Interpreted and Generated Signals

Lu Hong[*] and Scott Page[†]

May 17, 2006

Abstract

Models of incomplete information use signals to capture the uncertain values of variables. The relationship between these signals and the relevant variables can then be described by a joint probability distribution function. In this paper, we contrast *generated signals* and *interpreted signals*. Generated signals are distortions of the true values or outputs of some process. Interpreted signals are predictions based on inputs of a process or attributes of an object. These two types of signals produce distinct statistical signatures. In particular, under rather mild assumptions, interpreted signals will be negatively correlated in both the accuracy, be it conditional or unconditional, whereas generated signals will be independent conditional on the state. Thus, our findings limit the contexts in which results pertaining to information aggregation in auctions, markets, and voting may apply.

[*]Department of Finance, Loyola University Chicago
[†]Departments of Political Science and Economics, University of Michigan

1 Introduction

When making economic or political decisions, agents rarely have the benefit of full information. They rely on partial or distorted information or on crude predictive models. Social scientists model this incomplete information in the form of signals. A joint probability distribution characterizes the statistical relationships between these signals and the variable of interest: the better the information, the tighter the link between the signals and the underlying variable of interest. In this paper, we show that assumptions about these joint probability distributions, such as independence or independence conditional on the state, constrain the set of environments that could produce the signals. The statistical properties of signals influence strategic choices and therefore outcomes. For this reason, analyses of economic and political institutions hinge on statistical assumptions. This paper helps to connect those assumptions, and therefore those results, to specific environments. In doing so, the paper lays a foundation for understanding the design and the performance of markets and political mechanisms.

To show how properties of the joint distribution function depend on the micro level foundations of the signals, we contrast two frameworks for producing signals. One of these frameworks assumes that signals are generated by a process. The other assumes that agents construct predictive models that produce signals. Both types of the signals, of course, can be modeled by joint probability distributions. However, the assumptions applicable to the two types differ. Binary signals generated by a process are typically independent. Binary signals based on predictive models cannot be. They must be negatively correlated. And moreover, the degree to which they are negatively correlated is uniquely determined by their accuracy.

In what follows, we refer to signals produced by some process as *generated signals*. Two agents' generated signals differ because the agents draw different samples or experience idiosyncratic shocks or distortions. Alternatively, agents can produce signals by filtering a high dimensional reality onto dimensions or into categories that they believe important. They can then construct predictive models based on those filterings (Fryer and Jackson 2003, Page 2007, Nisbett 2003).[1] We call these *interpreted signals*. Two agents' interpreted signals differ if the agents' predictive models differ in how they categorize or classify objects, events, or data.[2]

Given our constructions, generated signals have a probabilistic relationship to the relevant outcome variable. No number of generated signals, independent or non

[1] Organizations also differentially filter reality. See for example, Stinchecombe 1990.

[2] Our introduction of the term interpreted signals is new to the literature. Similarities exist between our interpreted signal framework and models of causal inference (Pearl 2000), fact free learning (Aragones, Gilboa, Postlewaite, and Schmeidler 2005), and complexity (Al-Najjar, Casadesus-Masanell and Ozdenoren 2003). Interpreted signals also resemble predictions from PAC learning (Valiant 1984) and classification theory (Barwise and Seligman 1997). The acronym PAC refers to Probably Approximately Correct. We differ from PAC learning and classification models in that we consider multiple agents each with his/her own classification rule.

independent, pins down the outcome value with certainty. Moreover, any number of generated signals can be produced satisfying the same distributional and correlational assumptions. Neither of these statements need be true of interpreted signals. If all relevant attribute values can be learned, then the value of an outcome may be predicted with certainty. If a poker player knows all of the cards in her opponents' hands, she knows whether she wins or loses. Given that interpreted signals depend on attributes, the number of distinct interpreted signals is bounded by the power set of the set of attributes. This latter implication implies that any model that assumes a large number of signals assumes generated signals and rules out interpreted signals, an insight we take up in more depth in the discussion at the end of this paper.

In this introductory paper, we focus on *independent interpretations* – interpretations that consider distinct attributes of a common representation, or perspective (Hong and Page 2001, Page 2007). This captures situations in which agents look at distinct but relevant pieces of information when making predictions. We also consider the special case where the outcome of interest takes on only binary values. Doubtless, the model generalizes, but the binary case provides the cleanest entry into the issues of concern.

Given independent interpretations, we show that some common distributional assumptions are not consistent with independent interpreted signals. Many models of information aggregation in common value auctions, voting, and in information cascades assume the following signaling structure that involves two states G (good) and B (bad) and two signals g and b: The two states are equally likely, the two signals predict the two states with equal likelihood, the signals are independently correct (i.e. knowing that one signal is correct tells nothing about the probability that the other is correct) and the signals are informative, i.e. they are correct more than half of the time. These assumptions are inconsistent with independent interpreted signals, which must negatively correlated conditional on at least one outcome. Moreover, these signals cannot be independently correct. They must be negatively correlated in their correctness both unconditionally and conditionally on at least one outcome.

The negative correlation result bodes well for information aggregation but has a contingent normative implication for auctions depending on whether we're aligned with the buyers or the seller. Were our result only that the correlation is negative, it would not be of much use. However, if as often assumed, good and bad predictions are equally likely, the probability of a correct interpreted signal uniquely determines the extent of this negative correlation. Thus, we have a benchmark correlation assumption for interpreted signals, but it is not zero.

In sum, the statistical properties of generated and interpreted signals differ substantially. These differences call into question the generality of many existing results. A close reading of the auction, voting, or signaling literature reveals that some models make assumptions consistent with generated signals, while others make assumptions consistent with interpreted signals. Particular assumptions often seems to be based on tractability and not on descriptive realism. In that part of the auction literature that relates to the derivation of optimal bidding strategies, the assumptions align with our interpreted signal framework (Klemperer 2004). However, in that part concerned with information aggregation in common value auctions, particularly those papers with large numbers of bidders, the assumptions are inconsistent with interpreted signals. Similarly, the bulk of the political science literature, including almost all jury models and election models (Ladha 1992, Feddersen and Pesendorfer 1997), makes assumptions consistent with generated signals.

The remainder of the paper is organized as follows. In Section 2, we provide examples of generated and interpreted signals and compare them. In Section 3, we introduce a framework for interpreted signals and discuss three notions of independence: *independent interpretations*, *independent signals*, and *independently correct signals*. In Section 3, we prove our main results, introduce the idea of signals based on overlapping interpretations and connect features of interpreted signals to the complexity of the outcome function. In Section 4, we compare interpreted signals to generated signals and demonstrate how to construct non independent interpreted signals that satisfy the conditional independence. In Section 5, we show how interpreted and generated signals provide alternative lenses through which to view the *affiliation* assumption (Milgrom and Weber 1982). The affiliation assumption requires stronger restrictions on generated signals than on interpreted signals. Moreover, if we allow for endogenous interpretations, i.e. strategically chosen predictive models, the affiliation becomes even weaker. In the conclusion, we discuss the implications of our results and future research questions.

2 Generated and Interpreted Signals: Examples

We begin with an example that highlights the statistical differences between generated and interpreted signals. Consider two venture capitalists who receive signals about the quality of an investment, say a restaurant. This restaurant can either be a good investment ($V = G$) or a bad one ($V = B$). Each outcome occurs with equal probability. We first describe the standard *generated signal* model.

2.1 Generated Signals

Generated signals can be thought of as noisy glimpses or distortions of an outcome value V. Imagine, for example, that two potential investors eat meals prepared at the restaurant. We can think of those meals as generated signals denoted by s_1 and s_2 that either take value g (for good) or b (for bad). Conditional on the restaurant being a good investment, i.e. conditional on state G, we assume that the probability of getting a good meal, i.e. of receiving the signal g, equals 2/3. Similarly, conditional on B, the probability of getting the signal b equals 2/3. Each of these signals is further assumed to be an independent draw from this distribution. Thus, we can write the joint probability distributions of signals conditional on the restaurant's value as follows:

Generated Signals Conditional on G

s_1/s_2	b	g
b	1/9	2/9
g	2/9	4/9

Generated Signals Conditional on B

s_1/s_2	b	g
b	4/9	2/9
g	2/9	1/9

Using this information, we can compute the probability that the restaurant is a good investment conditional on the signals of the two investors.

Probability of a Good Investment
Conditional on Generated Signals

s_1/s_2	b	g
b	1/5	1/2
g	1/2	4/5

The above table can be read as follows. If both investors get the signal b then the probability that the restaurant is a good investment equals $\frac{1}{5}$.

2.2 Interpreted Signals

We now turn to interpreted signals. Interpreted signals are predictions based on *interpretations*. Interpretations create partitions (or categorizations) of the set of possible restaurants. Interpretations partition the set of attributes that define a restaurant. In this example, we consider those attributes to be the restaurant's *location* and its *prices*. We assume that each investor sees only one of these attributes and bases her prediction on that attribute's value.

Interpreted signals require a *outcome function*, V, that maps the restaurant's *attributes*, into a probability that the restaurant is a good investment. Here, we assume that the location ℓ and the prices \$ can be either good 1 or bad 0, and that each combination of attribute pairs is equally likely. We assume the following functional form for the outcome function.[3]

[3] In the paper, we consider deterministic outcome functions. Any probabilistic outcome function can be transformed into a deterministic outcome function by adding attributes. In this example, we need only add a third independent attribute h that takes take three values, say 0, $\frac{1}{2}$, and 1, with equal probability. The outcome function can then be written as a deterministic function as follows:

$$V(\ell, \$) = \begin{cases} \frac{1}{3} & \text{if } (\ell + \$) \leq 1 \\ 1 & \text{otherwise} \end{cases}$$

In this example, we assume that the first investor looks only at the location and the second looks only at prices. The investors then construct *predictive models* based on their attributes' values. In brief, if the attribute's value equals 1 (resp. 0), the investor predicts the restaurant will be good (bad). The investors' *interpreted signals* equal the value of the investors' attributes. In the general framework, predictions can be based on multiple attributes and therefore are not identical to the attributes themselves. Given these interpreted signals, we can now write a joint probability distribution for the signals and the value of the outcome function.

