INDIVIDUAL ERROR, GROUP ERROR, AND THE VALUE OF INFORMATION

ITAI SHER

Northwestern University

Extended Abstract

Abstract. This paper studies the interaction of error and information both in a single-person setting and in an interactive setting. In contrast to Blackwell's Theorem, which says that more information is always good, the perspective of this paper is that while a lot of information is beneficial, a little information can be harmful. The main achievements of this paper are: (1) A characterization of the class of signals which always benefit a decision-maker in all decision problems. The analysis is carried out in a model which allows for the possibility that the decision-maker makes a mistake. (2) A demonstration that there are public signals within this class which can nevertheless reduce the utility of a team (i.e., a collection of agents with a common objective), as well as a characterization of the class of signals which always benefit a team in every team game. (3) A theorem that shows that in decision problems, beyond a certain threshold of precision, the value of information is increasing in the precision of signals, and which also provides a characterization of this threshold.

1. Introduction

Formalization of the notion that information is beneficial to a decision-maker goes back to Blackwell [7, 8] and Ramsey [17]. As in most economic models, the analysis presupposes that the decision-maker has not made any mistake at the time that he receives the information. On the other hand, consider a decision-maker who has made a mistake. This means that the decision-maker's view is in some sense distorted, when compared with reality. Perhaps the decision-maker assigns a probability that is too high to some event, or the decision maker assumes a causal relationship where none exists. Suppose that the decision-maker receives some *true* information. How will this information interact with his error? In general, it is possible that this true information will make the decision-maker's outlook more distorted rather than less distorted for the purpose of

I am very grateful to David Austen-Smith, Eddie Dekel, Marciano Siniscalchi, and Asher Wolinsky for helpful comments and advice. I would also like to thank Melissa Koenig, Peter Sher, and Tomasz Strzalecki.

making a good decision. However, if the information is precise enough, it will be beneficial. For example, if the decision-maker is simply given a complete and accurate description of all aspects of the world pertinent to his decision, this description will correct his initial error, and will then be beneficial to the decision-maker. In general, it is not necessary to give the decision-maker a *completely precise* description, but there is a threshold of informativeness,[1] beyond which information is always beneficial, partially for the same reason as in Blackwell's theorem, although matters are a bit more complex. In the model of this paper, signals are only partially ordered in terms of informativeness, so the minimally precise signals which constitute the threshold, beyond which all signals are guaranteed to be beneficial, may themselves be unordered in terms of informativeness.

The perspective of this paper is that a lot of information is always good, but a little bit of information can be bad. It is always better to be a complete expert than to be ignorant, but ignorance may be preferable to a state of knowledge between expertise and ignorance. I believe that this is actually a more realistic perspective than the traditional perspective of decision theory, according to which more information is always good. The reason is that in reality, decision-makers cannot hope to have a completely accurate model of all aspects of reality pertinent to a decision. Therefore, it is necessary for a decision-maker to construct a simplified model of reality, which may contain some distortion. Perhaps, beyond some threshold, enough distortion is eliminated for information to become beneficial, but in the meantime, information can make things worse.

When these considerations arise in a group setting, matters are considerably more complex. Within an organization, the level of distortion in the outlooks of different people may vary. Even if a single person is brought to the point where information would only be beneficial to him, were he to have complete control of the decision, others may still lag behind. Thus information which is sufficiently informative to eliminate distortion in one person's outlook may not be sufficiently informative to eliminate distortion in another person's outlook. Even if the two people have a common objective, such information may be harmful to them, due to the interaction of their outlooks. For example, the information could cause the two people to *disagree* more sharply; such a disagreement might cause the person with the more erroneous view to take actions which undercut the attempts of the more informed person to further the goals of the organization.

[1] In this paragraph, I use *precision* and *informativeness* interchangeably. For for a formal definition of precision, see Section 4

The purpose of this paper is to study the value of information in a setting in which agents make mistakes. I study this issue both in the context of single-person decision problems, and in the context of team games, i.e. games in which all agents have the same objective. Before proceeding to my main results, I would like to discuss the related literature.

