Interpreted and Generated Signals

Lu Hong[*] and Scott Page[†]

May 12, 2006

Abstract

In economics and other social sciences, incomplete information about variables of interest is often modeled as signals. The relationship between the signals and the relevant variables is captured by a joint probability distribution function. The assumptions we make about the distribution function, therefore, implicitly restrict the set of environments that could produce these signals. In this paper, we construct and contrast two frameworks for two different signal producing environments: *generated signals* and *interpreted signals*. Generated signals are distortions of true values or outputs of some process. Interpreted signals are predictions based on inputs or attributes. By making this distinction, we connect characteristics of the signal producing environment to assumptions on the joint distribution function. We show that some common assumptions about the joint distribution are only consistent with interpreted signals while others are only consistent with generated signals. Our findings therefore limit the contexts in which some existing results apply.

Introduction

When making economic or political decisions, people rarely have the benefit of full information. Instead, they rely on partial or distorted information or on crude predictive models. Typically, this information and these predictions are modeled as signals and these signals can be distinct: each person receives his or her own signal. The statistical relationships between the signals and between the signals and the variable of interest are characterized by joint probability distribution functions. Assumptions about these joint probability distributions, such as independence or independence conditional on some state or variable, implicitly constrain the set of environments that could produce the signals. Signals that equal the true value plus a noise term differ in kind from signals based on idiosyncratic predictive models.

[*]Department of Finance, Loyola University Chicago

[†]Departments of Political Science and Economics, University of Michigan

In this paper, we formally construct and contrast two frameworks for producing signals. We demonstrate that even though both types of the signals can be modeled by joint probability distributions, the natural assumptions applicable to the two types of signals differ. Towards the end of this section, we summarize six differences. These differences speak to the generality and relevance of existing theoretical results in incomplete information environments. Analyses of economic and political institutions require assumptions about the statistical properties of signals: the economics and political science literature contains primarily conditional claims dependent upon these statistical properties. This paper helps to connect those results to specific environments. More importantly, the paper lays the foundation for understanding the design and the performance of markets and political mechanisms under two different types of incomplete information environments.

In the first framework, signals are generated by some process and passively received by the economic or political agents. We call these *generated signals*. Generated signals differ because people draw different samples or experience idiosyncratic shocks or distortions. In the second, the agents produce signals by applying predictive models based on data. We call these *interpreted signals*. Interpreted signals differ because people differ in how they categorize or classify objects, events, or data.

A specific example helps to clarify our distinction between the two types of signals. Imagine two executives assigned the task of choosing between catering companies for the annual company picnic. These executives might sample from the caterers' menus. These taste tests would be generated signals: they are outputs of some process that depends upon the caterers' qualities. We might even model these signals as equal to the true quality of the caterer plus some error term. Alternatively, the executives might look at the caterers' menus and at their employees' credentials. The executives might plug this information into a model, be it formal or informal, to come up with some prediction of the caterers' qualities. These signals would be based on attributes of catering companies, not on something they produce. These would be interpreted signals.

Our generated signal framework should be familiar to most readers; whereas, our interpreted signal framework may not be. Though not a core assumption in economics, the notion that people see the world differently and that we interpret situations differently, has long been a standard assumption of psychology (Fryer and Jackson 2003).[1] People categorize and make predictions based on those categorizations. Our framework maps these diverse categorizations into a standard signaling framework.[2] Our interpreted signal framework resembles models of causal inference (Pearl 2000), fact free learning (Aragones, Gilboa, Postlewaite, and Schmeidler 2005), and complexity (Al-Najjar, Casadesus-Masanell and Ozdenoren 2003). It most closely resembles PAC learning theory (Valiant 1984) and classification theory (Barwise and Seligman 1997).

[1] Organizations do something similar (Stinchecombe 1990).

[2] Among psychologists, however, substantial disagreement exists about the extent to which those differences depend upon our life experiences, natural abilities, culture and incentives (Nisbett 2003).

We differ from PAC learning and classification models in two respects. First, while they consider an individual agent who formulates an optimal classification rule, we consider multiple agents each with his/her own classification rule. We allow our agents' rules to be chosen in strategic contexts. Thus, they need not be the optimal rule for a single decision maker. Instead, incentives would seem to lead strategic agents to choose different classification rules. Second, the natural extension of these PAC models would permit diversity over attributes only. One analyst might predict crime based on age, another might predict based on gender. We go further, we also allow agents to have different perspectives (Hong and Page 2001). By this we mean that people can differ in how they represent the set of the possible alternatives. One person may identify a point in the plane using Cartesian coordinates. Another person may use polar coordinates.

