

# Composition of Markets with Conflicting Incentives

Stanko Dimitrov  
University of Michigan  
Department of Industrial and Operations  
Engineering  
sdimitro@umich.edu

Rahul Sami  
University of Michigan  
School of Information  
rsami@umich.edu

## ABSTRACT

We study information revelation in scoring rule and prediction market mechanisms in settings in which traders have conflicting incentives due to opportunities to profit from the market operator's subsequent actions. In our canonical model, an agent Alice is offered an incentive-compatible scoring rule to reveal her beliefs about a future event, but can also profit from misleading another trader Bob about her information and then making money off Bob's error in a subsequent market. We show that, in any weak Perfect Bayesian Equilibrium of this sequence of two markets, Alice and Bob earn payoffs that are consistent with a minimax strategy of a related game. We can then characterize the equilibria in terms of an *information channel*: the outcome of the first scoring rule is as if Alice had only observed a noisy version of her initial signal, with the degree of noise indicating the adverse effect of the second market on the first. We provide a partial constructive characterization of when this channel will be noiseless. We show that our results on the canonical model yield insights into other settings of information extraction with conflicting incentives.

## Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

## General Terms

Economics, Theory

## Keywords

prediction markets, scoring rule, market scoring rule, conflicting incentives

## 1. INTRODUCTION

Scoring rules and markets are commonly deployed incentive mechanisms to incentivize participants to reveal their

private information. A recent class of prediction markets, Market Scoring Rules (MSRs), can be viewed as a sequential form of a scoring rule, and retain the incentive properties of individual scoring rules. Thus, they incentivize risk-neutral agents maximizing their expected reward to reveal their beliefs about future events. This property has led to MSR prediction market deployments in many corporate settings [7]. In this paper we initiate the study of how the composition of conflicting scoring-rule based incentive mechanisms affect information revelation in the individual markets.

Due to their attractive incentive properties, prediction markets are promoted as a fast and accurate method of forecasting future events. Empirical research has confirmed their accuracy in practice [18]. This in turn has led to proposals of prediction market to be used widely as a decision tool, for subjects ranging from national policy [11] to corporate governance [1]. However, there is one important challenge in making this leap: Theoretical analysis, and most empirical studies, of prediction market accuracy has been done in settings where prediction markets are not transparently integrated into a decision process, and hence participants do not have an outside interest in the prediction market prices. If and when prediction markets are integrated into decision processes, traders may have very strong external incentives to achieve certain market prices. Additionally, if decision makers rely on information gathered from prediction markets for their own trading actions, they may induce an incentive for traders to mislead them.

In principle, if the external incentive structure is known, one can correct the scoring rule for this effect, for example by using the techniques proposed by Shi *et al.* [14]; another seemingly attractive alternative is to provide additional markets that extract all relevant information. In practice, we never have complete markets, and incentive methods are deployed by different entities in an ad-hoc rather than coordinated fashion. Further, budgetary constraints may rule out the deployment of a scoring rule with stakes sufficient to counteract an external incentive. In this paper, we adopt the point of view that the existence of conflicting incentive markets is unavoidable. We therefore take a descriptive approach to characterize the effects of these conflicting incentives, before deriving prescriptive insights. In this way, we hope to lay a foundation for a more general theory of *composition* of scoring-rule mechanisms, and incentive mechanisms more broadly.

In this paper, we formulate and analyze a simple model of a class of conflicting incentive situation. Our core model is a three-stage game with two agents, each with private infor-

mation, and two incentive mechanisms. In the first stage, Alice is offered a scoring rule or a market scoring rule to elicit her private information relevant to forecasting a random variable  $U$ . Upon observing Alice’s report on  $U$ , Bob updates his belief on Alice’s private information. In the second stage, Bob uses this information to trade in a market scoring rule market for forecasting a different random variable  $V$ . Finally, Alice upon observing Bob’s report in this second market has the opportunity to trade in it by revising Bob’s forecast. Thus, Alice has a conflicting incentive in the first-round market: She would maximize her immediate profit by reporting honestly, but may be able to gain more by damaging Bob’s forecasts and then making a profit by correcting them. Bob is aware of this conflict of interest, and may update his beliefs defensively.

This basic analysis unit provides insight into other situations. It applies directly to situations in which prediction market operators, or other traders, may rely on the information gleaned to trade in other markets. It also applies to the analysis of strategies within a single market: In this case, a security  $V$  is identical to the security  $U$ , but represents a future round of trading. As such, our model generalizes the model in Chen *et al.* [5]. Our model also applies to the analysis of settings in which the conflict of interest is not a market at all, but rather, some unspecified consequence of Bob’s decision. These applications are described in detail in Section 6.

We characterize the equilibria of the 3-stage game. Note that this is not a constant-sum game, because different strategies for Alice lead to different values of expected profit from the  $U$ -market, and hence to different values of total profit. Somewhat unexpectedly, we find that in all weak-perfect Bayesian equilibria, the payoff earned by Alice in the  $U$ -market is the same. Consequently, the total payoffs earned by Alice and Bob are the same in every equilibrium. There may be multiple equilibria, but the payoffs are unique, and the amount of information revealed in the compromised market is unique across all equilibria.

This uniqueness result enables an information-theoretic view of the conflict of interest situation. The existence of the second market induces Alice and Bob to behave (in the first two stages) as if Alice had received a noisy signal. The noise level is unique in equilibrium, and depends on the nature of the forecast variables  $U$  and  $V$ , as well as the relative scale. We formulate an easy-to-compute necessary condition for this information channel to be noiseless, and derive characterizations for combinations of forecast variables.

The structure of the rest of this paper is as follows: In Section 2 we discuss the related work. In Section 3, we describe the model we consider in the two sequential market setting; we analyze this model in Section 4. In Section 5, we consider the role of information in equilibrium. In Section 6 we describe different applications and extensions of our result, and we conclude and discuss directions for future work in Section 7.

## 2. RELATED WORK

In this paper we considered market scoring rule (MSR) markets introduced by Hanson [10]. Hanson proved that risk-neutral traders that do not consider future payoffs have an equilibrium strategy of reporting their true beliefs. A recent strand of literature [5, 6, 8] examines what happens when traders trade multiple times in an MSR market and

do consider future payoffs. The authors show that there is a class of information distributions in which the honest strategy is not an equilibrium, and another class in which the honest strategy is an equilibrium of the multi-stage game. The 3-stage model described by Chen *et al.* [5] can be viewed as a special case of the model in this paper, in which the two forecast variables  $U$  and  $V$  are identical.

The role of manipulation in prediction markets has been studied empirically and experimentally; we refer readers to the survey paper by Wolfers and Zitzewitz [18] for further pointers to this literature. In summary, Wolfers and Zitzewitz find that manipulation does not have any noticeable effect on price, except for a short period of time. Hanson *et al.* [9] consider the effect of external incentives on manipulation in prediction markets. The authors show that even with these external incentives, “manipulation has no significant effect on the accuracy of price.”

Shi *et al.* [14] proposed the principal-aligned prediction market mechanisms to remove any external incentives traders may have on effecting the outcome of the traded event. This is in a sense orthogonal to our approach of describing and analyzing the effect of an unavoidable conflict between incentive schemes. The approach we present can help to get around the weaknesses of the principal-aligned market mechanism: it may lead to unbounded subsidies in general, and it may be difficult to implement if the external incentives are not precisely known.

Market scoring rule markets are based on proper scoring rules first introduced by Brier [4], who introduced the quadratic proper scoring rule to measure the accuracy of weather forecasters. The logarithmic scoring rule, the basis of the logarithmic market scoring rule, was later introduced again to assess weather forecasters [16, 17]. Lambert *et al.* [12] show that a larger class of functions may be used to elicit different properties of a distribution, rather than just a point estimate of the mean. As observed by Agrawal *et al.* [2], a prediction market can be constructed by a larger class of functions. The techniques we present here can be extended to a subset of the class discussed by Agrawal *et al.* because we only use the concavity of the  $\log(\cdot)$  function to show our results.