1.1. **Related Literature.** In this paper, I model a mistake simply as a false belief. Thus from the perspective of modal logic, this corresponds to a situation in which agents' information is represented using a logic of belief (KD45) rather than a logic of knowledge (S5). Geanakoplos [12] characterized the deviations from standard partitional (S5) information structures for which more information is always better. However, Geanakoplos's characterization consists of comparisons between an agent who has *not* made a mistake and another agent who has more information but who has made a mistake; the mistakes are thought of as occurring through faulty information processing. Geanakopolos characterizes one agent as having more information than another if the *beliefs of* (rather than the information received by) the former are logically stronger than the beliefs of the latter at each state of the world. This notion of "more information" assumes away circumstances in which new information induces agents to update their beliefs *nonmonotonically.* In other words, this notion rules out the possibility that more information induces agents to *change their minds.* Morris [14] extends the analysis of Geanakoplos in several ways, and studies iterated updating. Morris does not explicitly study the properties of non-monotonic belief revision in this setting, and Morris's treatment of the decision-maker's welfare differs from that found in this paper.

The perspective of this paper is different. Following the classic work on *belief revision* by Alchourrón, Gärdenfors and Makinson [1] (henceforth AGM), I assume that agents update their beliefs by adding new information and closing under logical consequence when consistent, but by discarding some old information otherwise. There has been a growing literature that either applies AGM-type theories of belief revision to game theory or develops such theories with an eye to such applications (Battigalli and Siniscalchi [5, 6], Board [9, 10], Bonnano [11], Stalnaker [23, 24], Tallon et al. [25]). This paper continues along these lines. However, to my knowledge the value of information has not been studied in a setting with nonmonotonic updating of beliefs. I use the simple device of a conditional probability system (Renyi [18]) to implement AGM-type belief revision. The relationship between AGM and conditional probability systems has been studied by Spohn [22].

Previously, Myerson [15] and Battigalli and Siniscalchi [5, 6] used conditional probability systems to model belief revision in a game-theoretic setting.

This paper is also related to strands of the large literature on the value of information. It has been shown that Blackwell's theorem fails when the decision-maker is not an expected utility maximizer (Wakker [26], Safra and Sulganik [19], Schlee [20]). In contrast, in this paper, I model agents as expected utility maximizers. Here, Blackwell's theorem fails because agents have an incorrect view of the world. Thus unlike the papers just cited, an agent would never attempt to avoid information, even if it would harm him. This paper also has relevance to the value of information in games, which has been studied by several authors (Bassan et al. [3], Bassan et al. [4], Gossner [13], Neyman [16]). It is well known that in games, information may sometimes be harmful. This paper looks at team games, in which all players have the same objective. In such games, if error is absent, information will always benefit the team at the efficient equilibrium. However, in the presence of error, the interaction of information with disagreement can cause information to have harmful effects. Thus, this paper shows that disagreement, like conflict of interest, can render information harmful.

1.2. **Overview of Main Results.** The main achievements of this paper are:

(1) A characterization of the class of signals which always benefit a decision-maker in all decision problems. The analysis is carried out in a model which allows for the possibility that the decision-maker makes a mistake.
(2) A demonstration that there are public signals within this class which can nevertheless reduce the utility of a team, as well as a characterization of the class of signals which always benefit a team.
(3) A theorem that shows that in decision problems, beyond a certain threshold of precision, the value of information is increasing in precision of signals, and which also provides a characterization of this threshold.

The plan of the remainder of this abstract is as follows. Section 2 provides an example of a situation in which more information is harmful in order to give the reader perspective on the results of this paper. Section 3 discusses the formal model. Section 4 provides a measure of error which is used in the theorems. Section 5 discusses the results for decision problems, namely (1) and (3) above. Section 6 discusses results in (2) relating to teams. Section 7 concludes. Proofs have been omitted in the interests of length. All proofs are included in the full version of the paper.