Probability of a Good Investment
Conditional on Interpreted Signals

$\ell/\$$	0	1
0	1/3	1/3
1	1/3	1

We can then calculate the joint probability distribution of the interpreted signals conditional on the restaurant's quality. As with the generated signals, the probability of an agent getting the good (bad) signal conditional on the restaurant being of good (bad) quality equals $\frac{2}{3}$.

Interpreted Signals Conditional on G

$\ell/\$$	0	1
0	1/6	1/6
1	1/6	1/2

Interpreted Signals Conditional on B

$\ell/\$$	0	1
0	1/3	1/3
1	1/3	0

$$V(\ell, \$, h) = \begin{cases} 1 & \text{if } h = 1 \text{ or : } (\ell + \$) > 1 \\ 0 & \text{otherwise} \end{cases}$$

Comparing the generated and interpreted signals reveals several differences. First, the interpreted signals are not conditionally independent. Conditional on the restaurant being a bad investment, having a good location *reduces* the likelihood that the prices are also good. In fact, the probability of having good prices equals zero. Second, the generated signals are independently correct, but the interpreted signals are not. Again, conditional on the restaurant being a bad investment, if one interpreted signal is incorrect, the other *must* be correct. Third, with interpreted signals, the correlation between the signals and values depends on the outcome function. With generated signals, whatever correlation exists is just assumed. Finally, in the case of interpreted signals, we are limited to two attributes: location and prices. The only constraint on the number of generated signals is the chef's time.[4]

3 The Interpreted Signal Framework

We now describe the interpreted signal framework. To avoid confusion, we also distinguish among several types of independence. In constructing interpreted signals, we differentiate between between the set of objects or events and their outcome values. We define *the environment*, Ω, to be a finite collection of objects or events with a cardinality equal to N. Each of these events or objects has associated with it an outcome. We denote the set of outcomes by S and the deterministic mapping from events to outcomes as an outcome function $\widetilde{O} : \Omega \to S$. [5]

The problems we consider are equivalent to binary classification problems in which an agent has to place the N objects into $|S|$ bins representing possible outcome values and are related to the problem of selecting regressors (Aragones, et al 2005). Here, we restrict attention to cases in which the cardinality of S equals two. To create an interpreted signal an agent partitions the environment into non-overlapping sets. We denote the partition of agent i, Π^i, to be the sets $\{\pi^i_1, \pi^i_2, \ldots, \pi^i_{n_i}\}$, where n_i is the number of sets in agent i's partition. Π^i is agent i's representation of the environment and this representation is incomplete as long as not all sets in the partition are singletons. When an agent sees an object or event, she associates it with the set in her partition that contains this object. We call these partitions **interpretations**.

Let $P : \Omega \to [0,1]$ be the probability distribution over Ω where $P(\omega)$ denotes the probability that event ω arises. Given this distribution over events and an interpretation of the environment, an agent makes predictions about the outcome. For example, an agent might use a Bayesian approach to making these predictions and

[4]This example creates a clean distinction between between generated and interpreted signals, whereas often the differences can be subtle. Consider a university whose quality q in $\{0,1\}$ depends on its faculty's abilities: each faculty member i has an ability a_i in the set $\{0,1\}$ drawn from some distribution F. Assume that this university produces students at regular intervals and that each graduate j receives an added value x_j in $[0,1]$. The graduating students' added values would be generated signals and the faculty's abilities would be interpreted signals. Thus, their statistical properties would differ in the ways that we describe.

[5]As previously mentioned, the framework extends to include probabilistic mappings.

assume that the most probable outcome arises conditional on the set in her interpretation. We refer to these as **experience generated predictions**. At this point, though, we do specify how predictions are generated. Agent i's **prediction** $\widetilde{\phi}_i$ is simply defined as a function from Ω to S with the restriction that $\widetilde{\phi}_i$ is measurable with respect to agent i's interpretation Π^i. We refer to these predictions as **interpreted signals**.[6] The following example illustrates the main components of the interpreted signal framework:

Example 1 *The environment,* $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$. *All events are equally likely,* $P(\omega_i) = \frac{1}{6}$. *The set of outcomes,* $S = \{G, B\}$. *The outcome function maps the first three events to* G *and the rest to* B. *Let* $\Pi = \{\pi_1, \pi_2\}$ *be an agent's interpretation of the environment where* $\pi_1 = \{\omega_1, \omega_2, \omega_4\}$ *and* $\pi_2 = \{\omega_3, \omega_5, \omega_6\}$. *If this agent makes experience generated predictions, then her predictions can be described by the following function*

$$\widetilde{\phi}(\omega_i) = \begin{cases} G & \text{for } \omega_i \in \pi_1 \\ B & \text{for } \omega_i \in \pi_2 \end{cases}$$

This example corresponds to the following specification in the standard signal model. There are two states $\{G, B\}$, *each state is equally likely. The agent gets a noisy signal* x. *We denote the signals with lower case letters* g *and* b. *The probability distribution of these signals conditional on state is given by:*

Conditional Probability Distribution of x

State	*Signal*	
of World	g	b
G	2/3	1/3
B	1/3	2/3

In the example and generally, the outcome function $\widetilde{O} : \Omega \to S$ together with a prediction $\widetilde{\phi} : \Omega \to S$ induce another random variable, $\widetilde{\delta}$, called the **correctness of predictions**, which can defined as follows: $\widetilde{\delta} : \Omega \to \{c, i\}$ such that

$$\widetilde{\delta}(\omega) = \begin{cases} c & \text{for } \omega \in \left\{\omega' \in \Omega : \widetilde{\phi}(\omega') = \widetilde{O}(\omega')\right\} \\ i & \text{for } \omega \in \left\{\omega' \in \Omega : \widetilde{\phi}(\omega') \neq \widetilde{O}(\omega')\right\} \end{cases}$$

where c means "correct" and i means "incorrect". Intuitively, the **accuracy** of a prediction $\widetilde{\phi}$ can be defined as the probability of its associated $\widetilde{\delta}$ taking c as its value. in the above example, $\widetilde{\delta}(\omega) = c$ for $\omega \in \{\omega_1, \omega_2, \omega_5, \omega_6\}$ and $\widetilde{\delta}(\omega) = i$ for $\omega \in \{\omega_3, \omega_4\}$. Thus, with probability $\frac{2}{3}$, $\widetilde{\phi}$ makes correct predictions.

[6]Unless in direct contrast with generated signals, we call them predictions as opposed to interpreted signals throughout the paper.

3.1 Definitions and Relationships of Types of Independence

We now turn to our discussion of various types of independence that relate to interpreted signals. Seveal of these definitions and results we describe are standard. We include them to place the new definitions and results in their proper context. We consider only pairwise independence. Extensions to include any finite number of interpretations are trivial.

definition 1 *Interpretations,* Π^1 *and* Π^2*, are* **independent interpretations** *if*

$$Prob\left(\pi_i^1 \cap \pi_j^2\right) = Prob\left(\pi_i^1\right) \times Prob\left(\pi_j^2\right)$$

for all $i \in \{1, ..., n_1\}$ *and* $j \in \{1, ..., n_2\}$

Two interpretations are independent implies that knowing how one agent interprets an event provides no information about how the other agent interprets that same event. Note that interpretations are not predictions, they are the sets within which agents place singular events. We now state a surprising result: *independent interpretations imply that the environment can be written as a product space and the interpretations written as projections onto variables.*

To make the logic that drives this result as transparent as possible, we first assume that each event in Ω is equally likely. Consider the following trivial observation: if Ω can be represented as a product of attribute spaces and if agents' make partitions by looking at subsets of attributes, then the agents have independent interpretations. For example, if the environment is written as a two by two lattice and one agent considers the row and the other considers the column, then these interpretations are independent. This is not surprising. Knowledge of an event's row, tells us nothing about its column.

We show that the converse also holds. If two interpretations are independent, then the event space can be mapped into a coordinate system (a two attribute model) where each event is represented by (x, y), and one interpretation considers only the x attribute and the other is along the y attribute. This result implies that any two independent interpretations can be rewritten as projections onto different attributes of *the same perspective* (Hong and Page 2001, Page 2007).[7] We provide the formal statement of this result below. It's proof along with all proofs of many subsequent claims is contained in the appendix.

[7]Recall that a perspective is an representation of the entire space of possibilities. Two agents use different perspectives if they represent the set of the possible alternatives with different languages. These languages can be basis. For example, one agent may identify a point in the plane using Cartesian coordinates (x, y). Another agent may use polar coordinates (r, θ). The natural interpretations differ for these two perspectives. In the former, someone might partition the space into points in which $x \leq 5$ and points in which $x > 5$. In the latter, an agent, might partition the space into points such that $r \leq 10$ and points in which $r > 10$.

Claim 1 *Assume that each event in* Ω *is equally likely. Let* $\Pi^1, \cdots, \Pi^n$ $(n \geq 2)$ *be non-trivial interpretations of* Ω*. If they are independent, then* Ω *can be represented by an* n*-attribute rectangle such that* Π^i *is along the* i*th attribute. Thus* $N = \prod_{h=1}^{n} a_h$ *for some larger-than-1 integers* a_h, $h = 1, ..., n$[8]

Intuitively, Claim 1 implies a bound on the number of independent interpretations. It cannot exceed the number of primes in the factorization of N. As stated in the Corollary below, a finite set of events admits few independent interpretations.

Corollary 1 *Assume events are equally likely. Let* $\prod_{i=1}^{k} p_i$ *be the unique prime factorization of* N, *that is,*

$$N = \prod_{i=1}^{k} p_i$$

where each p_i *is a prime. Then, the maximum number of independent non-trivial interpretations is* k.