## 3. CANONICAL MODEL

In this section, we detail our particular model, with two sequential markets.

There are two players, Alice and Bob. Alice has an information signal described by a random variable  $X$ , with values from a set  $\mathcal{X}$ ; let  $x$  denote the actual realization of  $X$ . Likewise, Bob has an information signal described by a random variable  $Y$ , with values from a set  $\mathcal{Y}$ ; let  $y$  denote the actual realization of  $Y$ .

There are two other random variables of forecasting interest. Random variable  $U$  has an unknown value  $u \in \{0, 1\}$ . Random variable  $V$  has an unknown value  $v \in \{0, 1\}$ . Further, there is a prior joint probability distribution over  $X, Y, U$ , and  $V$ . Alice and Bob know the prior probability distribution; at the start of our analysis, they also receive their own signal realization. However, they do not know the realization of the other party’s signal.

We assume that there are correlations between the information signals and the forecast variables, so that making an optimal forecast requires information acquisition/aggregation. ( $X$  and  $Y$  may be independent of each other, however.)

Given the structure of the prior distribution, there will be an optimal aggregate. We define  $f_{xy} \stackrel{\text{def}}{=} \mathcal{P}_0[u = 1|X = x, Y = y]$ . Likewise, define  $f_x \stackrel{\text{def}}{=} \mathcal{P}_0[u = 1|X = x]$  and  $f_y \stackrel{\text{def}}{=} \mathcal{P}_0[u = 1|Y = y]$ . For variable  $V$ , we can similarly define values  $g_{xy}, g_x, g_y$ . Finally, we let  $f_0 \stackrel{\text{def}}{=} \mathcal{P}_0[u = 1]$  and  $g_0 \stackrel{\text{def}}{=} \mathcal{P}_0[v = 1]$  denote the prior expectations of  $U$  and  $V$ .

We make one additional technical assumption, that the signal  $Y$  is generically  $V$ -distinguishable:

**DEFINITION 1.** *The signal  $Y$  is generically  $V$ -distinguishable if, given any probability distribution  $p(x)$  over  $\mathcal{X}$ , the expected value of  $V$  is different for any two different values  $y_1, y_2$ :*

$$\forall y_1 \neq y_2 \quad \sum_x p(x)g_{xy_1} \neq \sum_x p(x)g_{xy_2}$$

This is a nontrivial assumption that allows us to exclude degenerate cases in our analysis. We explore the impact of this genericity assumption not holding in Section 6.1.

**Markets.** Now, let us consider two markets that are deployed to forecast  $U$  and  $V$  respectively. Both markets are organized as market scoring rules: In the  $U$ -market, a trader who changes the forecast from  $r$  to  $q$  earns a positive or negative profit  $\pi_U(u, r \rightarrow q)$  after the true outcome  $u$  is revealed, according to a market scoring rule:

$$\pi_U(u, r \rightarrow q) = \begin{cases} \lambda_U[\log q - \log r] & \text{if } u = 1 \\ \lambda_U[\log(1 - q) - \log(1 - r)] & \text{if } u = 0 \end{cases}$$

where  $\lambda_U$  is a positive constant. We extend this notation to reflect the expected profit given a probabilistic outcome: define

$$\pi_U(p, r \rightarrow q) \stackrel{\text{def}}{=} p\pi_U(1, r \rightarrow q) + (1 - p)\pi_U(0, r \rightarrow q)$$

Similarly, for a mover in the  $V$ -market, the profit is given by

$$\pi_V(v, r \rightarrow q) = \begin{cases} \lambda_V[\log q - \log r] & \text{if } v = 1 \\ \lambda_V[\log(1 - q) - \log(1 - r)] & \text{if } v = 0 \end{cases}$$

where  $\lambda_V$  is (possibly different) positive constant. As above, we extend this to define a function  $\pi_V(p, r \rightarrow q)$  for all  $p \in [0, 1]$ .

Without loss of generality, we assume that  $\lambda_U = 1$ ; only the ratio between  $\lambda_U$  and  $\lambda_V$  is important for our analysis.

**Sequence of Trade.** We initially assume that the following sequence of actions takes place:

- Stage 1: Alice trades in the  $U$ -market
- Stage 2: Bob trades in the  $V$ -market
- Stage 3: Alice has an opportunity to trade in the  $V$ -market

**Strategies and Payoffs.** The strategic choices that are made are as follows:

- In stage 1, Alice can move the  $U$ -market price from  $f_0$  to a new value  $r \in [0, 1]$ . The move is in general dependent on the signal  $x$  Alice observed, and perhaps randomized. Thus, we can denote Alice's strategy in

stage 1 by a function  $\sigma : \mathcal{X} \rightarrow \mathcal{D}$ , where  $\mathcal{D}$  is the space of distributions over  $[0, 1]$ .

We make the following technical assumption on the strategies played by Alice: We assume that, for all  $x$ , the support of  $\sigma(x)$  is contained in a pre-specified finite set  $\mathcal{R}$ . In other words, we restrict the set of moves available to points in  $\mathcal{R}$  instead of all of  $[0, 1]$ . We make this assumption for analytical convenience; of course, any real market is likely to have a finite number of possible positions as well. The set  $\mathcal{R}$  can be arbitrarily large, and is not chosen for its strategic properties (*e.g.*, we do not assume that we can rule out certain points to reduce deviations from honesty.) We denote by  $\sigma_x(r)$  the probability of moving to  $r$  when the drawn signal is  $x$ .

For any given strategy  $\sigma$ , and any position  $r$  in the support of  $\sigma$ , we let  $\mathcal{P}_{\sigma, r}$  denote the posterior probability distribution of the random variables, conditioned on knowledge of  $\sigma$  and  $r$ .

- In stage 2, Bob moves market  $V$  from  $g_0$  to  $t \in [0, 1]$ . As Bob trades only once in our model, and as the log-scoring rule is honest for myopic traders, Bob's optimal moves  $s$  and  $t$  are completely determined by Bob's belief about Alice's signal after observing her move  $r$ . Thus, we describe Bob's strategy indirectly by his *belief function*  $\mu : r \mapsto \mu_r \in \mathcal{D}(\mathcal{X})$ , where  $\mathcal{D}(\mathcal{X})$  is the set of distributions over  $\mathcal{X}$ . Given a strategy  $\mu$ , observed first-stage move  $r$ , and his own signal  $y$ , Bob will move the  $V$ -market to  $t_{ry} = [\sum_x \mu_r(x)g_{xy}]$ .
- In the third stage, Alice's move will depend on the extent to which she can infer Bob's move from the second-stage market. Assuming generic  $V$ -distinguishability (Definition 1), Alice can distinguish all of Bob's signals with knowledge of Bob's second-round strategy  $\mu$ . In this case, Alice's third-stage move is to move the  $V$ -market from  $t_{ry}$  to  $g_{xy}$ .

Given a strategy pair  $(\sigma, \mu)$ , Bob's payoff can be written as:

$$\pi_B(\sigma, \mu) = \sum_x \sum_y \mathcal{P}_0(X = x, Y = y) \left[ \sum_{r \in \sigma(x)} \sigma_x(r) \pi_V(g_{xy}, g_0 \rightarrow t_{ry}) \right]$$

In any setting in which Alice knows  $\mu$  in the final stage, she will correct the price in the  $V$ -market to the optimal value  $g_{xy}$ . We define

$$\hat{\pi}_A(\sigma, \mu) \stackrel{\text{def}}{=} \text{Eff}_U(\sigma) + \text{Eff}_V - \pi_B(\sigma, \mu) \quad (1)$$

and note that Alice's profit  $\pi_A$  will be equal to  $\hat{\pi}_A$  if Alice knows  $\mu$  in the final stage; this latter condition will be true in equilibrium. Here,  $\text{Eff}_U(\sigma)$  is the efficiency of the  $U$ -market (*i.e.*, expected payoff in the  $U$  market) conditional on the strategy  $\sigma$ . It is given by:

$$\text{Eff}_U(\sigma) = \sum_{x \in \mathcal{X}} \sum_{r \in \mathcal{R}} \sigma_x(r) \lambda_U [\pi_U(f_x, f_0 \rightarrow r)]$$

Likewise,  $\text{Eff}_V$  is the efficiency in the  $V$  market, given by:

$$\text{Eff}_V(\sigma) = \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \lambda_V [\pi_V(g_x, g_0 \rightarrow g_{xy})]$$

Under the genericity assumption, Alice always moves the  $V$  market to the optimal value, and hence  $\text{Eff}_V$  is independent

of  $\sigma$  and  $\mu$ . For notational simplicity, in the remainder of this section, we abuse notation by using  $\pi_A(\sigma, \mu)$  to denote the quantity on the RHS of equation 1.