We first restrict our attention to independent interpretations – interpretations that consider distinct attributes. Considering independent interpretations simplifies the analysis, allowing us to characterize properties of interpreted signals. Though restrictive, independent interpretations are the obvious place to begin an investigation into the nature of interpreted signals, especially given that agents often have incentives to look at different attributes. Later, we relax that assumption and allow overlap on some attributes. We also consider the special case where the variable of interest takes on binary values. We establish several surprising results. For example, we show that some common signal models are not consistent with the interpreted signal framework. Consider the following standard set of assumptions in a model with two states G (good) and B (bad) and two signals: (i) the states are equally likely ($p(G) = p(B) = 1/2$) (ii) the two signals s_1 and s_2 predict the two states with equal likelihood ($\text{Prob}(s_i = G) = 1/2$) and (iii) the signals are independently correct (i.e. knowing that one signal is correct tells nothing about the probability that the other is correct) and (iv) the signals are informative, i.e. they are correct more than half of the time. In the context of generated signals, these assumptions are perfectly reasonable, however the results in this paper imply that these assumptions are *not* compatible with independent interpreted signals. The above model is pervasive in the economics and political science literature. Models of information aggregation in common value auctions and in majority or proportional voting and models of information cascades and herd behavior all utilize this model. It could be argued that, in many of the environments considered in those papers, signals should be more appropriately modeled as interpreted rather than generated.

In what follows, we demonstrate six important differences between the two types of signals. Some differences are straightforward and are implied by the definitions of the two types of signals. Other differences are novel and are proven in the paper. First, the appropriate independent assumption for generated signals is if the signals are independent conditional on the value. Independent interpreted signals, in contrast, cannot be independent conditional on every possible value. Second, independent generated signals are independently correct – the correctness of one signal does not depend on the correctness of any other. But, independent interpreted signals cannot be independently correct. In fact, independent interpreted signals must be negatively correlated in their correctness both unconditionally and conditionally on at least one value. If we are interested in aggregating information, we like this result, but if we are interested in the revenue generated by an auction, whether we see this as good news or bad news depends on whether signals are negatively correlated in the high value state or the low value state.

Third, generated signals can be symmetric. If signals are outputs of some process they may well be symmetrically distributed. In contrast, interpreted signals would be symmetric only in special cases. We would not expect the inputs to a process to have symmetric effects on the value or the quality of the process. Fourth, the correlation of generated signals with the value of the relevant variable alone determines the signals' collective accuracy. The collective accuracy of interpreted signals depends on a function that connects the attributes to the value. The more nonlinear and interacting terms, the less accurate the signals both individually and collectively.[3]

Fifth, in their standard construction, generated signals have a probabilistic relationship to the relevant variable. No number of generated signals, independent or non independent, pins down the true value with certainty. The same may not be said of independent interpreted signals. If all of the attributes are known, then the value may be known with certainty. If a geologist knows all of the attributes of an oil deposit and has the right predictive model, she could know the value of the oil deposit. Finally, any number of independent generated signals often can be produced. In contrast, the number of independent interpreted signals is bounded by the number of combinations of attributes, and the number of distinct interpreted signals of any type is bounded by the power set of the set of attributes, provided that we stick to a common set of attributes – a point we take up in a moment. This last difference would seem to imply that any model that assumes a large number of signals implicitly assumes generated signals and rules out interpreted signals. We discuss the implications of this intuition in more detail at the end of the paper.

As these six differences make clear, the statistical properties of generated and interpreted signals differ substantially. These differences call into question the generality and relevance of existing results. Even a cursory 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. The reason for particular assumptions often seems to be based on tractability and not on descriptive realism. For example, in that part of the auction literature that relates to the derivation of optimal bidding strategies, the assumptions align with our interpreted signals 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

[3]By collective accuracy, we mean here how accurate a prediction could be made knowing the values of all signals.

and election models (Feddersen and Pesendorfer 1997), makes assumptions that are not consistent with interpreted signals. Yet, the stories used to motivate those papers rely on jury members and voters who use predictive models.

Clearly the distinction between the two types of signals allows us to make better connections between theory and reality. It also allows us to unpack assumptions more finely. In the penultimate section, we show how interpreted and generated signals provide alternative lenses through which to view the affiliation assumption (Milgrom and Weber 1982). Many believe affiliation to be an extremely strong, and perhaps unrealistic, assumption. We show that assuming interpreted signals weakens the affiliation assumption substantially. The affiliation assumption becomes even weaker if we allow for endogenous interpretations.