The implications of this corollary sink in when applied to a specific example such as the set of possible independent interpretations of all of the 300 million people who in the United States. Such interpretations, the parsing of people into categories like soccer moms or NASCAR, are used to construct predictive models for economic, political, and social outcomes. The corollary implies that there exist fewer than *thirty* independent interpretations for the entire US population.[9]

Independent interpretations are distinct from independent predictions. Saying two agents's predictions are independent means that knowing one agent's prediction about the outcome of an event provides no information about the other agent's prediction.

definition 2 *Predictions,* $\widetilde{\phi}_1$ *and* $\widetilde{\phi}_2$*, are* **independent predictions** *if they are independent random variables.*

Note that if two agents have independent interpretations, then their predictions are independent as well.

[8]The result above is established with the assumption that all events are equally likely. This assumption is not essential. We can show that if events in the original space Ω do not have equal probability, there exists an equally probable event space Ω' that has greater cardinality (the least common denominator of probabilities of original events expressed in fractions) such that Ω' can be represented by an n-attribute rectangle and the independent interpretations of the original event space Ω correspond to interpretations of the new event space Ω' along different attributes. The key is that for independent interpretations, probabilities have the rectangle property, i.e.,

$$Prob(\pi_i^1 \cap \pi_j^2) = Prob(\pi_i^1) \times Prob(\pi_j^2)$$

[9]This result assumes that each attribute is binary such as { male, female }. To be precise, $2^2 8$ is slightly less than 300 million and $2^2 9$ exceeds it by a substantial margin.

Observation 1 *Independent interpretations imply independent predictions.*

Proof: Since each agent's prediction is measurable w.r.t. her interpretation, the claim follows.

In contrast, predictions can be independent without agents having independent interpretations. This result should come as no big surprise, but it drives a conceptual wedge between the two types of independence.

Observation 2 *Independent predictions may not imply independent interpretations.*

Proof: Due to the measurability requirement, an agent's prediction defines a partition on the event space that is in general coarser than her interpretation. Recall Example 1. We can add a second agent whose interpretation is

$$\Pi^2 = \{\{\omega_1, \omega_2\}, \{\omega_4, \omega_5\}, \{\omega_3, \omega_6\}\}$$

and whose prediction is

$$\widetilde{\phi}_2(\omega) = \begin{cases} G & \text{for } \omega \in \{\omega_1, \omega_2, \omega_3, \omega_6\} \\ B & \text{otherwise} \end{cases}$$

It can be shown that this agent's prediction is independent of the prediction made by the agent in Example 1, even though the two interpretations are not independent.

Relatedly, we say that two predictions are independently correct, if knowing that one agent's prediction of an event is correct gives no information about whether the other's prediction of the same event is correct.

definition 3 *Predictions, $\widetilde{\phi}_1$ and $\widetilde{\phi}_2$, are* **independently correct predictions** *if $\widetilde{\delta}_1$ and $\widetilde{\delta}_2$ are independent random variables.*

Our next observation reveals the absence of a causal linkage between independent predictions and independently correct predictions. Even though seeing the world independently implies predicting independently, it need not imply being independently correct.

Observation 3 *Independent predictions need not be independently correct.*

Proof: Consider the following example:

Predictions	g	g	b	b
g	G	G	G	B
g	G	G	B	G
b	G	B	B	B
b	B	G	B	B

In this example, each of sixteen events is equally likely. The upper case letters represent outcomes of events. The lower case letters in the first column and in the first row are predictions of the row agent and the column agent respectively. By construction, the two predictions are independent. However, they are not independently correct. The joint probability of both agents making correct predictions is $\frac{1}{2}$ while the multiplication of the probabilities of each agent making correct predictions equals $\frac{3}{4} \times \frac{3}{4} = \frac{9}{16}$. They are not equal.

Finally, the correctness of predictions need not be independent even if the interpretations are.

Observation 4 *Independent interpretations may not lead to independently correct predictions.*

4 Results

We now present results within the interpreted signal framework. We first show independent predictions to be inconsistent with predictions being independently correct. Since independent interpretations imply independent predictions, these inconsistency results also apply to independent interpretations. We establish these results for the case of binary outcomes in which agents' predictions have identical probability distributions and equal accuracy.[10]

As before, we let upper case letters, G and B, refer to outcomes and lower case letters, g and b refer to predictions. Let $P(G)$ and $P(B)$ denote priors, assumed to be common among all agents. $P(g)$ and $P(b)$ denote the probabilities of predicting g and b respectively. $P(g,g)$ denotes the probability of both predicting g. $P(b,b)$, $P(g,b)$ and $P(b,g)$ are similarly defined. $P(c)$ and $P(i)$ denote the probabilities of making correct and incorrect predictions, which are also assumed to be the same for both agents. Finally $P(c,c)$, $P(i,i)$, $P(c,i)$ and $P(i,c)$ denote joint probabilities of both correct, both incorrect, agent 1 correct but agent 2 incorrect and agent 1 incorrect but agent 2 correct respectively. We impose the following symmetry assumptions: $P(g,b) = P(b,g)$ and $P(c,i) = P(i,c)$.

4.1 Reasonable and Informative Predictions

Given an interpretation, an agent need not make the best possible predictions. For example, an agent who categorized agents by gender could predict that women are taller than men. To impose some degree of competence we assume that predictions are either *reasonable* or *informative*.

[10] We have derived a set of results that do not assume that agents's predictions have identical distributions.

definition 4 *A prediction is* **reasonable** *if it is correct at least half of the time, i.e.,* $P(c) \geq \frac{1}{2}$.

definition 5 *A prediction is* **informative** *if it is correct more than half of the time, i.e.,* $P(c) > \frac{1}{2}$.

In the binary outcome case, an experience generated prediction must be reasonable. Further if at least one prediction is correct more than half of the time, then an experience generated prediction is also informative. However, as we observe next, an informative prediction need not predict correctly half of the time conditional on every state.

Observation 5 *An informative prediction need not be reasonable conditional on a state.*

This observation may be obvious to some. Nevertheless, we provide an example because the underlying logic is central to our analysis.

An Informative Prediction

Prediction	*Outcomes*		
g	G	G	B
g	G	G	B
g	G	G	B
g	G	G	B
b	B	B	G

Assuming each outcome to be equally likely, $P(c) = \frac{2}{3}$, implying that the prediction is informative. However, conditional on outcome B, the probability of making correct prediction equals $P(c \mid B) = P(b \mid B) = \frac{1}{3}$, implying that the prediction is not reasonable conditional on the bad state. This example provides insight into why when predicting rare events, agents may not make reasonable predictions.

We now state two lemmas that build to our results about negative correlation of interpreted signals. The lemmas that follow are interesting in their own right. They reveal a tension between the accuracy of predictions and the correlation of their correctness. When predictions are independent, the higher the accuracy, the less correlated their correctness is. In fact, highly accurate predictions must be negatively correlated in their correctness. In what follows, we relax the assumption that the probability of a good prediction and a bad prediction are equally likely. Without loss of generality, we assume $P(g) \geq \frac{1}{2}$.

Lemma 1 *The correctness of independent and reasonable predictions exhibit positive (zero, negative) correlation if and only if* $P(g) > [=, <]\ P(c)$.

The intuition that drives this result is straightforward. If the probability of predicting the good outcome is large relative to the probability of being correct, then both agents often predict good outcomes at the same time whether or not the prediction is correct. Thus, the correctness of their predictions must also be positively correlated.

Next, we reverse the assumption. We require that the predictions be independently correct. We can then show that the predictions themselves are independent or negatively or positively correlated depending again on the relationship between the probability of the more frequently picked prediction and the probability of being correct.

Lemma 2 *Independently correct and reasonable predictions exhibit positive (zero, negative) correlation if and only if* $P(c) > [=, <]\ P(g)$.

Our claim follows from Lemma 1 .

Claim 2 *Independent informative predictions that predict good and bad outcomes with equal probability must be negatively correlated in their correctness.*

We can restate this claim as follows:

Corollary 2 *Any independent and independently correct predictions that predict good and bad outcomes with equal probability cannot be informative.*

Since independent interpretations imply independent predictions, we also have the following:

Corollary 3 *If interpretations are independent and if their associated predictions are informative and predict good and bad outcomes with equal probability then the correctness of their predictions must be negatively correlated.*

In the special case, where the probability of good and bad outcomes are equal and, not only do we know that the interpreted signals must be negatively correlated, we also know the exact value of that correlation. In other words, the accuracy of independent interpreted signals uniquely determines their negative correlation.

Claim 3 *Independent interpreted predictions that predict good and bad outcomes with equal likelihood and are correct with probability* p *exhibit negative correlation defined by the following expression*

$$\rho = 1 - \frac{1}{4(p - p^2)}$$

pf. Define variable χ_i as the indicator of the correctness of signal s_i. In other words, $\chi_i = 1$ if s_i predicts correctly and otherwise $\chi_i = 0$. Then the correlation coefficient of the correctness of the signals, ρ, is defined as

$$\rho = \frac{Cov\ (\chi_1, \chi_2)}{\sqrt{Var\ (\chi_1)}\sqrt{Var\ (\chi_2)}}$$

Consider the following relationship between events and their probabilities. Let p_i denote the probability that s_i is correct. Then

$$\begin{aligned} p_1 &= \Pr(1 \text{ agrees with } 2 \text{ and } 1 \text{ is correct}) + \Pr(1 \text{ disagree with } 2 \text{ and } 1 \text{ is correct }) \\ &= \Pr(1 \text{ and } 2 \text{ agree and both correct}) + \Pr(1 \text{ and } 2 \text{ disagree and } 2 \text{ is incorrect}) \end{aligned}$$

And, similarly,

$$p_2 = \Pr(1 \text{ and } 2 \text{ agree and both correct}) + \Pr(1 \text{ and } 2 \text{ disagree and } 2 \text{ is correct})$$

Add the two equations, we have

$$p_1 + p_2 = 2\Pr(\text{both agree and correct}) + \Pr(1 \text{ and } 2 \text{ disagree})$$

The above equation hold without any specific assumptions. Now we assume that $p_i = p_j = p$ and that interpreted signals s_1 and s_2 are independent. Then

$$\Pr(1 \text{ and } 2 \text{ disagree}) = 2p(g)p(b) = 2p(g)(1 - p(g))$$

So,

$$\Pr(\text{both agree and correct}) = p - p(g)(1 - p(g))$$

We can now compute the correlation coefficient of the correctness of interpreted signals, ρ.

$$\begin{aligned} Var(\chi_i) &= E\chi_i^2 - (E\chi_i)^2 = p - p^2 \\ Cov\ (\chi_1, \chi_2) &= | : E(\chi_1\chi_2) - (E\chi_1)(E\chi_2) = p - p(g)(1 - p(g)) - p^2 \end{aligned}$$

Therefore,

$$\rho = \frac{p - p(g)(1 - p(g)) - p^2}{p - p^2} = 1 - \frac{p(g)(1 - p(g))}{p - p^2}$$

To complete the proof, note that $p(g) = \frac{1}{2}$, then

$$\rho = 1 - \frac{1}{4(p - p^2)}$$

This last two claims together with the corollaries reveal a fundamental conflict between an assumption that agents see the world independently (independent interpretations) and an assumption that the interpretation based predictions they make are independently correct.