$\text{Eff}_U(\sigma)$  is linear in  $\sigma$ .  $\pi_B(\sigma, \mu)$  is also linear in  $\sigma$  for any fixed  $\mu$ . Thus,  $\pi_A$  (as given in equation 1) is linear in  $\sigma$ .

The solution concept we use is that of weak Perfect Bayesian Equilibrium (wPBE). Informally,  $(\sigma, \mu)$  form a wPBE profile if along any possible path of play that is played with positive probability, the beliefs  $\mu$  are consistent with Alice's strategy  $\sigma$ , and  $\sigma$  is the optimal strategy given the beliefs  $\mu$ . For a technical definition we refer the reader to the book by Mas-Colell *et al.* [13].

#### 4. STRATEGIC ANALYSIS OF MODEL

In this section, we show that the wPBE of the 3-stage conflict of interest game can be characterized in terms of the minimax value of a related game. The main results, in Theorems 5 and 6, show that the equilibrium payoffs are unique and is consistent with minimax play.

The first observation we make is that the space of pairs  $(\sigma, \mu)$  is closed, convex, and compact, as it is the product space of probability simplices. This shows the existence of a *maximin pair*, defined as follows:

**DEFINITION 2.** A maximin pair  $(\sigma, \mu)$  is a pair that solves the following optimization:

$$\begin{aligned}\sigma &\in \operatorname{argmax}_{\sigma'} \min_{\mu'} \pi_A(\sigma', \mu') \\ \mu &\in \operatorname{argmin}_{\mu'} \pi_A(\sigma, \mu')\end{aligned}$$

**DEFINITION 3.** A minimax pair  $(\sigma, \mu)$  is a pair that solves the following optimization:

$$\begin{aligned}\mu &\in \operatorname{argmin}_{\mu'} \max_{\sigma'} \pi_A(\sigma', \mu') \\ \sigma &\in \operatorname{argmax}_{\sigma'} \pi_A(\sigma', \mu)\end{aligned}$$

**LEMMA 1.** The minimax and maximin values of  $\pi_A$  are the same. There exists a pair  $(\sigma, \mu)$  that is both a minimax pair and a maximin pair.

**PROOF.** We start by observing  $\pi_A(\cdot, \cdot)$  is continuous. We now show that for a fixed  $\mu$ ,  $\pi_A(\sigma, \mu)$  is a concave function of  $\sigma$ , and, for a fixed  $\sigma$ ,  $\pi_A(\sigma, \mu)$  is a convex function of  $\mu$ . Then using Sion's minmax theorem [15], we can say  $\max_{\sigma'} \min_{\mu'} \pi_A(\sigma', \mu') = \min_{\mu'} \max_{\sigma'} \pi_A(\sigma', \mu')$ , showing the desired result.

We start by writing down  $\pi_A(\sigma, \mu)$ :

$$\begin{aligned}\pi_A(\sigma, \mu) &= \text{Eff}_U(\sigma) + \text{Eff}_V - \pi_B(\sigma, \mu) \\ &= \sum_{x \in \mathcal{X}} \sum_{r \in \mathcal{R}} \sigma_x(r) \lambda_U [f_x, f_0 \rightarrow r] \\ &\quad - \sum_x \sum_y \mathcal{P}_0(X = x, Y = y) \\ &\quad \left[ \sum_{r \in \sigma(x)} \sigma_x(r) \pi_V(g_{xy}, g_0 \rightarrow t_{ry}) \right] + \text{Eff}_V\end{aligned}$$

$$\begin{aligned}&= \sum_{x \in \mathcal{X}} \sum_{r \in \mathcal{R}} \sigma_x(r) \lambda_U \left[ f_x \log \frac{r}{f_0} \right. \\ &\quad \left. + (1 - f_x) \log \frac{1 - r}{1 - f_0} \right] \\ &\quad - \sum_x \sum_y \mathcal{P}_0(X = x, Y = y) \left[ \sum_{r \in \sigma(x)} \sigma_x(r) \lambda_V \right. \\ &\quad \left. \left[ g_{xy} \log \frac{t_{ry}}{g_0} + (1 - g_{xy}) \log \frac{1 - t_{ry}}{1 - g_0} \right] \right] + \text{Eff}_V \\ &= \sum_{x \in \mathcal{X}} \sum_{r \in \mathcal{R}} \sigma_x(r) \lambda_U \left[ f_x \log \frac{r}{f_0} \right. \\ &\quad \left. + (1 - f_x) \log \frac{1 - r}{1 - f_0} \right] \\ &\quad - \sum_x \sum_y \mathcal{P}_0(X = x, Y = y) \left[ \sum_{r \in \sigma(x)} \sigma_x(r) \lambda_V \right. \\ &\quad \left[ g_{xy} \log \frac{\sum_x \mu_r(x) g_{xy}}{g_0} \right. \\ &\quad \left. \left. + (1 - g_{xy}) \log \frac{1 - \sum_x \mu_r(x) g_{xy}}{1 - g_0} \right] \right] + \text{Eff}_V\end{aligned}$$

Note that in terms of the variables,  $\text{Eff}_V$ , is a constant in  $\pi_A(\sigma, \mu)$ . More over, for a fixed  $\mu$  value,  $\pi_A(\sigma, \mu)$  is linear in  $\sigma$ . This means, for a fixed  $\mu$ ,  $\pi_A(\sigma, \mu)$  is concave  $\forall \sigma \in \Sigma$ . It remains to be shown that for a fixed  $\sigma$ ,  $\pi_A(\sigma, \mu)$  is convex  $\forall \mu \in M$ . We recall the following [3]: The sum of convex/concave functions is convex/concave, the composition functions  $f(x) = h(l(x))$  is concave if  $h(x)$  is concave and non-decreasing and  $l(x)$  is concave.

We note that for a fixed  $\sigma$ :

$$\sum_{x \in \mathcal{X}} \sum_{r \in \mathcal{R}} \sigma_x(r) \lambda_U \left[ f_x \log \frac{r}{f_0} + (1 - f_x) \log \frac{1 - r}{1 - f_0} \right]$$

is constant in  $\mu$ , therefore if we show:

$$- \sum_x \sum_y \mathcal{P}_0(X = x, Y = y) \left[ \sum_{r \in \sigma(x)} \sigma_x(r) \lambda_V \left[ g_{xy} \log \frac{\sum_x \mu_r(x) g_{xy}}{g_0} + (1 - g_{xy}) \log \frac{1 - \sum_x \mu_r(x) g_{xy}}{1 - g_0} \right] \right]$$

is convex in  $\mu$ , then we have shown the desired result. Equivalently, we will show

$$\sum_x \sum_y \mathcal{P}_0(X = x, Y = y) \left[ \sum_{r \in \sigma(x)} \sigma_x(r) \lambda_V \left[ g_{xy} \log \frac{\sum_x \mu_r(x) g_{xy}}{g_0} + (1 - g_{xy}) \log \frac{1 - \sum_x \mu_r(x) g_{xy}}{1 - g_0} \right] \right]$$