2. Information Can Harm an Individual in the Presence of Error: An Example

To understand the results of this paper, it is useful to consider a situation in which information is harmful to a decision maker. Consider a decision-maker who must choose one of two actions a and b. If the decision maker has no information, a delivers the highest expected utility. Consider two events E and F, and let $\neg E$ be the complement of E. Suppose that conditional on the intersection $E \cap F$, b is the best action, but in any state outside of $E \cap F$, a is the best action. Moreover, conditional on E alone or F alone, a is the best action. Now suppose that the event $(\neg E) \cap F$ occurs, so that by assumption, a is the best action. Suppose further that the decision-maker makes a mistake, perhaps due to bad information from a trusted source or to mismeasurement, and believes that E has occurred. This will not harm the decision-maker, for he still correctly believes that a is the best action. Now suppose that the decision-maker gets the true information that F occurred. Then the belief $E \cap F$ will induce the action b, which will harm the decision-maker. Thus, the value of the *true* information that F occurred is negative, and corresponds to the loss generated by the choice of b rather than a. It is clear that this information could have a negative value *ex ante* either if the loss in this instance is large enough or if the belief in E is adopted at the *ex ante* stage before the realization that determines whether the event E occurs.[2] The second possibility could occur if the decision-maker stubbornly holds on to a theoretical or ideological assumption regardless of reality.

To summarize the example, whenever two conditions are necessary to warrant overturning the status quo, and the decision-maker mistakenly believes that the first condition obtains, true information that the second condition obtains will harm the decision-maker. The assumptions of the example are clearly easy to satisfy. To take just one example, any decision-maker who is overly optimistic may be harmed by information which would rationalize an action under the assumption that things will go well.

3. The Model

Assume that there is a finite set of agents I and a finite set of states Ω. Each agent i has a partition t_i over Ω. The cell of i's partition containing state ω is written as $t_i(\omega)$, which represents

[2]Notice that if the belief that E occurs is adopted *before* the realization that determines whether E occurred, then the value of information can be negative regardless of the magnitude of the loss. For instance, suppose that the decision-maker is informed that F occurs exactly if F occurs. Then whenever F occurs the decision maker will choose b. However, above it is assumed that conditional on F, a yields a higher expected utility than b.

his private information at ω. In addition to this partition, each agent has a conditional probability system p_i. As in Myerson [15], $p_i(\cdot|F)$ is assumed to be defined for every nonempty subset F of Ω.

Definition 3.1. A *conditional probability system* is a mapping $p(\cdot|\cdot)$ from $2^\Omega \times 2^\Omega \backslash \{\emptyset\}$ to $[0,1]$ such that for all events E, and for all nonempty events F and G:

(1) $p(\cdot|F)$ is a probability measure on Ω.

(2) $p(F|F) = 1$.

(3) If $E \subseteq F \subseteq G$, then $p(E|G) = p(E|F)p(F|G)$.

For any event E, i *believes* E at ω if $p_i(E|t_i(\omega)) = 1$. Belief revision conditional on partitions finer than t_i is determined by p_i. It is easy to check that this definition induces a model that simultaneously satisfies both KD45 and the AGM postulates for belief revision.[3] Moreover, it induces higher order beliefs about other agents of all orders at every state of the world and conditional on any information. This is a simple generalization of an Aumann structure [2], and resembles many other constructions in the literature. After receiving their private information $t_i(\omega)$, the agents receive a public signal $X(\omega)$, and update accordingly. Formally, I model a signal X as a partition on Ω, where $X(\omega)$ is the cell of X containing ω. After receiving the public signal X, agents' probabilistic beliefs are given by $p_i(\cdot|t_i(\omega) \cap X(\omega))$. Since X is a partition, it represents true information.

I make two assumptions about probabilistic beliefs:

Assumption 1. (**Quasi-Common Prior**) There is a probability measure p^* on Ω, such that for all agents i and all events E:

$$p_i(\cdot|E) = p^*(\cdot|\text{supp}(p_i(\cdot|E))).$$ [4]

This assumption is without loss of generality if there is only one agent. The justification for this assumption is that I want to focus on the effects of false propositional beliefs and belief revision rather than on probabilistic errors. Intuitively, this means that agents may disagree about which propositions are true and may have different belief revision policies, but all disagreements in probabilistic beliefs are due to these two sources of disagreement. Since this assumption only

[3] The AGM postulates are not satisfied for all propositions, but they are satisfied for all propositions that the agent may have occasion to come to know in this model. For the relation between conditional probability systems and the AGM postulates, see Spohn [22].