4.2 Signal Accuracy, Correlation, and Function Complexity

We next discuss some findings that relate signal accuracy and correlation with the complexity of the outcome function. Intuition would suggest that the more complex the outcome function, the less accurate the signals and the less correlated the correctness of the signals. We show that while both of these intuitions hold true if we average across all possible functions, neither intuition holds for a large class of complex functions. Specifically, we show that these relationships depend on where the complexity lies in the space of attributes.

We restrict attention to the special case of a two dimensional attribute space in which each attribute takes on one of 2K values. We further assume that good and bad outcomes as well as good and bad interpreted signals are equally likely. This implies that we can write each outcome as a vector (x, y), where x and y take values in $\{1, ..2K\}$. Moreover, without loss of generality, we can assume that if $x \leq K$ (resp $> K$, then the interpreted signal based on the first attribute equals g (resp b) and that the same conditions hold for the interpreted signal based on the second attribute.

To proceed, we need some measure of the complexity of the outcome function. Complexity has many meanings (Miller and Page 2007) Here, by complexity we mean nonlinearities and interaction terms in the outcome function. Given that our function only takes two values, we can appeal to a crude but simpler measure that we call *attribute based value changes*, which we denote by ΔV. For each attribute value, we can count the number of times the value changes as that attribute ranges from 1 to K. For each attribute, we denote this value as $\delta(z)$To compute ΔV, we compute the sum of the $\delta(z)$'s. We show two examples below for the case $K = 3$.

	y_1	y_2	y_3	y_4	y_5	y_6	$\delta()$
x_1	1	1	1	1	1	1	**0**
x_2	1	1	1	1	1	1	**0**
x_3	1	1	1	1	0	0	**1**
x_4	1	1	0	0	0	0	**1**
x_5	1	0	0	0	0	0	**1**
x_6	1	0	0	0	0	0	**1**
$\delta()$	0	1	1	1	1	1	$\Delta = 9$

	y_1	y_2	y_3	y_4	y_5	y_6	$\delta()$
x_1	0	1	1	0	1	1	**3**
x_2	1	0	1	1	0	0	**3**
x_3	1	1	1	0	1	1	**2**
x_4	0	1	1	0	0	0	**2**
x_5	1	0	0	0	1	0	**3**
x_6	0	1	1	0	0	0	**2**
$\delta()$	**4**	**4**	**2**	**2**	**5**	**3**	$\Delta = 35$

Note from this second example, that switching attributes 2 and 3 lowers Δ.

	y_1	y_2	y_3	y_4	y_5	y_6	$\delta()$
x_1	0	1	1	0	1	1	**3**
x_3	1	1	1	0	1	1	**2**
x_2	1	0	1	1	0	0	**3**
x_4	0	1	1	0	0	0	**2**
x_5	1	0	0	0	1	0	**3**
x_6	0	1	1	0	0	0	**2**
$\delta()$	**4**	**4**	**2**	**2**	**3**	**1**	$\Delta = 31$

Given a outcome function V, we can define $\Delta^*(V)$ to be the minimum value of $\Delta(V)$ given any permutation of the attributes such that interpreted signals for the first K attribute values equal g for both attributes. We can partition the set of outcomes into two sets: the *agreement set* and the *disagreement set*.

definition 6 *The* ***agreement set*** $A = \{(x, y) \text{ s.t. } s_1(x, y) = s_2(x, y)\}$.

definition 7 *The* ***disagreement set*** $D = \{(x, y) \text{ s.t.} s_1(x, y) \neq s_2(x, y)\}$.

We can now state the following claims whose proofs are straightforward.

Claim 4 Changes in Δ in D have no effect on accuracy or the correlation of accuracy

Claim 5 Increases in Δ in A decrease accuracy and decrease the correlation of accuracy

These claims imply that all complexity relevant to correlation is captured in signal accuracy.

4.3 Overlapping Interpretations

We conclude our investigation of interpreted signals by considering cases in which agents overlap in the attributes that they include in their predictive models. These overlaps can result in correlated predictions. For example, if profits contribute to a firm's value, then two investors who both consider profits when making their predictions may make unconditionally positively correlated predictions. However, if we take this common attribute into account, we are left with independent interpretations. Therefore, all of our results for independent interpretations can be interpreted as independent conditional on common attributes.

We first show that most of the results for independent interpretations apply to overlapping interpretations as well. We then relate the properties of overlapping interpretations to the complexity of the outcome function. In our example, we assume five binary attributes determine whether an outcome is good or bad. One interpretation considers the first two attributes and the other considers only the third. These interpretations and the associated outcomes can be represented with rectangles. In each box, let (G, B) be the vector denoting the number of good and bad outcomes respectively so that $(3, 1)$ refers to three good outcomes and one bad outcome. Each box contains four outcomes as there are four possible values that the other two binary variables can take.

Non Overlapping Interpretations

1^{st} *and* 2^{nd} /3^{rd}	*# # 0*	*# # 1*
00#	(3,1)	(4,0)
10#	(2,2)	(1,3)
01#	(3,1)	(2,2)
11#	(2,2)	(0,4)

The row interpretation predicts good outcomes for the first and third rows. The column interpretation predicts good outcomes for the first column. The probability that the row player predicts correctly equals $\frac{23}{32}$, the probability that the column player predicts correctly equals $\frac{19}{32}$ and the probability that they are both correct equals $\frac{13}{32}$ which as we know is less than the product of the probabilities that each is correct given the negative correlation.[11]

We now assume that the column player also considers the second attribute. This generates the following rectangular representation. Some of the cells are now empty because the two interpretations conflict in those cells.

Overlapping Interpretations

1^{st} *and* 2^{nd} /2^{nd} *and* 3^{rd}	*# 0 0*	*# 0 1*	*# 1 0*	*# 1 1*
00#	3,1	4,0		
10#	2,2	1,3		
01#			3,1	2,2
11#			2,2	0,4

The column interpretation now predicts good outcomes everywhere but in the last column and is correct with probability $\frac{21}{32}$. The probability that both interpretations predict correctly equals $\frac{14}{32}$. The probability that they are both correct is still less than the product that each is correct.[12] Thus, in this example, the overlap in interpretations does not create positive correlation in the correctness of their predictions.

This example also reveals a diagonal structure to the outcomes. Once we take into account the common attributes, the interpretations are again independent. This intuition can be stated formally.

[11] $\frac{19*23}{32*32}$ is approximately $\frac{13.6}{32}$.
[12] $\frac{21*23}{32*32}$ is approximately $\frac{15.1}{32}$.

Observation 6 *Conditional on the values of their overlapping attributes, any two interpretations based on the same perspective are independent.*

By implication, if agents are aware of jointly considered attributes, then the correctness of their predictions will be negatively correlated conditional on the values of those attributes. All of the results that we derive for independent interpretations also hold for overlapping interpretations provided that the overlapping attributes are common knowledge. Therefore, the assumption of independent interpretations may not be especially restrictive.[13]

Unconditional on the values of the common attributes, the correctness of predictions can be positively correlated. This would be more likely the greater the predictive power of the common attributes. Common attributes of high predictive power imply less variance in outcomes within boxes along the diagonal and more variance in values across the diagonal boxes. For example, if profits are a crucial determinant of firm value, then in the high profit box, predictions are likely to be that the firms have high values, and in the low profit box, predictions are likely to be that the firms have low value. The overlap would then create positive correlation in both correctness and prediction

We can explore this intuition more generally. It suffices to consider a case in which the column interpretation adds an attribute already considered by the row interpretation. By symmetry, we can further restrict attention to the case where the column interpretation now predicts good outcomes in some cases where it previously predicted bad. Let X denote the set of outcomes previously predicted as bad outcomes but now predicted as good outcomes by the column interpretation. X can be partitioned into two sets X_G and X_B which denote the good and bad outcomes within X. Let lower case letters, x_G and x_B denote the cardinality of these sets. With the addition of the new attribute, the column interpretation is now correct in $(x_G - x_B)$ more instances. This number must be positive otherwise the correct prediction for the column interpretation would be bad rather than good.

Let X_G^{rG} denote the set of outcomes in X_G that the row interpretation predicts good outcomes. Define X_B^{rB} similarly. Then, the change in the number of outcomes where both interpretations are correct equals $(x_G^{rG} - x_B^{rB})$. Contrary to intuition this number can be negative and when positive, it can be larger than $(x_G - x_B)$.

Observation 7 *When an overlapping attribute is added to an interpretation, the number of outcomes in which both interpretations are correct can either decrease or increase. The increase in the number of outcomes in which both are correct can exceed the increase in the number of outcomes for which the altered interpretation is correct.*

Recall that the change in the number of cases that both are correct equals $x_G^{rG} - x_B^{rB}$. First, we show that both can be correct less often. This can happen so

[13]If we consider more than two interpretations and the overlap among these is not common, then independent interpretations can exist across pairs.

long as $x_G^{rG} < x_B^{rB}$. This inequality holds provided that in the set X, the row interpretation predicts bad outcomes correctly more often than it predicts good outcomes correctly. This condition is easily satisfied. Consider the example below. By definition, X denote the second column. Prior to adding the new attribute, the column interpretation predicted bad outcomes in X. Now it predicts good outcomes.

