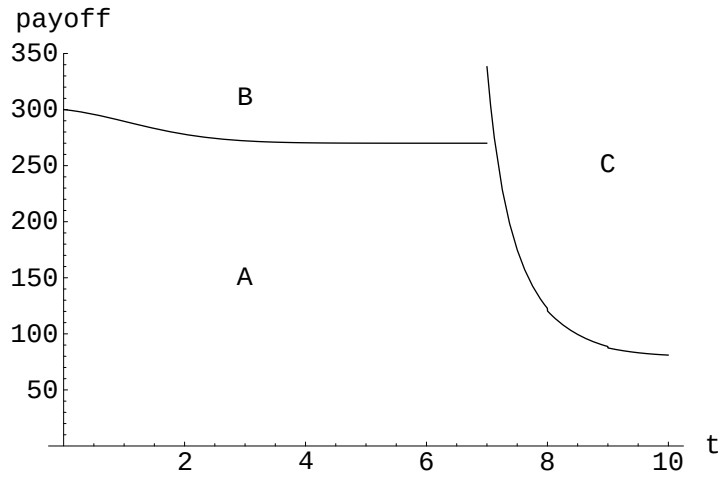

Figure 2: Equilibrium regimes for parameters $n = 10$, $p_1 = 0.1$, $p_2 = 0.6$, $\alpha = 0.5$, $c_p = 1.0$, $c = 3$, $\gamma = 0.2$, $b = 10$ as threshold t and payoff A vary.

minimum number (the threshold) of individuals that can activate revelation of a given forecaster's identity. If the threshold is too low, it will be activated by too few individuals and therefore people will be less likely to participate. On the other hand, if the threshold is too high then the moral hazard will be less likely to be detected as it requires a large number of many suspicious observers.

In addition, one must consider the utilities accrued by both the forecasters and the organization, since that determines the choices for the mechanism.

In order to illustrate these tradeoffs, we present an example to show how the behavior changes as the parameters are varied.

One example of the tradeoffs is shown in Fig. 2. In this case, suitable choices for payoff and threshold allow a regime in which people participate in the forecast but do not work against the group. By contrast, Fig. 3 shows a situation in which such an ideal outcome does not exist.

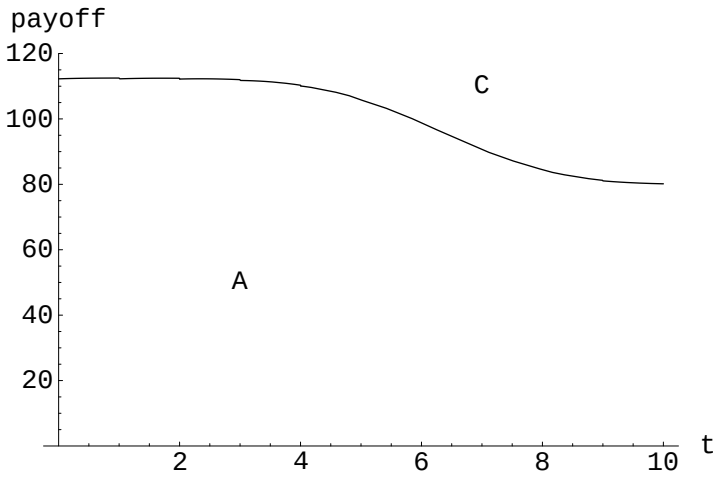

Figure 3: Unavoidable moral hazard: in this case, equilibrium B is never selected by the group. The parameters are $n = 10$, $p_1 = 0.1$, $p_2 = 0.6$, $\alpha = 0.5$, $c_p = 1.0$, $c = 3$, $\gamma = 0.02$, $b = 10$ as threshold t and payoff A vary.

5 Discussion

In this paper, we introduced a mechanism that avoids both the moral hazard associated with forecasting and the social stigma associated with negative forecasts. The mechanism involves the notion of identity escrow, whereby one's identity is hidden in a crowd in such a way that can only be revealed if sufficient members of that community decide to do so. This can balance the requirements of privacy, to encourage truthful reporting of negative forecasts, with the ability to identify conflicts of interest.

We showed that under general conditions, three possible regimes exist. In the first one, members contribute to production but do not participate in the forecasting. For different values of the payoffs members contribute to both production and forecast, while in a third regime members forecast but work against the group's production. We compute the conditions for the appearance of each of these regimes and show how one can adjust the payoffs so as to be in the desirable state, i.e. one in which forecasters contribute to both accurate

predictions and production.