[4] For any probability measure p, supp p means the support of p.

has consequences in an interactive setting, I defer further discussion on its content and justification to section 6, in which team games are discussed. It is important to note that Assumption 1, in conjunction with the fact that $p_i(\cdot|E)$ is assumed to be defined for every event, implies that for all states ω, $p^*(\omega) > 0$.

Assumption 2. p^* is the objective probability measure on Ω.

This assumption specifies the way in which agents' probabilistic beliefs are mistaken. Like Assumption 1, this assumption restricts the analysis to situations in which the agents have false propositional beliefs, as opposed to more general situations in which agents assign the wrong relative probabilities to states. From my perspective, it is best to think of this assumption in terms of its ramifications for mistaken decisions. Given that an agent chooses the action which maximizes his expected utility given his beliefs, Assumption 2 implies that conditional on i's strongest belief (i.e., $\text{supp}(p_i(\cdot|t_i(\omega)))$), i always takes the optimal action; the problem is that this belief may be false. In the setting of decisions, Assumption 2 rationalizes the use of p^* for making definitive conclusions about welfare from an *ex ante* perspective. An alternative assumption with a more subjective flavor, which plays a similar role is presented in Sher [21]. Sher [21] makes the more general assumption that welfare is measured by $p^*(\cdot|E)$ for some event E, rather than by p^*.

It is possible to construct a simple story which generates Assumptions 1 and 2. In particular, suppose that each agent started out with correct probabilistic beliefs, and that his error resulted from listening to some trusted information sources, where different agents may have different information sources. These information sources can be ordered in terms of the logical strength of their information. If the agent makes an observation which shows that he was mistaken, he rejects all information from sources which are inconsistent with his observation, but keeps trusting the rest.[5] It can be shown that processes which are consistent with this story will generate exactly the conditional probability systems which are consistent with the assumptions of the model of this paper.

4. Measuring Error

In this section, I will briefly discuss the way in which error is to be measured, and characterize the signals that reduce error. *Agent i's error at ω* is the set of events E such that $\omega \notin E$ but i

[5]Since the information sources are ordered by logical strength, the set that he rejects is unique.

believes E. An agent's error can be measured either before or after the public signal X is sent. i's error at ω can then be thought of as the *false part* of i's theory. Formally, the expressions:

$$\begin{aligned}\text{Error}(i,\omega) &= \{E \in 2^{\Omega} : \omega \notin E, p_i(E|t_i(\omega)) = 1\} \\ \text{Error}_X(i,\omega) &= \{E \in 2^{\Omega} : \omega \notin E, p_i(E|t_i(\omega) \cap X(\omega)) = 1\}\end{aligned}$$

define the agent's error at ω, respectively, before and after the signal X is sent. These are local measures of error. A global measure can also be given. Define *i's error* as the mapping from ω to i's error at ω. Formally, *i's error* is defined as $(\text{Error}(i,\omega))_{\omega\in\Omega}$ before the signal, and $(\text{Error}_X(i,\omega))_{\omega\in\Omega}$ after the signal. Given this definition, it is possible to check whether a signal increases or decreases an agent's error at a state ω, according to whether the false part of i's theory grows or shrinks when measured by the inclusion relation. A signal X is said to *decrease i's error* if X (weakly) decreases i's error at every state (according to $\subseteq$). Notice that a signal strictly decreases i's error at ω if and only if the signal eliminates i's error at ω. This is because in order for a signal to decrease i's error at ω, it must be informative. But if it is informative and does not eliminate i's error, then the conjunction of any erroneous beliefs that remain and the information given by the signal constitutes a new false belief. Notice that this argument applies locally at a single state of the world, but globally, it is possible to reduce an agent's error without eliminating it.