Overlapping Interpretations: Decreased Correlation

1^{st} *and* 2^{nd} /2^{nd} *and* 3^{rd}	*# 0 0*	*# 0 1*	*# 1 0*	*# 1 1*
00#	0,4	3,1		
10#	0,4	3,1		
01#			0,4	0,4
11#			0,4	0,4

The two bad outcomes in X were previously predicted correctly by both interpretations, now they are predicted correctly only by the row interpretation. All other outcomes that both predicted correctly are unchanged, so the total number of outcomes that both predict correctly falls. As a result, the correctness of predictions changed from being positively correlated to being negatively correlated after the column interpretation added the new attribute.

It remains to show that the change in the number of outcomes that both predict correctly can exceed the change in the number of outcomes that the column interpretation predicts correctly. Consider the following variant of the previous example. By definition, X consists of the first and second columns. The column interpretation used to predict bad outcomes in X, but now it predicts good outcomes.

Overlapping Interpretations: Increased Correlation

1^{st} *and* 2^{nd} /2^{nd} *and* 3^{rd}	*# 0 0*	*# 0 1*	*# 1 0*	*# 1 1*
00#	3,1	3,1		
10#	3,1	3,1		
01#			0,4	0,4
11#			0,4	0,4

The increase in the number of outcomes that the column interpretation predicts correctly equals $12 - 4 = 8$. The row interpretation only predicts good outcomes in X, so previously none of the outcomes in X were predicted correctly by both interpretations. That number now equals 12. Therefore, the increase in the number of outcomes that both are correct exceeds the increase in the number of outcomes that the column interpretation is correct. As a result, the correctness of predictions changed from being negatively correlated to being positively correlated after the column interpretation added the new attribute.

An implication of this observation is that we can never be certain of the effect of adding an overlapping attribute to each interpretation. Doing so can increase the amount of positive correlation in accuracy. But, the opposite can also occur: Even though the probability that each prediction is correct cannot decrease, the probability that both predictions are correct could fall. Which of these outcomes occurs depends upon the functional relationship between outcomes and attributes.

Note also the striking difference in the correlations of the correctness of predictions as a result of the new overlapping attribute in the two examples used in the proof of the previous claim. These differences can be attributed to the difference in the predictive power of the common attribute for the row player. In the first example, the common attribute does not have any predictive power for the row player while in the second example, it has a substantial impact.

From our many examples, it should be clear that the outcome function implicitly defines the statistical properties of the signals. We want to clarify how more complex functions can create independent signals even with overlap. Assume that the outcome function defined over five binary attributes equals one if and only if three or more of the attributes have value one. Consider two interpretations, each of which looks at four attributes. These two interpretations must overlap on at least three of the attributes. A straightforward calculations reveals that the predictions of the agents are, on average, positively correlated.

Next, we consider a more complicated function defined over these five attributes. For this function if the sum of the first three attributes is even, then the probability that the function takes value one equals $0.5 * (x_4 + x_5)$, but if the sum of the first three variables is odd, then the probability that the function takes value one equals $1 - 0.5 * (x_4 + x_5)$. It is a simple exercise to show that if one interpretation looks at the first four attributes and the other interpretation looks at the first three attributes and the fifth attribute, then the predictions are independent. This example shows that overlapping interpretations can be consistent with independent interpreted signals but only if the outcome function includes interactive terms.

5 Interpreted vs Generated Signals

We now discuss some of the differences between interpreted and generated signals. We first embed this interpreted signal framework within the standard signal framework where each signal is described by conditional (on states or outcomes in our terminology) probability distributions. We then explore, in a multi-agent signal model, the implication of the standard assumption of conditional independence between signals, keeping in mind that these signals are in fact interpreted. We show that the standard assumption of conditional independence, which is arguably quite reasonable for generated signals, implies a positive correlation structure on the interpreted signals.

5.1 Conditionally Independent Interpreted Signals

We first show how to construct a collection of signals which are independent conditional on the value of the outcome within our interpretation framework. We restrict attention to cases where the objects have equal probability and that good and bad outcomes are equally likely. Let $p = \frac{r}{m}$ denote the probability that a signal is correct. We assume that $2r > m > r$ so that this probability lies in the open interval $(0.5, 1)$. The fraction $\frac{r}{m}$ will also equal the probability that a signal is correct conditional on each state. To construct K interpreted signals that are conditional independent on the state we set N, the number of objects of equal probability equal to $2m^K$.

We can denote a object as a vector of $K+1$ attributes in which the first attribute takes one of two values, for convenience, 0 or 1, and each of the remaining K variables take values in the set $\{1, 2.....m\}$. As a matter of convention, we write an object as a vector of attributes $(\theta, x_1, x_2, ...x_m)$. We construct the payoff function so that if an even number of the last K attributes are greater than r, the value of the function equals θ. Otherwise, the value equals $(1-\theta)$.

$$
\begin{aligned}
f(\theta, \vec{x}) &= \theta \text{ if } \quad \mid \{i : x_i > r\} \mid = 2j \text{ for some } j \\
&= 1-\theta \quad \text{else}
\end{aligned}
$$

We define the interpretations and the interpreted signals as follows. Interpretation i considers every attribute except attribute i, i.e. $(\theta, \vec{x_{-i}})$. The *interpreted signal*, s_i based on this interpretation equals θ if an even number of the attributes other than i take values greater than r and equals $1-\theta$ otherwise:

$$
\begin{aligned}
s_i(\theta, \vec{x_{-i}}) &= \theta \text{ if } \quad \mid \{j \neq i : x_j > r\} \mid = 2j \text{ for some } j \\
&= 1-\theta \quad \text{else}
\end{aligned}
$$

Given our assumption that all outcomes are equally likely, the probability that this signal is correct equals $\frac{r}{m}$. Further, that is also true conditional on either state. A straightforward exercise establishes that the interpreted signals s_i and s_j are also independent conditional function's value.

Unlike generated signals, in which an agent gets a signal that is correct with some probability, these interpreted signals correspond to models that leave out one piece of information – the value of an attribute. Most of the time, the realization of that value will not change the value of the outcome, but sometimes it will. Thus, it is possible for conditionally independent signals as occurring provided each agent constructs a predictive model that ignores a different attribute and if the functional form has the property that each attribute has a proportionally similar effect on the outcome regardless of the values of the other attributes.

As we now show, this last restriction rules out all functions over attributes except those isomorphic to the class of examples we just described.

Claim 6 *Any conditionally independent signal can be mapped into this class of examples by combining events*

Let f be an outcome function, taking values 0 and 1 with equal probability. Let s_1, s_2,...,s_K (where $K \geq 2$) be any collection of signals that are independent conditional on the value of the outcome function in our framework. For each signal s_i, we assume that the probability of being correct is $\frac{r_i}{m_i}$ where $2r_i > m_i > r_i$. **We also assume that $\frac{r_i}{m_i}$ equals the probability that signal s_i is correct conditional on each outcome.** (This seems to be a crucial assumption for the proof to work.). Then f and s_1, s_2, ..., s_K can be mapped into our interpretation framework in the following way.

Define $K+1$ attributes, $(\theta, x_1, ..., x_K)$, where $\theta \in \{0,1\}$ and $x_i \in \{1, ..., m_i\}$. Call each realization of attributes $(\theta, x_1, ..., x_K)$ a state. Let each state be equally likely. Then the given f and s_1, s_2, ..., s_K can be reduced forms of the following: f takes the value θ if an even number of the last K attributes are greater than r_i respectively. Otherwise, the value equals $(1-\theta)$.

$$\begin{aligned} f(\theta, \vec{x}) &= \theta \text{if } \mid \{i : x_i > r_i\} \mid := 2k : \text{ for some } : k \\ &= 1-\theta \quad \text{else} \end{aligned}$$

For any $i \in \{1, ..., K\}$, it considers every attribute except attribute x_i, i.e. $(\theta, \vec{x_{-i}})$. The *interpreted signal*, s_i based on this interpretation equals θ if an even number of the attributes other than i take values greater than r_j respectively and equals $1-\theta$ otherwise:

$$\begin{aligned} s_i(\theta, \vec{x_{-i}}) &= \theta \text{if } \mid \{j \neq i : x_j > r_j\} \mid := 2k : \text{ for some } : k \\ &= 1-\theta \quad \text{else} \end{aligned}$$

An example is provided here. Let there be two signals. Each signal is correct with probability $\frac{2}{3}$. In each cell below, the number that is not in the parenthesis represents the value of f. The first (second) number in the parenthesis represents the value of s_1 (s_2).