is concave. As sum of concave functions is concave, we will only show

$$\mathcal{P}_0(X = x, Y = y) \left[ \sum_{r \in \sigma(x)} \sigma_x(r) \lambda_V \left[ g_{xy} \log \frac{\sum_x \mu_r(x) g_{xy}}{g_0} + (1 - g_{xy}) \log \frac{1 - \sum_x \mu_r(x) g_{xy}}{1 - g_0} \right] \right]$$

is concave. Note that  $l(\mu) = \frac{\sum_x \mu_r(x) g_{xy}}{g_0}$  is linear in  $\mu$ , and  $h(x) = \log x$  is concave and non-decreasing. This means  $h(l(\mu)) = \log \frac{\sum_x \mu_r(x) g_{xy}}{g_0}$  is concave in  $\mu$ , similarly we can show  $\log \frac{1 - \sum_x \mu_r(x) g_{xy}}{1 - g_0}$  is concave in  $\mu$ . Recalling that the

sum of concave functions is concave, we note that

$$\sum_x \sum_y \mathcal{P}_0(X=x, Y=y) \left[ \sum_{r \in \sigma(x)} \sigma_x(r) \lambda_V \left[ g_{xy} \log \frac{\sum_x \mu_r(x) g_{xy}}{g_0} + (1 - g_{xy}) \log \frac{1 - \sum_x \mu_r(x) g_{xy}}{1 - g_0} \right] \right]$$

is concave. This means that for a fixed  $\sigma$ ,  $\pi_A(\sigma, \mu)$  is convex  $\forall \mu \in M$ .

□

LEMMA 2. Let  $\mathcal{P}_{\sigma, r}(x)$  be the posterior probability of  $x$ . If  $\mu_r(x) \neq \mathcal{P}_{\sigma, r}(x)$  for  $r$  in the support of  $\sigma$  then  $\mu$  does not maximize  $\pi_B$  given  $\sigma$ .

PROOF. For a given  $\sigma$ ,  $\pi_A(\sigma, \mu) + \pi_B(\sigma, \mu) = \text{Eff}(\sigma)$  is fixed independent of  $\mu$ . Thus, the  $\mu'$  that minimizes  $\pi_A(\sigma, \mu')$  must also maximize  $\pi_B(\sigma, \mu')$ . By the myopic honest property of the market scoring rule, if  $\mu' \neq \mathcal{P}_{\sigma, r}(x)$  means Bob has an incentive to deviate by setting  $\mu' = \mathcal{P}_{\sigma, r}(x)$  and maximize his reward giving us the desired result. □

LEMMA 3. For any two strategy pairs  $(\sigma, \mu)$  and  $(\sigma', \mu')$  both satisfying Lemma 1, we must have the total reward from every report by Alice and Bob be the same. i.e.,  $\text{Eff}_U(\sigma) = \text{Eff}_U(\sigma')$  and  $\pi_B(\sigma, \mu) = \pi_B(\sigma', \mu')$ .

PROOF. Here we consider the case where there are two or more minimax-maximin strategy pairs. If there is only one, then clearly the statement of the lemma holds. Assume there are two minimax-maximin strategy pairs, those that satisfy Lemma 1, say  $(\sigma, \mu)$  and  $(\sigma', \mu')$ , where  $\text{Eff}_U(\sigma) \neq \text{Eff}_U(\sigma')$ . As both are minimax-maximin pairs it follows  $(\sigma, \mu')$  is also a minimax-maximin strategy pair. Assume not, then there are two cases, both leading to contradiction:

$\pi_A(\sigma', \mu') < \pi_A(\sigma, \mu')$  meaning  $(\sigma', \mu')$  is not a minimax strategy pair.

$\pi_A(\sigma, \mu) > \pi_A(\sigma, \mu')$  meaning  $(\sigma, \mu)$  is not a maximin strategy pair.

Recall that the cases above are exhaustive as  $\pi_A(\sigma, \mu) = \pi_A(\sigma', \mu')$ . As  $(\sigma', \mu')$  and  $(\sigma, \mu')$  are both minimax-maximin strategy pairs and, by Lemma 1,  $\pi_A(\sigma, \mu)$  is linear in  $\sigma$  for a fixed  $\mu$  it follows that  $(0.6\sigma + 0.4\sigma', \mu')$  and  $(0.4\sigma + 0.6\sigma', \mu')$  are both minimax-maximin strategies. Moreover, both  $(0.6\sigma + 0.4\sigma', \mu')$  and  $(0.4\sigma + 0.6\sigma', \mu')$  have the same support. The support is the same for both strategy pairs because each place non-negative weights on the same  $r$  values. Let  $r_\sigma$  be the support of  $\sigma$  and  $r_{\sigma'}$  be the support of  $\sigma'$ . The support of any convex combination of  $\sigma$  and  $\sigma'$ , will be over  $r_\sigma \cup r_{\sigma'}$ . As  $\text{Eff}_U(\sigma) \neq \text{Eff}_U(\sigma')$  it follows from the linearity of  $\text{Eff}_U(\cdot)$  that  $\text{Eff}_U(0.6\sigma + 0.4\sigma') \neq \text{Eff}_U(0.4\sigma + 0.6\sigma')$ . As  $\pi_A(0.6\sigma + 0.4\sigma', \mu') = \pi_A(0.4\sigma + 0.6\sigma', \mu')$ , it follows that  $\pi_B(0.6\sigma + 0.4\sigma', \mu') \neq \pi_B(0.4\sigma + 0.6\sigma', \mu')$ . Because both  $(0.6\sigma + 0.4\sigma', \mu')$  and  $(0.4\sigma + 0.6\sigma', \mu')$  have the same support, and the reward of Bob differ it follows that the posterior probability of  $x$ ,  $\mathcal{P}_{\sigma, r}(x)$ , under the two strategies differ. However,  $\mu_r(x) = \mathcal{P}_{\sigma, r}(x)$  for all minimax-maximin strategies. By Lemma 2 it follows that either  $(0.6\sigma + 0.4\sigma', \mu')$  or  $(0.4\sigma + 0.6\sigma', \mu')$  is not a minimax-maximin strategy, thus contradicting the assumption  $\text{Eff}_U(\sigma) \neq \text{Eff}_U(\sigma')$ . This means that all minimax-maximin strategies must have the same expected reward at every stage of play.

□

LEMMA 4. There is a maximin-minimax pair  $(\sigma, \mu)$  such that, for all  $r$  in the support of  $\sigma$ , and all  $x$ ,  $\mu_r(x) = \mathcal{P}_{\sigma, r}(x)$ . In other words, there is a maximin pair in which Bob's beliefs are consistent with Alice's strategy after every realized move  $r$ .

PROOF. Fix  $\sigma$ . Then,  $\pi_A(\sigma, \mu) + \pi_B(\sigma, \mu) = \text{Eff}(\sigma)$  is fixed independent of  $\mu$ . Thus, the  $\mu'$  that minimizes  $\pi_A(\sigma, \mu')$  must also maximize  $\pi_B(\sigma, \mu')$ . Based on the myopic honest property of the market scoring rule, this is guaranteed to be the case if and only if  $t_{ry}$  is accurate given the posterior probability of  $x$ ,  $\mathcal{P}_{\sigma, r}(x)$ . By picking  $\mu'_r(x) = \mathcal{P}_{\sigma, r}(x)$ , Bob can ensure this. If the original  $(\sigma, \mu)$  satisfied this, then it must be the case that, for all  $r$  with positive support, the change from  $\mu$  to  $\mu'$  did not change either  $s_{ry}$  and  $t_{ry}$ . For other  $r$ , there is no change from  $\mu$  to  $\mu'$ ; hence, for any  $\sigma'$ ,  $\pi_A(\sigma', \mu) = \pi_A(\sigma', \mu')$ . Thus,  $(\sigma, \mu')$  is minimax as well as maximin. □

THEOREM 5. Consider a minimax-maximin pair  $(\sigma, \mu)$  that satisfies belief-consistency (Lemma 4). Then, generically, it is a weak Perfect Bayesian Equilibrium for Alice to play  $\sigma$  in the first round, Bob to play  $\mu$  in the second round, and Alice to correct the price in the  $V$ -market in the third round.