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 formally present our framework for interpreted signals. Various notions of independence are discussed. It is also in this section that we establish a surprising result that when people have independent interpretations, they can be thought of as having the same representation of reality but focusing on different attributes. In Sections 4 and 5, we prove the main statistical results, demonstrating the different properties of generated and interpreted signals. In Section 6, we discuss the possibility of overlapping interpretations and connect features of signals to the complexity of the outcome function. In Section 7, we examine the affiliation assumption in common value auction models under the two types of signals. Section 8 concludes.

Generated and Interpreted Signals: Examples

We begin by constructing examples of generated and interpreted signals. These examples highlight the statistical differences between them. They also help to solidify the distinction between the two types of signals. We describe these examples in the familiar context of two investors who get signals about the quality of an investment. To frame the discussion, we consider the investment to be a restaurant. The restaurant can either be of high value ($V = 1$) or of low value ($V = 0$). We first describe generated signals, as they are both less complicated and more familiar and then turn to interpreted signals. In fact, many readers will find our discussion of generated signals to be trivial. Nevertheless, it is necessary to provide a point of comparison for interpreted signals.

Generated Signals

Generated signals are based on outputs. Often, we can think of them as blurry glimpses or distortions of the value V. Here, we consider the case of investors in a restaurant. As signals of the restaurant's value, we assume that each of the two investors eats a sample meal prepared by the chef. We assume that these signals also take on binary values. If the restaurant will be of high value, i.e. if $V = 1$ then each investor independently gets a good meal with probability p, where $p > \frac{1}{2}$. For notational convenience, hereafter, we denote the high value restaurant with a G and the low value restaurant as B to signify good and bad investments. We assume G and B to be equally likely. We denote the meals that the investors sample with the lower case letters g and b to signify good and bad meals. These signals are *generated*. 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. We further assume that these generated signals are conditionally independent given the restaurant's quality, i.e. each is an independent draw. Given these assumptions, 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	$\frac{1}{9}$	$\frac{2}{9}$
g	$\frac{2}{9}$	$\frac{4}{9}$

Generated Signals Conditional on B

s_1/s_2	b	g
b	$\frac{4}{9}$	$\frac{2}{9}$
g	$\frac{2}{9}$	$\frac{1}{9}$

We can also compute the expected value of the restaurant conditional on the signals of the two actors.

Expected Quality Conditional on Generated Signals

s_1/s_2	b	g
b	$\frac{1}{5}$	$\frac{1}{2}$
g	$\frac{1}{2}$	$\frac{4}{5}$

Interpreted Signals

We now turn to interpreted signals. The precise interpreted signal framework will be developed in the next section. Here, we give an informal description. Interpreted signals are predictions based on *interpretations*. Interpretations are defined as partitions (or categorizations) of reality. An example can be that an agent partitions reality along a particular attribute, one of many attributes that represent reality.

To construct an analog of our restaurant example but with interpreted signals we consider investors who look at attributes of the restaurant. These attributes are the restaurant's *location* and its *prices*. We assume that the investors see these attributes perfectly. To complete the example, we must assume a *value function*, V, that maps the restaurant's *attributes*, into a probability that the restaurant is of high quality. Here, we assume that the location ℓ and the prices \$ can be either good 1 or bad 0,each is equally like. We also assume that these attributes are statistically independent. This creates four equally likely pairs of attributes. We assume the following functional form for V[4]

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

We assume that the first investor looks only at the location and the second looks only at the prices. Based on the attributes that they see, these investors construct *predictive models*. In this simple example, the predictive models take a common form: if the attribute is good, then the restaurant is more likely to be of high quality. In our more general model, predictions can be based on multiple attributes and therefore are not identical to the attributes themselves.

We can write the value function as follows:

Quality Conditional on Attributes

$\ell/\$$	0	1
0	$\frac{1}{3}$	$\frac{1}{3}$
1	$\frac{1}{3}$	1

We can then calculate the joint probability distribution of the two attributes conditional on the restaurant's quality

Attributes Conditional on G

$\ell/\$$	0	1
0	$\frac{1}{6}$	$\frac{1}{6}$
1	$\frac{1}{6}$	$\frac{1}{2}$

[4]The value function above as a function of the two attributes is probablistic. This is not important as we can introduce a third independent attribute h (hidden or not observable) which can take three values, say 0, $\frac{1}{2}$, and 1, with equal probability. The value function can then be written as a deterministic function as follows:

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

Attributes Conditional on B

$\ell/\$$	0	1
0	$\frac{1}{3}$	$\frac{1}{3}$
1	$\frac{1}{3}$	0

Note that 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 $\frac{2}{3}$. What this example shows is that it is possible to construct either generated or interpreted signals to describe some situation. However, the statistical properties of these signals differ, as we shall see in a moment.