$\theta = 0$

$x_1 \setminus x_2$	1	2	3
1	0 $(0,0)$	0 $(0,0)$	1 $(1,0)$
2	0 $(0,0)$	0 $(0,0)$	1 $(1,0)$
3	1 $(0,1)$	1 $(0,1)$	0 $(1,1)$

$\theta = 1$

$x_1 \setminus x_2$	1	2	3
1	1 $(1,1)$	1 $(1,1)$	0 $(0,1)$
2	1 $(1,1)$	1 $(1,1)$	0 $(0,1)$
3	0 $(1,0)$	0 $(1,0)$	1 $(0,0)$

Given our assumption that all states are equally likely, it is straightforward to show that the probability that s_i is correct equals $\frac{r_i}{m_i}$. That is also true conditional on either of the outcome function's values. Now we show that these signals are independent conditional on the function's values. First, observe the following relationship among the three events.

$$\{f=0\}=\{\theta=0 \ \& \ \{i:x_i>r_i\} \text{ is even}\}\cup\{\theta=1 \ \& \ \{i:x_i>r_i\} \text{ is odd}\}$$

and the two events on the right of the equation are disjoint. Now for any $k\in\{1,...,K\}$ we have that

$$\begin{aligned}
&\{s_k=f=0\}\\
&=\{s_k=\theta=0 \ \& \ \{i:x_i>r_i\} \text{ is even}\}\cup\{s_k=0 \ \& \ \theta=1 \ \& \ \{i:x_i>r_i\} \text{ is odd}\}\\
&=\{x_k\leq r_k \ \& \ \theta=0 \ \& \ \{i:x_i>r_i\} \text{ is even}\}\cup\{x_k\leq r_k \ \& \ \theta=1 \ \& \ \{i:x_i>r_i\} \text{ is odd}\}
\end{aligned}$$

Therefore,

$$\begin{aligned}
&\Pr(s_k=f=0)\\
&=\Pr\left(x_k\leq r_k \ \& \ \theta=0 \ \& \ \{i:x_i>r_i\} \text{ is even}\right)+\Pr\left(x_k\leq r_k \ \& \ \theta=1 \ \& \ \{i:x_i>r_i\} \text{ is odd}\right)\\
&=\Pr\left(x_k\leq r_k \ \& \ \theta=0 \ \& \ \{i:x_i>r_i\} \text{ is even}\right)+\Pr\left(x_k\leq r_k \ \& \ \theta=0 \ \& \ \{i:x_i>r_i\} \text{ is odd}\right)\\
&=\Pr(x_k\leq r_k \ \& \ \theta=0)
\end{aligned}$$

Note that the second to the last equality utilizes the assumption that conditional on the value of the outcome function, the probability of being correct is the same. Similarly, we can show that for any subset $\{k_1,...,k_j\}\subseteq\{1,...,K\}$,

$$\Pr(s_{k_1}=0 \ \& \ ... \ \& \ s_{k_j}=0 \ \& \ f=0)=\Pr(x_{k_1}\leq r_{k_1} \ \& \ ... \ \& \ x_{k_j}\leq r_{k_j} \ \& \ \theta=0).$$

Similar relations hold for $f=1$ and $s_k=1$ hold with the corresponding adjustments that $\theta=1$ and $x_k>r_k$. This means that conditional on the value of the outcome function, the signals $s_1,...,s_K$ are identical to the signals based attributes $x_1,...,x_K$, each observing the value of the corresponding attribute, conditional on the value of θ. Given the rectangle structure of these signals conditional on θ, $s_1,...,s_K$ are therefore independent conditional on the value of f.

This claim implies that while it is possible to create conditionally independent interpreted signals, doing so implies a unique outcome function and requires that each agent neglects one piece of information. Thus, paradoxically, conditionally independent signals are not consistent with agents looking at different attributes of an arbitrary function but they are consistent with agents neglecting different pieces of information give a specific function mapping attributes to outcomes.

5.2 Mapping Interpreted Signals into Generated Signals

We can also perform the inverse translation. We can map interpreted signals into the generated signal framework. Relatively speaking, this task is much simpler . Let p denote the probability that an agent predicts g conditional on the true outcome being G and q denote the probability that an agent predicts b conditional on the true outcome being B, that is,

$$p=P(g\mid G)$$

$$q = P(b \mid B)$$

Interpreted signals can then be written as a typical binary signal model with the following conditional distributions (conditional on outcomes)

Conditional Probability Distribution of Signals

*outcome**signal*	g	b
G	p	$1-p$
B	$1-q$	q

Consistent with the notation from the previous section, we have the following unconditional distribution of predictions.[14]

$$P(g) = P(G)p + P(B)(1-q)$$

and

$$P(b) = P(G)(1-p) + P(B)q$$

We can now relate independent interpretations and predictions to signals that are independent conditional on outcomes. The first claim states that informative and experience generated predictions – therefore, reasonable predictions – that satisfy independence conditional on outcomes must be unconditionally positively correlated. In other words, the assumption that interpreted signals satisfy the standard assumption from signaling models (independence conditional on outcomes) implies that the predictions themselves are positively correlated and cannot come from independent interpretations.

Claim 7 *Experience generated and informative predictions that are independent conditional on outcomes must be positively correlated unconditionally.*

The intuition behind this claim can be seen in an example. Suppose that good and bad outcomes are equally likely and that conditional on the state each of two agents predicts correctly with probability $\frac{2}{3}$. Suppose, for example, that the environment consists of nine good outcomes and nine bad outcomes and that each is equally likely.

[14]Binary generated signalsare often assumed to satisfy the *Strong Monotone Likelihood Ratio Property (SMLRP).* Using the notation above, a signal satisfies the SMLRP if and only if

$$\frac{p}{1-q} > \frac{1-p}{q}$$

or equivalently

$$p + q > 1.$$

It can be shown that if predictions are experience generated and informative, then they satisfy the SMLRP.

For the predictions to be independent conditional on the state, two agents would have to both predict four of the good outcomes correctly and both predict one of the good outcomes incorrectly. Each would also have to predict two good outcomes correctly that the other predicted incorrectly. The same holds for the bad outcomes. We can represent this in a table.

Conditionally Independent Predictions

agent 1 /agent 2	$PredictsG$	$PredictsB$
$PredictsG$	$4G$ and $1B$	$2G$ and $2B$
$PredictsB$	$2G$ and $2B$	$1G$ and $4B$

Notice that the predictions are not independent. If agent 1 predicts G then agent 2 predicts G with probability $\frac{5}{9}$. If agent 1 predicts B then agent 2 predicts G with probability $\frac{4}{9}$. The predictions are positively correlated. Independence conditional on the state requires that the predictions be positively correlated.

We can now state the flip side of Claim 7.

Claim 8 *Assume that predictions are experience generated and are informative. If predictions are independent, then for at least one outcome, predictions conditional on that outcome are negatively correlated.*

Either of the above two claims implies an inconsistency between conditional independence and independent interpretations.

Claim 9 *Conditional independence of predictions is inconsistent with informative and experience generated independent predictions and therefore with independent interpretations.*

This final claim parallels our previous result showing a conflict between reasonable, informative, independent predictions and independently correct predictions. These two claims together reveal a fundamental incompatibility between standard signaling assumptions and independent interpretations: *Seeing the world independently, looking at different attributes, not only does not imply, it is inconsistent with, both conditional independence of signals and independently correct signals.*

6 Affiliation Assumption and Signal Type

We have established that the statistical properties of generated and interpreted signals differ. In this section, we use models of common value auctions as an example to illustrate the relevance of our results. In a companion paper, we apply interpreted signals in the context of voting models Our investigation of affiliated values here provides a brief, but powerful example of why the distinction matters.

In what follows, we consider generated and interpreted signals and their relationships to the affiliation assumption in a common value auction model. The literature devoted to characterizing optimal bidding strategies and the efficiency of auction mechanisms admits various assumptions about values, in particular whether they are public or private. The common value framework is the more relevant setting for thinking about generated and interpreted signals as they concern a single object with an unknown value. In a seminal paper, hereafter denoted as MW, Milgrom and Weber (1982) identified a symmetric Nash equilibrium for the second price auction in a general setting with both private and common value components. Their results require the assumption of affiliation. Affiliation is satisfied when if an agent obtains a high signal, it is likely that other agents also obtain high signals and that the unknown common value is likely to be high.

MW describes a common value auction in which each agent obtains a private signal. They assume that the joint distribution of the unknown common value and agents' signals satisfies the affiliation assumption. This view of the common value auction is natural if signals are generated. MW further shows that if agents' signals are independent conditional on true values and if they all satisfy the Monotone Likelihood Ratio Property, then the signals and the unknown common value also satisfy the affiliation assumption.

The affiliation assumption imposes restrictions on the nature of the signal. Using a simple example, we show that if the signals are attributes of the object with a common value, i.e. if they are *interpreted signals*, then the affiliation assumption is less restrictive. In this example, we consider two bidders and a common value object. The common value, denoted by v, can take two possible values, 0 or 1. Prior to bidding, each agent gets a signal x_i. x_i can either be 0 or 1. The joint distribution of signals and the value is described as follows: $p(x_1,\ x_2; v) = \frac{1}{4}$ if $x_1 = 0, x_2 = 0$, and $v = 0$; or $x_1 = 1, x_2 = 0$, and $v = 1$; or $x_1 = 0, x_2 = 1$, and $v = 1$; or $x_1 = 1, x_2 = 1$, and $v = 1$. Other possibilities all have 0 probability. Notice that $p(0,0;1) \cdot p(1,1;1) = 0 < \frac{1}{16} = p(1,0;1) \cdot p(0,1;1)$. This construction therefore violates affiliation.

Nevertheless, MW's strategy can be shown to be optimal. A strategy of bidding 0 upon receiving the signal 0 and bidding 1 upon receiving the signal 1 is a symmetric equilibrium.[15] If we treat the signals as interpreted, then the affiliated values assumption is satisfied. To show this we first write the relationship between signals and the common value given by the joint probability distribution of signals using an outcome function that depends on two attributes x_1 and x_2 :

$$v = f(x_1, x_2) = x_1 + x_2 - x_1 x_2$$

[15] This example makes clear that MW's assumptions are sufficient but not necessary for their result, but that is not our reason for discussing it.

where x_i can either be 0 or 1 with equal probability and x_1 and x_2 are independent. Prior to bidding, bidder i learns the value of x_i and bids accordingly. This is the bidder's interpreted signal. The utility function of each bidder is given by

$$u_i\left(x_1, x_2, v\right) = v = f(x_1, x_2) = x_1 + x_2 - x_1x_2$$

Instead of thinking of each bidder's utility u_i as an increasing function of the common value (which is what MW did), we think of it as an increasing function of the the bidder's interpreted signal. Since x_1 and x_2 are independent, they are trivially affiliated. Thus, when we write the variables as interpreted signals, the affiliation assumptions is trivially satisfied. This example belongs to the class of monotonic, multi attribute, binary valued functions. Our next claim states that no function in this class creates an affiliated joint distribution over the signals and the value. That is, if we treat the signals as generated, as in MW, the affiliation assumption cannot be satisfied.