PROOF. In the third round, conditioned on the play in the first two rounds, Alice makes the optimal correction; no one-stage deviation would be profitable. In the second round, as  $(\sigma, \mu)$  is a maximin strategy, Bob maximizes  $\pi_B$  (and minimizes  $\pi_A$ ) by following  $\mu$ , given that Alice is following  $\sigma$ . Finally, because it is a minimax strategy, no  $\sigma'$  would be a profitable deviation for Alice from  $\sigma$ . □

THEOREM 6. Let  $(\sigma^*, \mu^*)$  be a minimax-maximin pair. In any weak Perfect Bayesian Equilibrium  $(\sigma, \mu)$ , the payoff of Alice is exactly  $\pi_A(\sigma^*, \mu^*)$  and the payoff of Bob is exactly  $\pi_B(\sigma^*, \mu^*)$ .

PROOF. Consider Bob's trade in round 2 of the equilibrium strategy profile  $(\sigma, \mu)$ . By the equilibrium belief condition, Bob's optimal strategy is the myopic strategy conditioned on his beliefs. This is consistent with  $\sigma$  (by wPBE condition). Consider Alice's payoff. Given  $\mu$ ,  $\sigma$  must be the strategy that maximizes  $\pi_A(\sigma, \mu)$ . Thus, we have:  $\pi_A(\sigma, \mu) = \pi_A(\sigma^*, \mu^*)$ .

$$\forall \sigma' \neq \sigma, \quad \pi_A(\sigma', \mu) \leq \pi_A(\sigma, \mu). \quad (2)$$

As  $\mu$  is consistent with  $\sigma$ , it is the best response to  $\sigma$ . Thus:

$$\begin{aligned} \forall \mu' \neq \mu & \quad \pi_B(\sigma, \mu') \leq \pi_B(\sigma, \mu) \\ \text{Eff}_U(\sigma) - \pi_A(\sigma, \mu') & \leq \text{Eff}_U(\sigma) - \pi_A(\sigma, \mu) \\ \Rightarrow & \quad \pi_A(\sigma, \mu') \geq \pi_A(\sigma, \mu). \end{aligned} \quad (3)$$

Equations (2) and (3) imply that  $(\sigma, \mu)$  is a minimax-maximin pair. It directly follows that  $\pi_A(\sigma, \mu) = \pi_A(\sigma^*, \mu^*)$ . By Lemma 3 it follows that  $\pi_B(\sigma, \mu) = \pi_B(\sigma^*, \mu^*)$ . □

## 5. CHARACTERIZING INFORMATION REVELATION IN EQUILIBRIUM

In Section 4, we showed that the weak perfect Bayesian equilibria of the 3-stage game have payoffs consistent with the minimax value of Alice's payoff function  $\pi_A(\sigma, \mu)$ . In this section, we interpret this minimax value in information-theoretic terms, discuss approaches to computing the equilibrium strategies. We then define a related parameter, the

limiting ratio, of two given forecast variables  $U$  and  $V$ , that determines when Alice will be honest in the  $U$  market. We then characterize the limiting ratio of combinations of variables in terms of limiting ratios of the components.

## 5.1 Interpretation as a noisy channel

One important consequence of Theorem 6 is that the efficiency of the  $U$  market,  $\text{Eff}_U(\sigma)$ , has the same value at any strategy  $\sigma$  that is part of an equilibrium strategy profile  $(\sigma, \mu)$ . Thus, we can use the notation  $\text{Eff}_U(V, \lambda_u, \lambda_v)$  to denote the value of  $\text{Eff}_U$  in any equilibrium of the game.

Consider an equilibrium strategy pair  $(\sigma, \mu)$ , and assume that the space of possible moves  $\mathcal{R}$  is rich enough for  $\sigma$  to satisfy Lemma 4.  $\text{Eff}_U(\sigma)$  measures the increase in expected score in the  $U$  scoring rule that results from Alice's (perhaps randomized) move in the first stage. Due the use of the logarithmic scoring rule, this has a natural information-theoretic interpretation: It defines the expected reduction in entropy of the forecast variable  $U$  as a result of Alice's move.

Note that a strategy  $\sigma$  specifies, for each realized signal  $x$ , a distribution over a subset of  $\mathcal{R}$ . We can view this process as follows: The signal  $x$  serves as input to a noisy channel, which generates the randomized output signal  $r = \sigma(x)$ . We can now conceptually decompose Alice's original private signal  $x$  into two signals  $\sigma(x)$  and  $x/\sigma(x)$ , where  $\sigma(x) = r$  is the noisy output of the channel, and  $x/\sigma(x)$  is the additional information contained in the original  $x$  once  $r$  is known. The efficiency  $\text{Eff}_U(\sigma)$  captures the mutual information between  $U$  and the noisy signal  $\sigma(x)$ . If  $\sigma$  is the honest strategy, then  $\sigma(x) = x$ , and the efficiency measures the mutual information between  $U$  and  $x$ , *i.e.*, the reduction in entropy of  $U$  conditional on knowing  $x$ . If  $\sigma$  is not honest, *i.e.*, it does not reveal all information known to Alice about  $U$ , the mutual information between  $U$  and  $\sigma(x)$  will be strictly lower.

The decomposition of the signal  $x$  into two components provides another characterization of the equilibrium strategy  $\sigma$ , as follows. Suppose that the signal  $x$  was actually given to Alice in two parts  $\sigma(x), x/\sigma(x)$ . The first component would be sufficient for Alice to execute her strategy in stage 1, while the second component would be used in stage 3. Moreover, if Alice only saw the first component, it must still be optimal for her to follow the same strategy: if further randomization improved her payoff, then  $\sigma$  would not be a minimax strategy. Thus, if Alice received the signal  $\sigma(x)$  instead of  $x$ , she would honestly reveal it in the first round. Further, we note that, having revealed  $\sigma(x)$ , Alice cannot profit by choosing to reveal additional information in the first round. Thus, the equilibrium strategy  $\sigma$  has the property that  $\sigma(x)$  would be revealed honestly, but any refinement of  $\sigma(x)$  that conveys additional information about  $u$  would not be honestly revealed.

## 5.2 Computing Equilibrium Strategies

In this section, we illustrate how the problem of computing the equilibrium noise channel  $\sigma$  can be formulated as a pure optimization problem. The optimization problem is nonlinear, so this formulation serves as a starting point but not a complete solution to the problem of how to compute the equilibrium  $\sigma, \mu$  for a given setting. Understanding the computational properties of the equilibrium strategies is an important direction for future research.

We consider a special case in which each of the players only observes a binary signal, *i.e.*  $\mathcal{X} = \mathcal{Y} = \{1, 2\}$ . We

also assume that the support of Alice's strategy,  $\sigma$ , has only two points  $r_1$  and  $r_2$ . Further, we assume that these points are consistent with honest reporting, *i.e.*,  $r_1 = \mathcal{P}_0[u = 1|X = 1]$  and  $r_2 = \mathcal{P}_0[u = 1|X = 2]$ . This is equivalent to assuming that Alice has to report a value that is consistent with one of her potential signals; of course, she need not report this honestly.