Statistical Differences Between The Signals

If we compare these two examples, we see several differences. First, while it is natural to assume that the generated signals are independent conditional on the value of the restaurant, the nature assumption for the interpreted signals is that they are independent (not conditionally). Independent interpreted signals here are not conditionally independent. Conditional on the restaurant being of high value, if the location is good, this *reduces* the likelihood that the prices are also good. Second, the generated signals are correct independent of one another. That is not true of the interpreted signals. If one interpreted signal is correct, the other is *less* likely to be correct. Third, the correlation between the signals and values in the second example depends on the value function, whereas the correlation in the first example 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. In this example, the function that maps attributes to values is symmetric with respect to the attributes, that need not generally be the case as we shall see.

Our two examples make a clean distinction between generated and interpreted signals. That distinction can often be subtle. Consider a university with a large number of faculty whose quality q in $\{0, 1\}$ is a function of its faculty's abilities. We might assume that each faculty member i has an ability a_i in the set $\{0, 1\}$ and that these abilities are independently drawn from some distribution F. We might also assume that this university produces students at regular time intervals and that each student j has a value x_j in $[0, 1]$. To evaluate the quality of the university, people could look at either graduating students' values, the x_j's, or the faculty's abilities, the a_i's. In this example, the graduating students' 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 outline.

The Interpreted Signal Framework

In this section, we describe the interpreted signal framework and provide definitions of the various forms of independence and establish relationships between them. Along the way, we uncover a surprising result: seeing the world independently, in a formal mathematical sense, implies seeing the world in the same way, but looking at different parts of it. Stated formally, interpretations can only be independent if they rely on non-overlapping sets of attributes from a common perspective.

We assume that *reality*, Ω, consists of a finite collection 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. The mapping from events to outcomes can be formalized as an outcome function $\widetilde{O} : \Omega \to S$. This mapping is deterministic. The framework can be extended to include probabilistic mappings, and the results of this paper hold for probabilistic mappings as well. The problems we consider are equivalent to classification problems in which a person has to place the N objects into $|S|$ bins representing possible outcome values. This problem is related to the problem of selecting regressors (Aragones, et al 2005).

We also assume that people can have different and incomplete representations of reality. This extends Hong and Page (2001) in which representations were bijective maps of reality. In the current framework, each person partitions reality into non-overlapping sets. We denote the partition of person i, Π^i, to be the sets $\{\pi_1^i, \pi_2^i, \ldots, \pi_{n_i}^i\}$, where n_i is the number of sets in person i's partition. Π^i is person i's representation of reality and this representation is incomplete as long as not all sets in the partition are singletons. When a person 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 reality, a person makes predictions about what the outcome of an event will be. There are many ways in which predictions might be generated. For example, **experience generated predictions** assume that the most probable outcome arises conditional on the set in her interpretation. Suppose, for instance, that a person interprets the event ω_1 as the set $\{\omega_1, \omega_2, \omega_3\}$ in the person's interpretation. Further suppose that ω_1 and ω_2 map to outcome G, and ω_3 maps to outcome B and that $P(\omega_1) + P(\omega_2) > P(\omega_3)$. A person with experience would then predict outcome G for all three events. At this point though, we do not make it specific 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.

These predictions are **interpreted signals**.[5] The following example illustrates the main components of the interpreted signal framework:

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

Example 1 *In this example, reality consists of six events,* $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$*. 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 a person's interpretation of reality where* $\pi_1 = \{\omega_1, \omega_2, \omega_4\}$ *and* $\pi_2 = \{\omega_3, \omega_5, \omega_6\}$*. If this person 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 person 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 of World	Signal g	Signal b
G	$\frac{2}{3}$	$\frac{1}{3}$
B	$\frac{1}{3}$	$\frac{2}{3}$

We can now define the *correctness* of predictions. An outcome function $\widetilde{O} : \Omega \to S$ and a prediction $\widetilde{\phi} : \Omega \to S$ induce another random variable, $\widetilde{\delta}$, called **correctness of predictions**, which is 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.