Claim 10 *Let $X = \{0, 1\}^N$. If $f : X \rightarrow \{0, 1\}$ is onto and monotonic and if for all $i \in \{1, 2, ..., N\}$, $f(x) \neq x_i$ for some x, then for any distribution P over X such that $P(x) > 0$ for all $x \in X$, the corresponding joint distribution over the signals and the value violates the affiliation assumption.*

Returning to and extending our previous argument, if we model the attributes as interpreted signals, then every function in this class combined with a distribution P (over $(x_1, ..., x_n)$ only, not over $(x_1, ..., x_n,\ f(x))$) satisfies the conditions of the MW model. Thus, affiliation places much stronger restrictions on generated signals than on interpreted signals.

Up until now, we've assumed the function is monotonic in each attribute. We next show that we can relax that assumption and still satisfy MW's conditions. For example, if

$$v = f(x_1, x_2) = x_1 + x_2 - \theta x_1 x_2$$

where θ is a constant greater than 1, then the outcome function is not monotonic in each attribute because of the large negative interaction term. One can easily show that the symmetric bidding strategy

$$b_i(\overline{x}_i) = E(v \mid x_i = \overline{x}_i, x_j = \overline{x}_i) = 2\overline{x}_i - \theta(\overline{x}_i)^2$$

is not an equilibrium.

The failure to satisfy monotonicity need not undermine MW's results. This is particularly true if we allow for endogenous acquisition of information. Consider a modified example where

$$v = f(x_1, x_2) = x_1 + x_2 + x_3 + x_4 - \theta x_1 x_2$$

and $\theta > 1$. So again because of the large negative interaction term between x_1 and x_2, the outcome function is not monotonic in x_1 or x_2. But if prior to bidding, the first bidder can observe x_1, x_2 and x_3, the second bidder can observe x_1, x_2 and x_4, and this information is common knowledge among the bidders, then MW's assumptions hold as long as bidders take into account their common information. That is, $b_1(x_1, x_2, x_3) = x_1 + x_2 + 2x_3 - \theta x_1 x_2$ and $b_2(x_1, x_2, x_4) = x_1 + x_2 + 2x_4 - \theta x_1 x_2$ constitute an equilibrium.

Thus, even if the outcome function contains large negative interaction terms, as long as the attributes producing the interactions belong to the overlap of bidder's information, the large negative interaction terms can be absorbed, and the rest of Milgrom and Weber's results carry through.[16]

7 Discussion

In this paper, we have contrasted generated and interpreted signals and demonstrated important differences between them. These differences have implications for the generality of claims of the optimality of strategies and of institutional designs. At a minimum, our results suggest modelers should relate their assumptions to the specific context: are the signals generated, interpreted, or possibly both? In addition, the distinction between the two types of signals might enable us to better understand differences between experimental and real world results. In experiments, information is often generated using the standard conditionally independent signal model. In practice, it may not be. Therefore, testing our theory using experiments may not be testing one of the most important assumptions: the assumption of conditionally independent signals.

This insight also applies to attempts to calibrate computational models with standard models of signals. These efforts may also run up against this fundamental inconsistency. In an agent based model (Tesfatsion 1997, Holland and Miller 1991), the signals are often lower dimensional projections of a larger reality. In rich, fine detailed computer models, such as the trading agent competition (Wellman et al 2003), agents do not take into account all of the information in the environment. Instead, they monitor a lower dimensional world than the one within which they interact. In spatial models and network models, something close to dimensional reduction also occurs. Agents only see what happens in a local region, thereby creating interpreted signals.

Owing to its close connections to computer science, the interpreted signal framework can be seen as a computational approach to incomplete information. Ideally, computational and mathematical models inform and complement one another (Judd 1997, Judd and Page 2004). However, our ability to align computational and mathematical models is hindered if the assumptions about signals that we make in our

[16] Whether or not, in a strategic context, the endogenous acquisition of information of attribute values would result in this overlap is an open question.

mathematical models are not consistent with the information that agents realize in the computational implementations of those models.

The interpretation signal framework allows for the modeling of endogenous signals. If agents want to predict correctly individually, this could lead to correlated signals as they might all learn to look at the same attributes. If agents are concerned with collective performance, such as in the case of voting to aggregate information, agents have an incentive to look at different attributes. These insights are not surprising. What is surprising is that by looking at different attributes, the correctness of agents's predictions is often negatively correlated, so their information should aggregate rather well. However, the number of different attributes that can be considered depends upon the dimensionality of the problem. Therefore, large groups of agents may do worse than generated signal theory predicts because they necessarily have lots of overlap. Small groups, in contrast, may do better than the generated signal theory predicts.

The choice over which attributes to include in interpretations in competitive situations, such as auctions, is among the most interesting questions to consider. Competitive situations create both types of incentives: an incentive to be correct and an incentive to be different. As we saw in the example above on the common value auction, bidders may simplify the strategic environment by absorbing externalities. The question of whether such simplification, where possible, is also incentive compatible, is an open one.

Finally, the interpretation framework also permits more fine grained analysis of the link between complexity and uncertainty. Many of our examples involve nonlinear mappings that include interaction terms. The complexity-uncertainty link is also the focus of a paper by Al-Najjar, Casadesus-Masanell, and Ozdenoren (2003). They consider the continual addition of more and more attributes. As the number of attributes considered increases, the signals should improve. A problem is complex if no matter how many attributes are considered, the uncertainty never goes away. In our framework, within some sets in a partition, both good and bad outcomes can exist. Our formulation highlights a related notion of complexity - nonlinearity and interaction terms in the mapping from attributes to outcomes. As this mapping becomes more complex, the inference problem becomes more difficult. Moreover, anomalies, such as adding an overlapping attribute creates less correlated predictions, are more likely to occur. What we might call regularities in signals should be related in a systematic way to this second conception of complexity.

8 Appendix: Proofs

Proof of Claim 1. We prove the claim for $n = 2$. The proof for more general cases follows the same procedure.

Without loss of generality, assume that $\Pi^i = \left\{\pi_1^i, \cdots, \pi_{n_i}^i\right\}$ where for each $i = 1, 2$, $n_i \geq 2$. We write the event space, Ω, in the following form which helps to visualize the proof.

$$\begin{array}{cccc} \pi_1^1 \cap \pi_1^2 & \pi_1^1 \cap \pi_2^2 & \cdots & \pi_1^1 \cap \pi_{n_2}^2 \\ \pi_2^1 \cap \pi_1^2 & \pi_2^1 \cap \pi_2^2 & \cdots & \pi_2^1 \cap \pi_{n_2}^2 \\ . & . & . & . \\ . & . & . & . \\ . & . & . & . \\ \pi_{n_1}^1 \cap \pi_1^2 & \pi_{n_1}^1 \cap \pi_2^2 & \cdots & \pi_{n_1}^1 \cap \pi_{n_2}^2 \end{array} \tag{1}$$

Each cell above can be represented by a 2-dimensional rectangle with the property that cells so represented in the same row have the same height and cells in the same column have the same width.

To show this, we first show that for each $j = 2, ..., n_2$, that the number of events contained in each cell in any given column is proportional to the number of events in each cell in the first column:

$$\frac{\left|\pi_1^1 \cap \pi_j^2\right|}{\left|\pi_1^1 \cap \pi_1^2\right|} = \frac{\left|\pi_2^1 \cap \pi_j^2\right|}{\left|\pi_2^1 \cap \pi_1^2\right|} = \cdots = \frac{\left|\pi_{n_1}^1 \cap \pi_j^2\right|}{\left|\pi_{n_1}^1 \cap \pi_1^2\right|} \tag{2}$$

where $|\cdot|$ denotes the cardinality of a set. By independence (recall that each event in Ω is equally likely), for all $i = 1, ..., n_1$ and all $j = 2, ..., n_2$,

$$\frac{\left|\pi_i^1 \cap \pi_j^2\right|}{N} = \frac{\left|\pi_i^1\right|}{N} \cdot \frac{\left|\pi_j^2\right|}{N}$$

and

$$\frac{\left|\pi_i^1 \cap \pi_1^2\right|}{N} = \frac{\left|\pi_i^1\right|}{N} \cdot \frac{\left|\pi_1^2\right|}{N}$$

Therefore,

$$\frac{\left|\pi_i^1 \cap \pi_j^2\right|}{\left|\pi_i^1 \cap \pi_1^2\right|} = \frac{\left|\pi_j^2\right|}{\left|\pi_1^2\right|}$$

This proves (2) above.

For each $j = 2, ..., n_2$, let the ratio in (2) be equal to $\frac{u_j}{d_j}$ where both u_j and d_j are positive integers and $\frac{u_j}{d_j}$ can not be further simplified. That is, for each $i = 1, ..., n_1$, we can write the number of events in the ith row and jth column as $\frac{u_j}{d_j}$ times the number of events in the first column of the ith row.

$$\left|\pi_i^1 \cap \pi_j^2\right| = \frac{u_j}{d_j} \cdot \left|\pi_i^1 \cap \pi_1^2\right|$$

This implies that for each $i = 1, ..., n_1$, $|\pi_i^1 \cap \pi_1^2|$ is divisible by all d_j's, $j = 2, ..., n_2$. Let d be the smallest positive integer that is divisible by all d_j's. Then for each $i = 1, ..., n_1$, there exists a unique positive integer k_i such that

$$\left|\pi_i^1 \cap \pi_1^2\right| = k_i \cdot d.$$

Thus,

$$\left|\pi_i^1 \cap \pi_j^2\right| = k_i \cdot \left(u_j \cdot \frac{d}{d_j}\right)$$

for all $i = 1, ..., n_1$ and $j = 2, ..., n_2$. Notice that $\frac{d}{d_j}$ is a positive integer in the above expression.

The above argument proves that for any $i = 1, ..., n_1$ and $j = 1, 2, ..., n_2$, $\pi_i^1 \cap \pi_j^2$ can be represented by a 2-dimensional rectangle of k_i rows (height) and $u_j \cdot \frac{d}{d_j}$ columns (width). Here we have implicitly defined $u_1 = d_1 = 1$. Therefore, each cell in (2) can be represented by a 2-dimensional rectangle such that cells in row i all have the same height of k_i and cells in column j all have the same width of $u_j \cdot \frac{d}{d_j}$. Therefore, (2) can be represented by a 2-dimensional rectangle with $\sum_{i=1}^{n_1} k_i$ rows and $\sum_{j=1}^{n_2} u_j \cdot \frac{d}{d_j}$ columns. That means, $N = \left(\sum_{i=1}^{n_1} k_i\right) \cdot \left(\sum_{j=1}^{n_2} u_j \cdot \frac{d}{d_j}\right)$. It is obvious that the number in each parenthesis is larger than 1.