It can be shown that the signals that reduce i's error are the signals X that must satisfy one of the following three conditions at each state ω:

(1) The agent's error is null at ω prior to receiving the signal.

(2) X is uninformative.

(3) The signal induces the agent to take the weakest belief which is consistent with the signal and the information $t_i(\omega)$; since both $X(\omega)$ and $t_i(\omega)$ are true, this means that in this information state, no true information would surprise the agent after updating on the basis of X.

A signal which satisfies one of these three conditions at every state of the world is called an *honest signal*. Clearly, whether a signal is honest, depends on the agent's private information and his belief revision rule as embodied in p_i.

In order to formalize the the notion of an honest signal, it is necessary to introduce a little notation. Define level $1(i)$ to be the support of $p_i(\cdot|\Omega)$. Let level $n+1(i)$ be the support of

$p_i(\cdot|\Omega \setminus \bigcup_{k=1}^{n} \text{level } k(i))$. If it is obvious which agent is being discussed, I simply write level n. The interpretation of these levels is that when presented with information that the true state is either in a higher level or in a lower level, the agent will believe that it is in a lower level. Thus, states on a lower level receive a higher epistemic priority. I assume that for all i and ω, $t_i(\omega)$ and level $1(i)$ have a nonempty intersection. For the purposes of the theorems proved in this paper, this is without loss of generality.

Definition 4.1. A signal X is *honest for i* if and only if:

$$\forall \omega, \text{level } 1(i) \cap t_i(\omega) \not\subseteq X(\omega) \Rightarrow \exists n, X(\omega) \cap t_i(\omega) \subseteq \text{level } n(i)$$

This provides a characterization of error-reducing signals:

Theorem 4.2. *X reduces i's error if and only if X is honest for i.*

This characterization is useful both for proving the theorems below, and for checking whether a signal will reduce an agent's error.

5. The Value of Information in Decision Problems

In the case of decision problems, reduction of error is necessary and sufficient to *guarantee* increase in welfare. Fix an information structure $(\Omega, (t_i, p_i)_{i \in I})$. Formally, a decision problem is a pair (u, Z), where Z is a set of actions with generic element z, and u is a state dependent utility function that maps $Z \times \Omega$ into the real numbers. I assume that Z is a compact set, and that $u(\cdot, \omega)$ is continuous in z for each ω. A *strategy* is a mapping $\mathbf{z}$ from Ω to Z, which is measurable with respect to the agent's information (t_i before the signal, or $t_i \wedge X$ after the signal). The agent chooses the strategy that maximizes his expected utility, given his beliefs.[6] The following theorem compares his utility in the case that he does not receive the signal X to his utility in the case that he does receive the signal X.

[6]Formally, let $\mathbf{z}^*$ be the strategy that maximizes i's expected utility before the signal, and $\mathbf{z}^{**}$ the strategy that maximizes i's expected utility after the signal. These strategies are characterized by the conditions:

$$\begin{aligned} \mathbf{z}^*(\omega) &\in \arg\max_{z \in Z} \sum_{\omega' \in t_i(\omega)} u(z, \omega') p_i(\omega' | t_i(\omega)) \\ \mathbf{z}^{**}(\omega) &\in \arg\max_{z \in Z} \sum_{\omega' \in t_i(\omega) \cap X(\omega)} u(z, \omega') p_i(\omega' | t_i(\omega) \cap X(\omega)) \end{aligned}$$

Theorem 5.1. *A signal (weakly) increases an agent's utility in all decision problems if and only if it reduces his error.*[7]

This contrasts with the common view, supported by Blackwell's theorem, that true information can only be beneficial to a decision-maker. In general, information which does not reduce an agent's error may introduce distortions in his outlook which are harmful. It follows from the theorem that even if an initial signal reduces the decision-maker's error, a second signal may still harm him. Nevertheless, there is some threshold beyond which additional information becomes beneficial. This threshold is *not* the point at which all error has been eliminated, but is prior to that point.