Alice's strategy  $\sigma$  can then be represented as follows:

$$\begin{aligned}\sigma &\stackrel{def}{=} [\sigma(1), \sigma(2)] \\ \sigma(1) &\stackrel{def}{=} [\sigma_1(r_1), \sigma_1(r_2)] \\ \sigma(2) &\stackrel{def}{=} [\sigma_2(r_1), \sigma_2(r_2)]\end{aligned}$$

As  $\sum_r \sigma_x(r) = 1$ , we can rewrite  $\sigma$  as:

$$\sigma \stackrel{def}{=} [\sigma_1(r_1), \sigma_2(r_2)] \quad (4)$$

By Lemma 4, we know that  $\mu_r(x) = \mathcal{P}_{\sigma, r}(x)$ , meaning Bob's beliefs are consistent with Alice's strategy. Therefore, if Alice moves to  $r_1$ , Bob's belief that Alice received signal 1 is  $\mu(r_1) = \frac{\sigma_1(r_1)}{\sigma_1(r_1) + \sigma_2(r_1)}$ . With these beliefs, if Bob observed  $r_1$  in the  $U$ -market and he observed 1 as his signal, then his truthful report would be  $\frac{\sigma_1(r_1)}{\sigma_1(r_1) + \sigma_2(r_1)}g_{11} + \frac{\sigma_2(r_1)}{\sigma_1(r_1) + \sigma_2(r_1)}g_{21}$ . Recall in equilibrium  $\pi_A(\sigma, \mu)$  is given by:

$$\pi_A(\sigma, \mu) \stackrel{def}{=} \text{Eff}_U(\sigma) + \text{Eff}_V - \pi_B(\sigma, \mu)$$

$\text{Eff}_V$  is constant, therefore in equilibrium Alice will find  $\sigma$  such that:

$$\sigma \in \underset{\sigma'}{\text{argmax}} \text{Eff}_U(\sigma') - \pi_B(\sigma', \mu) \quad (5)$$

Note that as described above, we may rewrite  $\mu$  as a function of  $\sigma'$ . We can rewrite the optimization problem above for the binary signal setting, with notation  $\bar{x} = 3 - x$ :

$$\begin{aligned}\underset{\sigma_1(r_1), \sigma_2(r_1)}{\text{argmax}} \quad & \mathcal{P}_0[X = 1] \left[ \sum_{i=1}^2 \sigma_i(r_i) \pi_U(f_i, f_0 \rightarrow r_i) \right] \\ & + \mathcal{P}_0[X = 2] \left[ \sum_{i=1}^2 \sigma_i(r_i) \pi_U(f_i, f_0 \rightarrow r_i) \right] \\ & - \sum_{x=1}^2 \sum_{y=1}^2 \mathcal{P}_0(X = x, Y = y) \\ & \left[ \sum_{i=1}^2 \sigma_x(r_i) \pi_V(g_{xy}, g_0 \rightarrow \frac{\sigma_x(r_i)}{\sigma_x(r_i) + \sigma_{\bar{x}}(r_i)} g_{xy} \right. \\ & \left. + \frac{\sigma_{\bar{x}}(r_i)}{\sigma_x(r_i) + \sigma_{\bar{x}}(r_i)} g_{\bar{x}y} \right]\end{aligned}$$

In this way, using the results from Section 4, we formulate the problem above that corresponds to the two signal, two values in the support of  $\sigma$  setting. The values of the decision variables that correspond to the optimal solution dictate Alice's and Bob's equilibrium strategies along the equilibrium path, although it does not directly yield the values for Bob's beliefs off the equilibrium path. The complexity of computing this optimal solution is not addressed in this paper.

## 5.3 Characterizing the Limiting Ratio

Although it is not easy in general to compute the equilibrium strategies of a given game setting, we can define a

related concept, the limiting ratio, that is computable to any degree of precision.

For given  $U, V, \lambda_U, \lambda_V$ , consider the following question: *Is there an equilibrium in which Alice is completely honest?* If there is an equilibrium  $(\sigma, \mu)$  in which this is true, then it must be that the honest strategy is optimal for Alice, even knowing that Bob is updating beliefs assuming that Alice is being honest. A necessary condition for this to happen is that, for any signal  $x$  that Alice has, she cannot profit by deviating to any other point in the support of the honest strategy, *i.e.*, she cannot profit by pretending to have any other signal  $x' \neq x$ . This condition captures that a bluffing strategy is not a profitable deviation even if Bob is “trusting”. We conjecture that this may also be a sufficient condition for honesty to be an equilibrium strategy. However, we have not been able to formally show this (in principle, Alice may still have profitable deviations to other points that are not consistent with any of her signals), and hence it only gives us a partial characterization of honest equilibrium. This necessary condition can be efficiently checked for a finite pool of signals.

We can then capture a measure of how aligned two forecast variables  $U$  and  $V$  are, given an information distribution  $P_0$ , as follows:

**DEFINITION 4.** *The Limiting Ratio  $LR(U, V)$  is defined as the maximum ratio  $\lambda_V/\lambda_U$  such that, in the conflict game  $(U, V, \lambda_U, \lambda_V)$ , if Bob plays a strategy that assumes Alice is honest whenever Alice moves to a point consistent with one of her signals, Alice would not profit by pretending to have a signal  $x'$  that is different from her true signals.*

(If no maximum exists,  $LR(U, V)$  is defined to be  $+\infty$ ).

This quantity is well-defined because of the following observations. The strategic behavior is clearly unchanged if both  $\lambda_V$  and  $\lambda_U$  are scaled by the same constant. As  $\lambda_V \rightarrow 0$ , the payoffs in the  $V$  market become negligible relative to Alice’s loss from bluffing in the  $U$  market, and thus  $LR(U, V) > 0$ . Further, if  $\lambda_V$  is reduced while  $\lambda_U$  is held fixed, Alice’s payoff from the honest strategy is unchanged while her payoffs from other strategies (for a given strategy  $\mu$  of Bob) can only reduce, and the relative value of different strategies for Bob (given Alice’s honest strategy) is unchanged. This means that if the condition is satisfied for any given  $\lambda_V$ , it must also hold for all smaller  $\lambda_V$  (for the same  $\lambda_U$ ). If an upper bound on  $\lambda_V/\lambda_U$  of interest is known,  $LR(U, V)$  can either be shown to exceed the upper bound or it can be computed to any desired precision by binary search, by checking at each value of  $\lambda_V/\lambda_U$  to test if the condition is satisfied.

As a measure of “alignedness” of incentives, the limiting ratio satisfies some structural relationships that may be useful in indirect analysis of strategic behavior. We now consider a more general form of forecast variable: Instead of restricting  $U$  and  $V$  to be Boolean functions, we can consider all  $U, V$  with discrete range contained in  $[0, 1]$ . The scoring function is assumed to be linear in the realized value, and to be maximized in expectation at the true mean value. Such  $U$  and  $V$  could be constructed by linear extension of a proper scoring rule. For example, if  $U$  takes the value 0.4, the score given would be  $0.4 \times$  score when  $U = 1 + 0.6 \times$  score when  $U = 0$ . The proper scoring rule is then an incentive scheme to infer the mean value of  $U$ .

With this extended setting, we can characterize the limit-

ing ratio of general forecast variables in terms of the simpler  $[0, 1]$  variables: We show that the limiting ratio  $LR(U, V)$  of combinations of forecast variables satisfies the following properties:

**THEOREM 7.** *Fix the signal spaces and the prior joint distribution of signals. Then, for any given forecast variables  $U_1, U_2, V_1, V_2$  with range in  $[0, 1]$  we have:*

1. If  $U = \frac{U_1+U_2}{2}$ , then  $LR(U, V) \geq \frac{LR(U_1, V)+LR(U_2, V)}{2}$
2. If  $V = \frac{V_1+V_2}{2}$ , then  $LR(U, V) \geq \text{Harm}(LR(U_1, V), LR(U_2, V))$ , where the harmonic mean function  $\text{Harm}(a, b) \stackrel{\text{def}}{=} \frac{2ab}{a+b}$ .

**PROOF.** 1. Let  $k = \frac{LR(U_1, V)+LR(U_2, V)}{2}$ . We need to show that, for  $\lambda_U = 1, \lambda_V = k$ , the honest strategy is an equilibrium, but this is not true for any higher value of  $\lambda_V$ . Testing if bluffing is profitable reduces to checking a finite set of possible deviations for Alice, given that Bob assumes Alice is being honest: Alice could pretend to have signal 0 when she had signal 1, etc. Without loss of generality, we assume that the most attractive deviation, at the critical value  $LR(U, V)$ , is for Alice to pretend to have 1 when she has 0. Alice would deviate to pretend to have signal 1 only if the expected penalty (call it  $\text{Pen}(U)$ ) she incurs for this inaccuracy in the  $U$  market is less than the expected gain  $k\text{Gain}(V)$  from moving the  $V$ -market from  $g_{0y}$  to  $g_{1y}$ .