Proof of Corollary 1. By Claim 1, a necessary condition for n non-trivial interpretations to be independent is that N can be written as the multiplications of n larger-than-1 integers. Thus, the largest number of independent non-trivial interpretations is bounded by the number of prime factors which is k. Now we only need to show that there exist k many non-trivial interpretations that are independent. When $N = \prod_{i=1}^{k} p_i$, Ω can be represented by a k-dimensional rectangle where the ith dimension has a length of p_i. Let Π^i be the interpretation that can only identify events along the ith dimension. Showing that these k interpretations are independent is a straightforward exercise.

Proof of Lemma 1, Lemma 2 and Claim 2. First, observe the following identity:

$$P(g, b) + P(b, g) = P(c, i) + P(i, c)$$

Each side of this equation expresses the probability that agents disagree. Then by symmetry,

$$P(g, b) = P(c, i)$$

Second, notice that the function $x(1 - x)$ is a decreasing function of x for $x \geq \frac{1}{2}$. Therefore,

$$P(g)(1 - P(g)) < [=, >]\, P(c)(1 - P(c))$$

if and only if

$$P(g) > [=, <]\, P(c)$$

That is,

$$P(g)P(b) < [=, >]\, P(c)P(i)$$

if and only if

$$P(g) > [=, <]\, P(c)$$

Now we prove Lemma 1. If predictions are independent, then

$$P(g, b) = P(g)P(b)$$

Also, by definition, the correctness of predictions are positively correlated (independent or negatively correlated) iff

$$P(c, i) < [=, >]\, P(c)P(i)$$

Combine the above two equations with the identity at the beginning of the proof, we have the correctness of predictions are positively correlated (independent or negatively correlated) iff $P(g)P(b) < [=, >]\, P(c)P(i)$. The result then follows. Claim 2 is a special case of Lemma 1. Lemma 2 can be similarly proved.

Proof of Claim 7. We need to show

$$P(g)^2 < P(g, g)$$

Here, $P(g, g)$ denote the joint probability of both agents predicting g. We know

$$P(g) = P(G)p + P(B)(1 - q)$$

Now we compute $P(g, g)$. Since predictions are conditionally independent,

$$P(g, g) = P(G)p^2 + P(B)(1 - q)^2$$

Therefore,

$$\begin{aligned}
&P(g, g) - P(g)^2 \\
&= P(G)p^2 + P(B)(1 - q)^2 - [P(G)p + P(B)(1 - q)]^2 \\
&= P(G)P(B)\,(p + q - 1)^2
\end{aligned}$$

Since the predictions are experience generated and informative,

$$p + q > 1$$

This means that

$$P(g, g) - P(g)^2 > 0$$

Therefore, predictions are unconditionally positively correlated.

Proof of Claim 8. We prove this claim by way of contradiction. Suppose both conditional distributions of predictions are not negatively correlated. Then

$$P(g, g \mid G) \geq p^2$$

and

$$P(g, g \mid B) \geq (1-q)^2$$

Since predictions are independent,

$$P(g)^2 = P(g, g)$$

By definition,

$$P(g) = P(G)p + P(B)(1-q)$$

and

$$P(g, g) = P(G)P(g, g \mid G) + P(B)P(g, g \mid B)$$

Thus,

$$[P(G)p + P(B)(1-q)]^2 \geq P(G)p^2 + P(B)(1-q)^2$$

which simplifies to

$$P(G)P(B)\,(p+q-1)^2 \leq 0$$

Since

$$p + q > 1$$

a contradiction.

Proof of Claim 10. By assumption $P(x) > 0$ for any $x \in X$. Let p be the corresponding joint probability distribution over $(x; v)$, e.g. if $v = f(x)$, $p(x; v) = P(x)$, and if $v \neq f(x)$, then $p(x; v) = 0$. For any $x \in X$, let $K(x) = \{i : x_i = 1\}$ be the set of attributes that take value 1. Similarly, define $x(K) = (x_1, x_2, \ldots, x_n)$ where $x_i = 1$ iff $i \in K$. Then there exists a $K^* \subset N$ such that $f(x(K^*)) = 0$ and for any K that strictly contains K^*, $f(x(K)) = 1$. In general, for any given f, there can be multiple K^*'s with this property. We concentrate on any one set that has the largest cardinality and still call it K^* to keep the notation simple. By the assumption that f is onto and monotonic, $\mid K^* \mid$ is strictly less than N. We now consider two cases. First $\mid K^* \mid= N-1$ and second $\mid K^* \mid\leq N-2$.

(1) $\mid K^* \mid= N-1$. Without loss of generality, assume $K^* = \{1, ..., N-1\}$. This means $f(1, ...1, 0) = 0$ and $f(1, ...1, 1) = 1$. By monotonicity, $f(x_{-N}, 0) = 0$ for any $x_{-N} \in \{0,1\}^{N-1}$. By the assumption that for any i, $f(x) \neq x_i$ for some x, we know that there exists $x_{-N} \in \{0,1\}^{N-1}$ such that $f(x_{-N}, 1) = 0$. By monotonicity again, $f(0, ...0, 1) = 0$. We have so far established that the joint probability distribution over $(x; v)$ satisfies the following: $p(0, ..., 0, 0; 0) > 0$, $p(1, ..., 1, 1; 0) = 0$, $p(0, ..., 0, 1; 0) > 0$, and $p(1, ...1, 0; 0) > 0$. However this violates affiliation of the joint probability distribution because affiliation requires that

$$p(0, ..., 0, 1; 0) \cdot p(1, ...1, 0; 0) \leq p(0, ..., 0, 0; 0) \cdot p(1, ..., 1, 1; 0)$$

(2) $\mid K^* \mid\leq N-2$. We prove the claim by contradiction. Suppose that the joint probability distribution over $(x; v)$ satisfies affiliation. Choose $j, j' \notin K^*$. We show

that $f(x(\{j\})) = f(x(\{j'\})) = 1$. We prove this for j. The proof for j' is identical. Suppose $f(x(\{j\})) = 0$, then $p(x(\{j\}); 0) > 0$. Affiliation requires that

$$p(x(K^*); 0) \cdot p(x(\{j\}); 0) \leq p(x(\emptyset); 0) \cdot p(x(K^* \cup \{j\}); 0)$$

But by definition, $p(x(K^* \cup \{j\}); 0) = 0$ leading to a contradiction. Thus, we have that $p(x(\{j\}); 1) > 0$ and $p(x(\{j'\}); 1) > 0$. Affiliation requires that

$$p(x(\{j\}); 1) \cdot p(x(\{j'\}); 1) \leq p(x(\emptyset); 1) \cdot p(x(\{j, j'\}); 1)$$

But by the assumption that f is onto and monotonic, $p(x(\emptyset); 1) = 0$ which leads to a contradiction.

References

[1] Al-Najjar, N., R. Casadesus-Masanell and E. Ozdenoren (2003) "Probabilistic Representation of Complexity", *Journal of Economic Theory* 111 (1), 49 - 87.

[2] Aragones, E., I. Gilboa, A. Postlewaite, and D. Schmeidler (2005) "Fact-Free Learning", *The American Economic Review* 95 (5), 1355 - 1368.

[3] Barwise and Seligman, (1997) *Information Flow: The Logic of Distributed Systems* Cambridge Tracts In Theoretical Computer Science, Cambridge University Press, New York.

[4] Feddersen, T. and W. Pesendorfer (1997) "Voting Behavior and Information Aggregation in Elections with Private Information", *Econometrica* 65 (5), 1029 -1058.

[5] Fryer, R. and M. Jackson (2003), Categorical Cognition: A Psychological Model of Categories and Identification in Decision Making, Working Paper, California Institute of Technology.

[6] Holland, J. and J. Miller (1991) "Artificial Agents in Economic Theory", *The American Economic Review Papers and Proceedings* 81, 365 - 370.

[7] Hong L. and S. Page (2001) "Problem Solving by Heterogeneous Agents", *Journal of Economic Theory* 97, 123 - 163.

[8] Judd, K. (1997) "Computational Economics and Economic Theory: Complements or Substitutes?" *Journal of Economic Dynamics and Control.*

[9] Judd, K. and S. Page (2004) "Computational Public Economics", *Journal of Public Economic Theory* forthcoming.

[10] Klemperer, P. (2004) *Auctions: Theory and Practice* Princeton University Press.

[11] Ladha, K. (1992) "The Condorcet Jury Theorem, Free Speech, and Correlated Votes", *American Journal of Political Science* 36 (3), 617 - 634.

[12] Milgrom, P. and R. Weber (1982) "A Theory of Auctions and Competitive Bidding", *Econometrica* 50 (5), 1089 - 1122.

[13] Nisbett, R. (2003) *The Geography of Thought: How Asians and Westerners Think Differently...and Why* Free Press, New York.

[14] Page, S. (2007) *The Difference: How the Power of Diversity Creates Better Firms, Schools, Groups, and Societies* Princeton University Press.

[15] Pearl, Judea (2000) *Causality* New York: Oxford University Press.

[16] Stinchecombe, A. (1990) *Information and Organizations* California Series on Social Choice and Political Economy I University of California Press.

[17] Tesfatsion, L. (1997) "How Economists Can Get A-Life" in *The Economy as a Complex Evolving System II* W. Brian Arthur, Steven Durlauf, and David Lane eds. pp 533–565. Addison Wesley, Reading, MA.

[18] Valiant, L.G. (1984) "A Theory of the Learnable" *Communications of the ACM*, 17(11),1134-1142.

[19] Wellman, MP, A Greenwald, P. Stone, and PR Wurman (2003) "The 2001 Trading Agent Competition" *Electronic Markets* 13(1).