Definitions and Relationships of Types of Independence

We now define various types of independence. These definitions of independence are standard, but they recur throughout our subsequent analysis and discussion so we reiterate them here.

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* $j \in \{1, ..., n_2\}$

Saying that two people's interpretations are independent means that knowing how one person interprets an event provides no information about how another person interprets that same event. Similarly, saying two people's predictions are independent means that knowing one person's prediction about the outcome of an event provides no information about the other person's outcome prediction of the same event.

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

If two predictions are independently correct, then knowing that one person'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.*

The definitions above are given for only two interpretations. They can be extended trivially to include any finite number of interpretations.

We now present as observations some basic relationships between the various types of independence we have defined. We include them here to clarify and reinforce the differences between the types of independence as much of the subsequent discussion will refer to them. Our first observation makes an obvious point but one that bears mentioning nonetheless. If two people have independent interpretations, then the predictions that come from those interpretations must be independent.

Observation 1 *Independent interpretations imply independent predictions.*

Pf. Since each person's prediction is measurable w.r.t. her interpretation, the claim follows.

Our next observations states that predictions can be independent without people having independent interpretations. This should come as no big surprise, but it helps to drive a conceptual wedge between the two types of independence.

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

Pf. Due to the measurability requirement, a person'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 person 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 person's prediction is independent of the prediction made by the person in Example 1, even though the two interpretations are not independent.

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 correct independently.

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

Pf. 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, there are 16 events, each 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 person and the column person respectively. Clearly, the two predictions are independent. However, they are not independently correct. The joint probability of both people making correct predictions is $\frac{1}{2}$ while the multiplication of the probabilities of each person making correct predictions equals $\frac{3}{4} \times \frac{3}{4} = \frac{9}{16}$. They are not equal.

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

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

Thus, the only causal relationship that exists between independent interpretations, predictions, and the correctness of those predictions is that independent interpretations must result in independent predictions.

The Structure of Independent Interpretations

We now explore the mathematical structure of independent interpretations. To motivate our result, let's first consider the following trivial observation: if the event space Ω is represented by an attribute model and agents observe different attributes, then they have independent interpretations. We now consider the converse. We show that independent interpretations create a rectangular structure on the event space. We show that when two interpretations are independent, 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 is along the x attribute and the other is along the y attribute. In other words, independent interpretations can be viewed as interpretations along different attributes of *the same perspective* (Hong and Page 2001). This result is stated in the following claim. The proof of this result as well as proofs of all other results in the paper is contained in the appendix.

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 ith attribute. Thus $N = \prod_{h=1}^{n} a_h$ for some larger-than-1 integers a_h, $h = 1, ..., n$*

This result is at first surprising. It shows that whatever people's representations of the world are, if they have independent interpretations, they can be thought of as having the same representation of the world but focusing on different attributes.[6]

Intuitively, Claim **??** implies a bound on the number of independent interpretations. It cannot exceed the number of primes in the factorization of N. Therefore, there cannot be very many independent interpretations of a finite set of events.

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.

Correlation and Correctness of Predictions

We now present some general 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,

[6]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.,

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

these inconsistency results also apply to independent interpretations. The results in this section and the next two sections are established for the simplest case where there are only two possible outcomes G and B. Since our main goal is to demonstrate the difference between interpreted signals and the generated signals, the restriction to binary outcomes does not reduce the power of our arguments. For the same reason, we assume that people's predictions have identical probability distributions even though they may be based on different interpretations of reality.[7] The accuracy of predictions is assumed to be the same as well.

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. Recall that by assumption these are the same for both people. $P(g,g)$ 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 people. Finally $P(c,c)$, $P(i,i)$, $P(c,i)$ and $P(i,c)$ denote joint probabilities of both correct, both incorrect, person 1 correct but person 2 incorrect and person 1 incorrect but person 2 correct respectively. We further assume the following symmetry between two people: $P(g,b) = P(b,g)$ and $P(c,i) = P(i,c)$.

Reasonable and Informative Predictions

Given an interpretation, there is no guarantee that the person using that interpretation makes the best possible predictions. To impose some degree of experience we assume that predictions are either *reasonable* or *informative.*

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 is always reasonable. Further if there is no tie-breaking for at least one prediction, then an experience generated prediction is also informative. Consider the following example:

An Informative Prediction

[7] We have derived a set of results that do not assume that people's predictions have identical distributions. For logical clarity, we do not include them in the paper. In our interpretation framework, there is no reason to expect people's predictions to have identical probability distributions.