Assume that the decision-maker's private signal is uninformative so that at every state, $t_i(\omega) = \Omega$. A set E is said to have n levels if it intersects n sets of the form level m. The *bottom level* of E is the $E \cap$ level n, where n is the smallest number for which this intersection is nonempty. Then the threshold $\mathcal{T}$, beyond which signals become beneficial can be defined in terms of properties of the levels of the signals it contains.

A signal X is said to be *more precise* than a signal X', if at each state ω, $X(\omega) \subseteq X'(\omega)$. This is written as $X \leq X'$. Now I define a set $\mathcal{T}$ of signals, which have some interesting properties. $\mathcal{T}$ is referred to as the *threshold*. $\mathcal{T}$ is the set of all signals X which have the following properties:

(1) For each ω, $X(\omega)$ has at most two levels.

(2) If $X(\omega)$ has two levels, then the bottom level is a singleton.

(3) If $X(\omega)$ has one level, then if $X(\omega) \neq X(\omega')$, then $X(\omega) \cup X(\omega')$ has more than one level.

(4) If $X(\omega) = \{\omega\}$ and $\omega \in$ level n, then for all $n' > n$, and for all ω', it is not the case that $X(\omega') \subseteq$ level n'.

(5) If $X(\omega)$ has two levels n and n' such that $n < n'$, then for all ω', if $X(\omega')$ intersects level n', $X(\omega')$ has two levels.

Roughly speaking, the threshold contains signals which, among other things, induce the following properties in the agent's beliefs:

a: When the agent makes a mistake, his beliefs are as strong as possible.

b: The agent is not inclined to substitute one mistake for another logically unrelated mistake upon the receipt of further information.

[7] It is assumed that if a signal is uninformative at ω, then it does not cause i to change his action. This is relevant only when the optimal strategy (from i's perspective) is not unique.

c: When the agent has true beliefs, his beliefs are as weak as possible subject to satisfying **a** and **b**.

Define $\mathcal{U} = \{X : \exists X' \in \mathcal{T}, X \leq X'\}$. Thus, $\mathcal{U}$ is the class of signals which are at least as informative as some signal in $\mathcal{T}$.

Lemma 5.2. *$\mathcal{U}$ is the set of all signals that satisfy (1) and (2).*

The following theorem shows that $\mathcal{T}$ is the threshold beyond which additional information is beneficial.

Theorem 5.3. *$\mathcal{T}$ and $\mathcal{U}$ have the following properties:*

- **i:** *$\mathcal{T}$ is nonempty.*
- **ii:** *If $X, X' \in \mathcal{T}$, then X and X' are unordered in terms of precision.*
- **iii:** *The value of information is increasing on $\mathcal{U}$. That is to say, if $X, X' \in \mathcal{U}$ and X is more precise than X', then for any decision problem, the agent gets a higher utility if he receives X than if he receives X'.*
- **iv:** *For any signal X not in $\mathcal{U}$, there exists a decision problem and a signal X' more precise than X, such that X' would give the agent a lower utility than X. In other words, before getting to the threshold $\mathcal{T}$, there is always some decision problem for which more information is detrimental.*

6. The Value of Information in Team Games

It is well known that in interactive settings, Blackwell's theorem no longer holds. In general, information can have any effect. Bassan et al. [3] show that there are many possibilities. However, if agents do not make errors, and they have a common objective, then a version of Blackwell's theorem holds. A team game is a game in which players have a common objective. Formally, a *team game* is a set $(u, (Z_i)_{i \in I})$. Z_i is a compact set of actions for i, and $u : (\times_{i \in I} Z_i) \times \Omega \to R$ is a common state-dependent utility function; all agents want to maximize the same function u. for $z \in \times_{i \in I} Z_i$, it is assumed that $u(z, \omega)$ is continuous in z for each ω. A strategy for player i is a map $\mathbf{z}_i$ from Ω to Z_i, which is measurable with respect to t_i. A strategy profile $(\mathbf{z}_i^*)_{i \in I}$ is an *equilibrium* if for each player i and each state ω, $\mathbf{z}_i^*(\omega)$ is a best response to i's beliefs given $\mathbf{z}_{-i}^*$.[8] A *most efficient equilibrium* is an equilibrium $\mathbf{z}^*$ which maximizes $\sum_{\omega\in\Omega} u(\mathbf{z}(\omega), \omega)p^*(\omega)$ among all equilibria. Consider some information structure $q = (\Omega, (t_i, p_i)_{i\in I})$. Then let $\Gamma(q)$ be the set of all team games based on q for which an equilibrium exists.[9]