Note that  $\text{Pen}(U) = (\text{Pen}(U_1) + \text{Pen}(U_2))/2$ , because of the linear extension of the scoring rules. Further,  $k\text{Gain}(V)$  is the same, whether we are looking at game  $(U, V, 1, k)$ ,  $(U_1, V, 1, k)$ , or  $(U_2, V, 1, k)$ . We know by definition that  $LR(U_1, V)\text{Gain}(V) \leq \text{Pen}(U_1)$ , because the penalty would not be outweighed by the gain; likewise,  $LR(U_2, V)\text{Gain}(V) \leq \text{Pen}(U_2)$ . Averaging these yields  $k\text{Gain}(V) \leq \text{Pen}(U)$ , and hence the result.

2. Again, we can focus our attention on a single potential deviation, and testing for honest equilibrium reduces to testing if  $\lambda_V\text{Gain}(V) \leq \lambda_U\text{Pen}(U)$ . By definition, we have that  $\frac{1}{LR(U, V_1)}\text{Pen}(U) \geq \text{Gain}(V_1)$ , and  $\frac{1}{LR(U, V_2)}\text{Pen}(U) \geq \text{Gain}(V_2)$ . Taking the average, we have that  $\frac{1}{\text{Harm}(LR(U, V_1), LR(U, V_2))}\text{Pen}(U) \geq \frac{\text{Gain}(V_1)+\text{Gain}(V_2)}{2} = \text{Gain}(V)$ . This completes the result.

□

These results provides a building block to gain insight into situations in which there are multiple conflicting incentives.

## 6. APPLICATIONS AND EXTENSIONS

### 6.1 Extensions of the basic result

*Extension to other scoring rules.* The results we presented in Section 4 are based on the fact the expected reward Alice stands to earn in either market is concave in  $\sigma$  for a fixed  $\mu$  and convex in  $\mu$  for a fixed  $\sigma$ . Therefore, any scoring rule that satisfies these conditions will have the

properties outlined in Section 4 and Section 5. For example, the quadratic scoring rule has these properties and thus has the same properties.

**Extension to Non-Binary Outcomes.** As described above, the model we consider only considers binary market variables. The model can extend to any finite outcome settings, by extending the definitions of  $f_{xy}, f_x, f_y, g_{xy}, g_x, g_y, f_0,$  and  $g_0$ . We also need to expand the definition of  $\pi_U(\mathbf{p}, \mathbf{r} \rightarrow \mathbf{q})$  and  $\pi_V(\mathbf{p}, \mathbf{r} \rightarrow \mathbf{q})$  with respect to all possible outcomes. Here,  $\mathbf{p}$  is a vector of beliefs on the possible outcomes. For example, if there were 3 outcomes, then  $\pi_U(\mathbf{p}, \mathbf{r} \rightarrow \mathbf{q}) = \lambda_U \left[ p_1 \log \frac{q_1}{r_1} + p_2 \log \frac{q_2}{r_2} + (1 - (p_1 + p_2)) \log \frac{1 - (q_1 + q_2)}{1 - (r_1 + r_2)} \right]$ , where  $p_i, q_i,$  and  $r_i$  are the  $i$ th component of  $\mathbf{p}, \mathbf{q},$  and  $\mathbf{r}$  respectively. With these definitions, results similar to those in Section 4 can be obtained.

**Other Signaling Mechanisms.** The result presented in Section 4 stemmed from the fact that for a fixed  $\mu, \pi_A(\sigma, \mu)$  is concave  $\forall \sigma \in \Sigma$ . Further, we used the fact  $\text{Eff}_U(\sigma)$  is linear in  $\sigma$  to show the payoff in the U-market is the same for all strategies. We can think of  $\text{Eff}_U(\sigma)$  as capturing the ‘‘cost’’ of Alice misleading Bob. From our results, we note that any concave, in  $\sigma$ , signaling mechanism could be used in place of the U-market and the results we displayed will still hold. Moreover, if this signaling mechanism was strictly concave, we note Lemma 3 would be simplified by noting there is a unique  $\sigma$  for any  $\mu$ .

**Need for Distinguishability Assumption.** The genericity assumption described by Definition 1 allows us to only consider non-degenerate cases of our problem. If the assumption does not hold, there would exist belief distributions  $\mu$  for Bob for which his move in the V-market does not reveal his signal to Alice. If the minimax equilibrium strategies for the specified  $U, V$  happened to result in a  $\mu$  value for which this was true, Alice could not actually attain the payoff  $\text{Eff}_U(\sigma) + \text{Eff}_V - \pi_B(\mu, \sigma)$ , because she could not always move the V-market to the optimal point. Then, the minimax equilibrium would not correspond to a wPBE of the 3-stage game. Thus, we cannot abandon the V-distinguishability assumption in our analysis. Note, however, that an arbitrarily small random perturbation of all the  $g_{xy}$  will lead to V-distinguishability.

## 6.2 Conflicting incentives in two markets

The most direct application of our model is in situations where there are multiple markets with an overlapping trading community. For example, information that Bob gleans from a market to forecast interest rates might be very relevant to his beliefs about the price of oil futures. Before trading in the oil futures market based on this information, Bob should take into account the possibility that he is being misled about the common wisdom on interest rates in order to exploit him in the oil futures market.

Our results suggest that, in situations like this, there is a minimax strategy  $\mu^*$  that Bob can employ to update his beliefs in a non-exploitable way. Moreover, Bob can even publicly announce the strategy  $\mu^*$ , while retaining guarantees on the informativeness of information provided by rational risk-neutral traders in the interest-rate market: His safety is not based on obscurity of his belief-updating process. Fur-

ther, if Bob is actually operating the interest-rate market, he can even set the stakes of this market to optimize some function of the amount of trustworthy information gathered and the subsidy that must be provided to traders in the interest rate markets. (As a caveat, we reiterate that the difficulty of computing the minimax equilibrium  $(\sigma^*, \mu^*)$  has not been resolved, and hence we do not know which class of market models can be feasibly solved in this way).

## 6.3 Non-myopic strategies in a single market

Our model also sheds light on multi-round strategies within a single market scoring rule market. Consider a special case of our model, in which the forecast variable  $V$  is *exactly the same* as the forecast variable  $U$ . In this case, the three stages of our model can be interpreted as three rounds of trade in a market scoring rule market: Alice trades in the first round, but she may be trying to mislead Bob who trades next, in order to profit from Bob’s errors by trading in the third round. The one difference between our model and a true model of a three-stage market scoring rule is that, in the market scoring rule, Bob’s payoff in the second round would be measured relative to the price to which Alice moved the market, rather than the prior expected value of  $V$ . However, the equilibrium characterization is invariant under this change: As Bob trades only once, his optimal strategy is always to maximize the expected score of the price he moves the market to, regardless of where it was when he started. Thus, the wPBE of our model are also wPBE of the single 3-stage market.

In this setting (among others), Chen *et al.* [5] have shown that the existence of honest equilibria depends on whether the information signals are substitutes or complements. The results of our paper provide additional insight into this problem: We now know that, even if there are no honest equilibria, the payoffs in equilibrium are unique, and correspond to the minimax payoffs of our model.

## 6.4 Non-market incentive conflicts

Although our model is centered on incentive conflicts that arise due to future *market* actions by Bob, it can be used to derive preliminary insight into situations where the incentive conflict arises due to non-market conflicts as well. In particular, the second ‘‘market’’  $V$  can be a virtual proxy for any source of potential gain for Alice. Of course, general conflicts of interest are unlikely to take the elegant form of proper scoring rules, and so our results do not apply directly. As long as they can be approximated by a scoring function that is optimal at the true probability, and generates a payoff for Alice that is convex in  $\mu$ , our analysis technique may still provide useful insights. We now present an example of such an analysis.