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

Assume that each outcome is equally likely. Clearly, $P(c) = \frac{2}{3}$. Therefore, this prediction is informative. However, conditional on outcome B, the probability of making correct prediction is $P(c \mid B) = P(b \mid B) = \frac{1}{3}$. Thus, an informative prediction need not predict correctly even half of the time conditional on every state. This leads to the following observation.

Observation 5 *An informative prediction need not be reasonable conditional on some states.*

This observation, like several of the earlier observations, will be obvious to some. We call attention to it so as to make clear that assuming a prediction to be reasonable or informative is less restrictive than requiring that the prediction makes correct predictions at least half of the time conditional on each possible outcome. This suggests that the assumption that predictions are reasonable conditional on every state may often be unrealistic. In predicting rare events, people may not be capable of much accuracy conditional on the state.

Predictions and Correctness of Predictions

Earlier in the paper, we established that independent predictions do not imply that predictions are independently correct. The example from the proof of Observation **??** showed negative correlation in the correctness of predictions from independent predictions. In this subsection, we further investigate the relationship between the correlation of predictions and the correlation of their correctness. The claim below shows a surprising relation between the informativeness, the correlation of predictions and the correlation of their correctness in a typical model where good and bad outcomes are predicted with equal probabilities. This relation is implied by the two lemmas which 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, predictions cannot be independently correct when they are highly accurate - the correctness of predictions must be negatively correlated. In what follows, without loss of generality, we designate g to be the more frequently picked prediction, i.e., $P(g) \geq \frac{1}{2}$.

The first Lemma below states that if predictions are independent, then whether the correctness of predictions is positively correlated, negatively correlated, or independent, depends entirely on whether the probability of the more frequently picked prediction is larger, smaller, or equal to the probability of being correct.

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 people often predict good outcomes at the same time whether or not the prediction is correct. So the correctness of their predictions is positively correlated.

Next, we reverse the assumption. If we require that the predictions be independently correct, we can show that the predictions themselves are independent or negatively or positively correlated depending again upon 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* $P(c) > [=, <]\ P(g)$.

An immediate implication of Lemma **??** is the following intriguing result.

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

An alternative way to state this claim is the following corollary.

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.

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 is negatively correlated.*

This claim and its corollaries reveal a fundamental conflict between an assumption that people see the world independently (independent interpretations) and an assumption that the interpretation based predictions they make are independently correct. That is, if people interpret the world independently, then their predictions are unlikely to be independently correct. For example, the assumptions of the original Condorcet Jury Theorem (See Ladha 1992 for a formal statement of the theorem) cannot be satisfied if jury members rely on independent interpretations. Thus, making statistical assumptions about predictions in our interpretation framework requires extra care to avoid logical inconsistency.

Interpreted vs Generated Signals

The interpretation based predictions, what we call interpreted signals, differ from generated signals in important ways. To show this, 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 situation, 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 themselves. Independent interpretations can never lead to predictions (or signals) that satisfy the standard assumption of conditional independence.

As before, we assume two possible outcomes, G and B,and two agents. $P(G)$ and $P(B)$ denote their common priors. For simplicity and in accordance with the previous section, we assume that probability distributions of agents' predictions conditional on the true outcome are identical . 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,

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

and

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

Then our model can 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.

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

and

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

In a typical binary signal model, signals are 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

$$p + q > 1.$$

For an interpreted signal, if predictions are experience generated and are informative, then the prediction/signal satisfies the SMLRP. This is easily established when we notice that by definition, predictions are experience generated and informative if $P(G \mid g) \geq \frac{1}{2}$ and $P(B \mid b) \geq \frac{1}{2}$ with at least one inequality holding strictly.[8]

We can now relate independent interpretations and predictions to signals that are independent conditional on the state.[9] The first claim states that informative and experience generate predictions – therefore, reasonable predictions – that satisfy independence conditional on the state must be positively correlated in their predictions unconditionally. 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. This, as we shall soon show, implies an inconsistency with independent interpretations.

Claim 3 *Experience generated and informative predictions that are independent conditional on outcomes must be positively correlated unconditionally.*

We can see this in an example. Suppose that good and bad outcomes are equally likely and that conditional on the state each of two people predicts correctly with probability $\frac{2}{3}$. To make this example as simple as possible, suppose that there are nine good outcomes and nine bad outcomes and that each is equally likely. For the predictions to be independent conditional on the state, the two people 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 is true for the bad outcomes. We can represent this in a table.