In this setting, it is possible to say more about the quasi-common prior assumption. In particular, consider team games in which the conditional probability systems of the agents are allowed to differ in arbitrary ways. Then, conditional on certainty about the state of the world, all agents will agree about which profile of actions are optimal. On the other hand, when agents agree about their probability one beliefs they may still disagree about the optimal action profile. On the other hand, the quasi-common prior assumption implies that whenever two agents assign the same events probability 1, then they will agree on the optimal action profile. Thus, the quasi-common prior assumption can be thought of a stronger form of common interests than merely sharing the same state dependent utility function $u(z, \omega)$. I have formalized this notion elsewhere [21].

I say that the agents are *infallible* if for all i, $p_i(\cdot|\Omega)$ has full support. In this case, each agent's conditional probability system is in fact a probability measure, and one recovers the standard model in which no one makes a mistake and there is a common prior. A generalization of the set-up of this paper is that rather than being sent a public signal X, each agent is sent a private signal X_i. The model of this paper is recovered if $X_i = X_j$ for all agents i and j. With this in mind, the following is a simple generalization of Blackwell's theorem:

Proposition 6.1. *Assume that agents are infallible. Then, it is impossible to hurt the team at the efficient equilibrium by giving each of the players a private signal. Consequently it is impossible to hurt the team at the efficient equilibrium by giving it a public signal.*

[8]Formally, if no public signal X is sent, an equilibrium satisfies

$$\mathbf{z}_i^*(\omega) \in \arg\max_{z_i\in Z_i} \sum_{\omega'\in t_i(\omega)} u(z_i, \mathbf{z}_{-i}^*(\omega'), \omega')p_i(\omega'|t_i(\omega))$$

for each i and ω. If a public signal is sent, then equilibrium satisfies:

$$\mathbf{z}_i^*(\omega) \in \arg\max_{z_i\in Z_i} \sum_{\omega'\in t_i(\omega)\cap X(\omega)} u(z_i, \mathbf{z}_{-i}^*(\omega'), \omega')p_i(\omega'|t_i(\omega) \cap X(\omega))$$

[9]Note that since I simply assume existence, compactness of the strategy sets and continuity of the utility function are clearly *not* assumed for this purpose. In particular, these assumptions are used to guarantee existence of a most efficient equilibrium given the existence of some equilibrium, and in order to establish existence of equilibrium in the game played after a signal X from existence of equilibrium in the game played without the signal in Theorem 6.4.

Proof. Let γ be the game with less information and γ' be the game with more information. Now consider the decision problems $\hat{\gamma}$ and $\hat{\gamma}'$ (corresponding to γ and γ' respectively), in which a single agent selects the *strategies* of all players in an attempt to maximize the ex ante utility of the team. The resulting maximum is an efficient equilibrium in the corresponding game.[10] Since the strategy sets in γ are subsets of the strategy sets in γ', the theorem follows. □

Clearly, when agents are fallible the theorem no longer holds. On the other hand, imposing the condition from Theorem 5.1 on each of the agents individually will not restore the theorem.