Consider the following situation: A forecaster Alice has to predict the probability that the company sales will go up next month. Alice’s information from this month could lead her to one of two conclusions: either the sales will go up 70% next month (in which case we say signal is ‘1’) or sales will go up 30% next month (in which case we say that her signal is ‘0’). The manager, Bob, needs a reliable forecast of the sales growth, so he offers Alice a proper scoring rule based on her forecast. However, Alice has an outside conflict of interest: If the official forecast announced by Bob is  $x\%$ , but Alice got the low signal, *i.e.*, the indicators are in fact low, Alice will get a bonus of  $10(x - 30)$ . Bob has to decide how to interpret any forecast given by Alice, and Alice has

to decide what (perhaps random) forecast to make when she gets signal '1', and when she gets signal '0'.

At the surface, this is a substantially different setting from the two-market models we analyze. The particular differences are: (i) Alice's conflicting payoff does not take the form of a proper scoring rule; (ii) Alice profits in the second stage if Bob overestimates the true forecast probability when the forecast should be low, but not if Bob underestimates the true forecast probability when the forecast should be high; (iii) Bob's expected gain or loss in the second stage need not be precisely the negative of Alice's gain or loss in this stage. We now illustrate how these variants can be accommodated within the class of models that we analyze in this paper.

We first tackle the issue of the scoring function, restricting our attention to the case when Alice has received a '0' signal. In this case, supposing Alice forecasts  $r$ , let  $\mu_r$  be Bob's belief on the conditional probability that Alice has a '1'. Acting according to this belief, Bob will forecast  $30 + 40\mu_r$ , and so Alice's bonus will be  $400\mu_r$ .

Now, consider a "scoring function"  $S(\text{true } p, \text{forecast } q) = -1000|q - p|$  that determines the payoff to a forecaster from saying  $q$  when the true probability is  $p$ . Unlike the expected score of a proper scoring rule,  $S()$  is not implementable in expectation – if we did not know the true probability, we could not implement the payoffs it dictates. However, for the purposes of *analysis* alone, this does not matter. We can analyze Alice's and Bob's strategy as if they were paid off according to a market scoring rule based on scoring function  $S()$ , as long as we restrict our analysis to expected profit. Moreover, we note that  $S()$  satisfies the two properties required for the proof of Theorem 5 and Theorem 6: Alice's profit from correcting the market price when she has signal '0' is  $400\mu_r$ , which is linear (and hence convex) in  $\mu$ , and the value of  $q$  that maximizes  $S()$  is exactly  $q = p$ , so Bob's optimal strategy is the honest strategy given his beliefs. Thus, the results on the existence and uniqueness of equilibria still hold, as does the analysis of the limiting ratio for honesty.

Next, we turn to the second problem: In the market model, Alice earns a profit when she has signal '1' but Bob believes that she has '1' with probability  $\mu_r < 1$ . However, in the real situation we want to model, Alice receives no benefit in this case. Again, we can handle this problem by taking advantage of the fact that the scoring function in our model need not be implementable. In particular, consider the scoring function  $S(\text{true } p, \text{forecast } q) = -1000(q - p)$  when  $q \geq p$ , and  $-\epsilon(p - q)$  when  $p > q$ , for some small  $\epsilon > 0$ . Again, this function is strictly maximized at  $q = p$ . Further, we note that  $q \geq p$  whenever Alice has signal '0', and  $p \leq q$  whenever Alice has signal '1'. For a given strategy  $\sigma$ , and a given  $r$  in the support of  $\sigma$ , there are fixed posterior probabilities  $\sigma_0(r)$  and  $\sigma_1(r)$ , and thus Alice's expected payoff is proportional to  $\sigma_0(r)400\mu_r + \sigma_1(r)0.4\epsilon(1 - \mu_r)$ , which is linear in  $\mu$ . Note that for all  $\sigma, \mu$ , this is within  $\epsilon$  of Alice's payoff in the original scenario, and hence an equilibrium of this market model corresponds to an approximate equilibrium of the original situation.

Finally, we consider the fact that Bob's profit or loss from the forecast need not be the exact inverse of Alice's profit or loss. Note that this does not affect the minimax analysis, which is based on Alice's profit, but it could affect our analysis of weak-perfect Bayesian equilibrium of the original game. It may still be a reasonable assumption that

Bob's profit from a forecast is maximized when the forecast matches the true posterior probability, i.e., Bob is never worse off by calibrating his forecasts to Alice's actual mixing strategy. Under this assumption, a minimax equilibrium of our analytical model would still correspond to a weak Perfect Bayesian Equilibrium of the original situation.

This example only illustrates some ways in which problem of strategic forecasts with conflicting incentives could be analyzed within our model. In this paper, we have no results to suggest that such a reduction can be carried out for any general class of external conflict of interests. Characterizing conflicts that are amenable to this reduction is an important direction for future work.

## 7. CONCLUSION

In this paper, we initiate the study of composition of scoring-rule based incentive mechanisms. We formulated and analyzed a simple model that captures the conflict of interest at the core of such situations. We showed that the information revealed and participant payoffs in equilibrium are uniquely determined, and consistent with a minimax strategy profile.

Our results suggest several important directions for future work. First, it will be interesting to develop a computational theory that addresses how hard it is to compute the optimal strategy profiles for this class of games. Second, developing a theory of optimal or bounded-regret strategies when the exact nature of the conflict of interest is unknown but bounded would be very useful. Third, we consider only a three round game composed of two markets. It would be insightful to extend the analysis to a multi-round setting in which Bob need not honestly reveal his information in his first move.

## Acknowledgments

This material is based on work supported by the National Science Foundation under Grant No. CCF-0728768.

## 8. REFERENCES

- [1] Michael B. Abramowicz and M. Todd Henderson. Prediction markets for corporate governance. *Social Science Research Network Working Paper Series*, September 2006.
- [2] Shipra Agrawal, Erick Delage, Mark Peters, Zizhuo Wang, and Yinyu Ye. A unified framework for dynamic pari-mutuel information market design. In *EC '09: Proceedings of the tenth ACM Conference on Electronic Commerce*, pages 255–264, 2009.
- [3] Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004.
- [4] Glenn W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, January 1950.
- [5] Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel Reeves, David M. Pennock, Robin Hanson, Lance Fortnow, and Rica Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, May 2009.
- [6] Yiling Chen, Daniel Reeves, David M. Pennock, Robin Hanson, Lance Fortnow, and Rica Gonen. Bluffing and strategic reticence in prediction markets. In

- WINE'07, the third Workshop on Internet and Network Economics*, pages 70–81, San Diego, CA, December 2007. Springer LNCS.
- [7] Steven Cherry. Bet On It! *Spectrum, IEEE*, 44(9):48–53, 2007.
- [8] Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *EC '08: Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 200–209, 2008.
- [9] R. Hanson, R. Oprea, and D. Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, 60(4):449–459, August 2006.
- [10] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, January 2003.
- [11] Robin Hanson. The policy analysis market (a thwarted experiment in the use of prediction markets for public policy). *Innovations: Technology, Governance, Globalization*, 2(3):73–88, August 2007.
- [12] Nicolas S. Lambert, David M. Pennock, and Yoav Shoham. Eliciting properties of probability distributions. In *EC '08: Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 129–138, 2008.
- [13] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, USA, June 1995.
- [14] Peng Shi, Vincent Conitzer, and Mingyu Guo. Prediction mechanisms that do not incentivize undesirable actions. In *WINE'09, the fifth Workshop on Internet and Network Economics*, pages 89–100, December 2009.
- [15] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.
- [16] Robert L. Winkler. The quantification of judgment: Some methodological suggestions. *Journal of the American Statistical Association*, 62(320):1105–1120, 1967.
- [17] Robert L. Winkler. Scoring rules and the evaluation of probability assessors. *Journal of the American Statistical Association*, 64(327):1073–1078, 1969.
- [18] Justin Wolfers and Eric Zitzewitz. Prediction markets. *The Journal of Economic Perspectives*, 18(2):107–126, 2004.