Conditionally Independent Predictions

Person 1 /Person 2	*PredictsG*	*PredictsB*
PredictsG	$4G$ and $1B$	$2G$ and $2B$
PredictsB	$2G$ and $2B$	$1G$ and $4B$

[8] Similarly, if predictions are experience generated, then $p + q \geq 1$ and thus *MLRP* holds. The fact that the *MLRP* holds is an artifact of our two outcome assumption. Examples can be easily constructed where there are more than two outcomes, predictions are experience generated, but the *MLRP* does not hold.

[9] Even though the claims and corollaries that follow are stated for experience generated and informative predictions, they hold more generally. By examining the proofs in the appendix, one can be convinced that they hold as long as predictions satisfy the SMLRP.

In this example, the predictions are not independent. If Person 1 predicts G then Person 2 predicts G with probability $\frac{5}{9}$. If Person 1 predicts B then Person 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 **??**.

Claim 4 *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.*

An immediate implication of either of the above two claims is summarized in the following claim.

Claim 5 *Conditional independence of predictions is inconsistent with informative and experience generated independent predictions.*

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

Corollary 4 *An assumption of conditional independence of informative, experience generated predictions is inconsistent with independent interpretations.*

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

Overlapping Interpretations

Earlier we proved that if people have independent interpretations, then these interpretations can be thought of as being derived from disjoint subsets of attributes that define reality (Claim **??**). In many instances, people may overlap in the attributes that they consider. Overlapping interpretations can lead to correlated predictions. If profits are a key determinant in a firm's value, then two investors who both consider profits when making their predictions will likely make predictions that are unconditionally positively correlated. However, if we take the common attributes into account, we are left with independent interpretations. Therefore, all of our results for independent interpretations have relevance to cases in which interpretations overlap on some attributes.

In this section, we provide an informal and intuitive discussion about what happens if interpretations overlaps. We begin with an example that demonstrates why many of the results for independent interpretations apply to interpretations that overlap. We then relate the properties of overlapping interpretations to the complexity of the outcome function. In this 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. There are four outcomes in each box 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.[10]

We next assume that the column player also considers the second attribute. This generates the following rectangular representation. Notice that some of the cells are now empty. This occurs 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

[10] $\frac{19*23}{32*32}$ is approximately $\frac{13.6}{32}$.

than the product that each is correct.[11] 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.

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

By implication, if people are aware of what attributes are commonly considered, then the correctness of their predictions can still be negatively correlated in light of that information. In fact, 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.[12]

Unconditional on the values of the common attributes, the correctness of predictions can of course be positively correlated. This would appear to be more likely the greater the predictive power of the common attributes. Common attributes of high predictive power will 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. This creates 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 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. It follows that with the new attribute, the column interpretation is now correct in $(x_G - x_B)$ more instances. This 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 $(x_G^{rG} - x_B^{rB})$. Despite the intuition that this number should also be positive and that it should be smaller than $(x_G - x_B)$, as we observe here, it can in fact be negative and it can be positive and larger than $(x_G - x_B)$.

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

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

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.*

Here is why. Recall that the change in the number of cases that both are correct equals $x_G^{rG} - x_B^{rB}$. First, we show that it is possible that both are correct less often. This can happen so long as $x_G^{rG} < x_B^{rB}$. This only requires that in the set X, the row interpretation predicts bad outcomes correctly more often than it predicts good outcomes correctly. This 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. The row interpretation only predicts bad outcomes in X.

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. That is what we might often expect. 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. To show this, we assume that the outcome function is linear and places equal weight on each attribute. Specifically, we assume that there are five attributes and that the outcome of the function equals one if three or more of the attributes have value one. Suppose now that we have two interpretations, each of which looks at four attributes. Any two interpretations must overlap on at least three of the attributes. It is a simple exercise to show that the predictions of the agents are, on average, positively correlated. This is because the function is not complex.

Next, we consider a more complicated function defined over the 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 when with more complex value functions, overlapping interpretations can be consistent with independent signals.

The upshot of this is that if we assume that there is some overlap in the attributes that people consider in their predictive models, then we need not assume that the signals are positively correlated. If the outcome function is not complex, then they will be. But with more complex outcome functions, no such assumption need hold.