Our model involves parameters for some group behaviors that are not readily measured directly. These include the size of the threshold necessary to instill confidence that privacy of forecasts is protected, and the likelihoods of detection. As one approach to estimating appropriate values, the company could try various pilot systems with differing choices of forecast payoff, group sizes over which forecasts are made, and thresholds. Observing the resulting participation can give some indication of the the appropriate parameters.

Other indications for reasonable thresholds could involve noticing the typical clique sizes in the group, e.g., as estimated from organizational structure or informal networks revealed through web page links, etc. With such estimates, we could consider the typical clique size as a lower bound on the threshold, so individuals feel several independent decisions would be needed to reveal their identity. On the other hand, estimating how visible undesirable behavior is likely to be (i.e., when an individual works against the interests of the group) will indicate a plausible upper bound on the threshold.

While we considered a simple case of uniform groups, a more general possibility to account for different clique sizes: have a weighted threshold so more people from larger cliques are required, i.e., they have lower weight. This provides a flexible adjustment to the identity escrow to account for different interaction groups. Such groups could, for instance, be identified, at least approximately, through electronic information exchanges such as common web links.

The methodology presented in this paper is yet another example of the value that can accrued when combining cryptographic techniques with economic incentive designs [6]. In this case,our mechanism allows for limiting the revealed information to be just that required to maintain desired incentives. This is to be contrasted with the choice of forcing the mechanism to pick one of the extreme choices of full privacy or completely public information.

There are a number of possible extensions to this mechanism that will make even more suitable for realistic applications to organizational forecasting. For example, our discussion of individual behavior used the expected cost of being

caught. While this is appropriate for risk neutral individuals, an interesting direction for future work is to determine how the behaviors change if some group members are risk adverse (so less likely to join the forecast) or risk seeking.

Another question is the extension to continuous level of production and variation in accuracy of forecasts. For instance, by expending additional effort individuals may be able to improve their forecast accuracy, which requires the consideration of the equilibrium values of efforts devoted to production and forecasting. These could, of course, vary among members of the forecasting group, resulting in mixed strategies.

Regardless of the simplicity of the model, this approach, which resorts to both mechanism design theory and cryptography, offers a flexible methodology for tackling problems that involve precise control of information exchange and incentive issues for participation. We thus expect that similar methods can be applied to a variety of other organizational issues and the deployment of automated monitoring of group assessments while maintaining privacy.

Acknowledgements

We thank M. Franklin for helpful discussions.This material is based upon work supported by the National Science Foundation under Grant No. 9986651.

References

[1] A. A. Alchian and H. Demsetz. Production, costs and economic organization. *American Economic Review*, 62:444–458, 1972.

[2] L. Anderson and C. Holt. Information cascades in the laboratory. *American Economic Review*, 87:847–862, 1997.

[3] C. Camerer and K. Weigelt. Information mirages in experimental asset markets. *J. of Business*, 64:463–493, 1991.

[4] Kay-Yut Chen, Leslie R. Fine, and Bernardo A. Huberman. Forcasting uncertain events with small groups. In *Proc. of the 3rd ACM Conference on E-commerce*, pages 58–64, October 2001.

[5] K. Eisenhardt. Agency theory: An assessment and review. *Academy of Management Review*, 14(1):57–74, 1989.

[6] Bernardo A. Huberman, Matt Franklin, and Tad Hogg. Enhancing privacy and trust in electronic communities. In *Proc. of the ACM Conference on Electronic Commerce (EC99)*, pages 78–86, NY, 1999. ACM Press.

[7] Bernardo A. Huberman and Christoph H. Loch. Collaboration, motivation and the size of organizations. *Journal of Organizational Computing and Electronic Commerce*, 6:109–130, 1996.

[8] M. Noeth and M. Weber. Information aggregation with random ordering: Cascades and overconfidence. Technical report, Univ. of Mannheim, 1998. Presented at the Summer 1998 ESA Meetings.

[9] Bruce Schneier. *Applied Cryptography.* John Wiley, 2nd edition, 1996.

[10] A. Shamir. How to share a secret. *Communications of the ACM*, 24:612–613, 1979.

[11] S. Sunder. Markets for information: Experimental evidence. *Econometrica*, 60:667–695, 1992.