Affiliation Assumption and Signal Type

We have established that the statistical properties of generated and interpreted signals differ. These differences matter for how we view the real world applicability of our models. In this section, we use models of common value auction as an example to illustrate this point. We show how by distinguishing between these two types of signals, we can reach a deeper understanding about the typical assumptions (especially the affiliation assumption) in the common value auction model. First, we provide some background. The literature devoted to characterizing optimal bidding strategies and the efficiency of auction mechanisms admits various assumptions about values. Some of that literature assumes private values. Another part considers common values. The common value framework is the relevant setting for thinking about generated and interpreted signals as they concern a single object with an unknown value. In a seminal paper, Milgrom and Weber (MW 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.[13]

As a special case of their general model, MW describes a common value auction as one in which each agent obtains a private signal about this common value. They then 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 must also satisfy the affiliation assumption. Nevertheless, many view the affiliation assumption as too restrictive. Using a simple example, we argue below that if signals are attributes of the common value (interpreted signals), then the affiliation assumption may not be as restrictive as some think.

Consider the following example of a common value auction for a single object with two bidders. 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.

[13]See Milgrom and Webber (1982) for a formal definition of 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. 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. The reason for our discussion is that within the context of this example, we can show that if we simply change the way we describe the model by treating the signals as interpreted signals, then MW's assumptions are satisfied. To show this, we notice that relationship between signals and the common value given by the joint probability distribution of signals and the value in this example is equivalent to the following description of the model.

The common value of the object is a function of two attributes x_1 and x_2 :

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

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

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

In other words, instead of thinking of each bidder's utility u_i as a trivial increasing function of the common value (which is what MW did), we can also think of it as an increasing function of the interpreted signals. Since x_1 and x_2 are independent, they are trivially affiliated. Therefore, when we write the variables as interpreted signals, all of the assumptions of Milgrom and Weber hold. So MW's model is applicable to this example after all.

This example belongs to the class of monotonic binary valued functions defined over more than one binary attribute. Our next claim states that no function in this class leads to an affiliated joint distribution over the signals and the value. That is, if we treat the signals as generated like MW did, the affiliation assumption will not be satisfied.

Claim 6 *Let $X = \{0,1\}^N$. If $f : X \to \{0,1\}$ 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.*

In contrast, every function f in this class combined with an affiliated distribution P (over the signals only) satisfies all conditions of the MW model. Thus, the affiliation assumption in the MW model appears to be both strong and weak depending on how we think of the 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 value 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.

However, lack of monotonicity need not imply that their results do not apply 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 value 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, in the interpreted signal framework, even if the value function contains large negative interaction terms, as long as the attributes involved in the large interaction terms 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.

This leads to a much larger point. In strategic context, we might expect that the choice of the attributes to be included in predictive models is endogenous. If the overlapping attributes contain the large negative externalities, then we might well expect the resulting interpreted signals to satisfy the affiliation assumption.

Discussion

In this paper, we have demonstrated important differences between generated and interpreted signals. As we mentioned in the introduction, these statistical differences have implications for how we view derivations of optimal strategies and designs of institutions in environments with incomplete information. These results suggest that at a minimum, it should be incumbent upon modelers to defend their assumptions based upon the specific context: are the signals generated, interpreted, or both?

The distinction between the two types of signals might also explain why reality often differs from our models and may 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. There is just too much information to process. 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 can only see what happens in a local region. This is close to what we model here as 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 mathematical models are not consistent with the information that people realize in the computational implementations of those models.

Finally, the interpretation framework permits more fine grained analysis of the link between complexity and uncertainty and allows for the modeling of endogenous signals. We comment first on the latter implication. If people want to predict correctly individually, this could lead to correlated signals as they might all learn to look at the same attributes. If people are concerned with collective performance, such as in the case of voting to aggregate information, people 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 people's predictions is often negatively correlated, so their information will aggregate rather well. However, the number of different attributes that can be considered depends upon the dimensionality of the problem. Therefore, large groups of people 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 an externality. The question of whether such simplification, where possible, is also incentive compatible, is an open one.

The link between complexity and uncertainty is also the focus of future work, but the relationship should be clear from our examples. Most of our examples involve mappings that are complex, i.e. nonlinear and including interaction terms. The complexity-uncertainty link is the focus of the aforementioned 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 notion of complexity.

Appendix: Proofs

Proof of Claim ??. 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.

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

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 ??. By Claim **??**, 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 ??, Lemma ?? and Claim ??. 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 people 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 **??**. 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 **??** is a special case of Lemma **??**. Lemma **??** can be similarly proved.

Proof of Claim ??. We need to show

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

Here, $P(g,g)$ denote the joint probability of both people 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 ??. 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 ??. 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] Pearl, Judea (2000) *Causality* New York: Oxford University Press.

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

[16] 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.

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

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