Theorem 6.2. *A public signal that reduces the error of every member of a team may still make the team worse off.*

Example 6.3. This example establishes the theorem. Consider the state space $\Omega = \{\omega_1, \omega_2, \omega_3\}$. Suppose that the states are equiprobable according to p^*, and that there are two players. Assume that for all states ω, $t_1(\omega) = t_2(\omega) = \Omega$. Consider the state-dependent utility function given by the following table:

u	L	R
U	1	0
D	θ	10

where $\theta \in \{0, 30\}$. Let $\theta(\omega)$ give the value of θ at each state and assume:

$$\theta(\omega_1) = 30, \theta(\omega_2) = \theta(\omega_3) = 0$$

Assume that the players' conditional probability systems are given by:

$$\begin{aligned} p_1(\omega_2) &= 1 \\ p_1(\omega_1 | \{\omega_1, \omega_3\}) = p_1(\omega_3 | \{\omega_1, \omega_3\}) &= \frac{1}{2} \\ p_2(\omega_3) &= 1 \\ p_2(\omega_1 | \{\omega_1, \omega_2\}) = p_2(\omega_2 | \{\omega_1, \omega_2\}) &= \frac{1}{2} \end{aligned}$$

So, at every state of the world, player 1 believes that the state is ω_2 and player 2 believes that the state is ω_3. Note that in this situation, it is common belief that $\theta = 0$. It follows that there are

[10]This is the step at which the proof would break down if agents are not infallible.

two pure strategy equilibria: (U, L) and (D, R). (D, R) is the efficient equilibrium and gives an ex ante utility of 10 (from the omniscient observer's perspective).

It is easy to verify that every public signal X is honest in this environment, since any signal X is informative in this environment only if it is surprising, in which case the signal induces the agent to believe that the event $X(\omega)$ occurred, and to be uncertain about any event strictly contained in $X(\omega)$. In other words, any informative signal causes the agent to make the weakest assumption consistent with his information. It then follows from Theorem 4.2, that every signal in this environment (weakly) reduces both agents' error.

In particular, consider the public signal X given by the partition:

$$\{\{\omega_1, \omega_2\}, \{\omega_3\}\}$$

The best equilibrium in the game following this signal is worse than the best equilibrium in the game without the signal. To see this, note that at ω_3, the signal informs both players that the state is ω_3, and it is in fact common knowledge that $\theta = 0$. So at the efficient equilibrium, the players play (D, R) at ω_3, and both receive a payoff of 10. On the other hand, at both ω_1 and ω_2, upon receiving the signal, player 1 is certain that $\theta = 0$ and player 2 is uncertain whether $\theta = 0$ or $\theta = 30$, and 2 thinks that these two possibilities are equiprobable. Given this state of uncertainty, L is a dominant strategy for player 2. But then given that player 1 is certain that $\theta = 0$, U is a best response for player 1. It follows that at the efficient equilibrium, the ex ante utility of the team is 4, which is worse than 10, the utility to the team at the efficient equilibrium without the information.

Formally, it is possible to provide a characterization result for team games which is analogous to Theorem 5.1, but which adds an additional condition to the condition that the signal reduce the error of every member of the team:

Theorem 6.4. *Fix a an information structure q. A signal is (weakly) beneficial to a team in all games in $\Gamma(q)$ (at the efficient equilibrium) if and only if*

(1) *It reduces the error of each member of the team, and*
(2) *If at a given state ω, the signal is uninformative to an agent i, and i's error is not null at that state, then at ω, i believes that it is common belief that the signal is uninformative.*

The proof of this theorem shows that for any game γ in $\Gamma(q)$, any game γ' which results by adding a public signal X satisfying (1) and (2) has an equilibrium, and that in γ', the team attains a higher utility than in γ (at the most efficient equilibrium). Conversely, whenever X does not satisfy (1) and (2), one can find a game such that the efficient equilibrium when the game is supplemented by the signal X attains a lower utility than the efficient equilibrium when X is absent. This theorem shows that the knowledge of a *change* in the other player's strategy induced by the signal may amount to harmful information when an agent is making a mistake, for similar reasons that additional knowledge about the state may be harmful in the presence of a mistake. Altering an agent's information state without correcting his error *in any way* may harm him when he is mistaken.

Notice that in Example 6.3, although condition (1) of the theorem is satisfied, condition (2) is violated. In particular, in state ω_1, agent 1 is making a mistake, and the signal X is uninformative to him, but he knows that the signal is informative to agent 2. From this knowledge, he infers that from agent 2's perspective L will now become a dominant action. From agent 1's perspective, the signal is harmful to agent 2. Thus, agent 1 decides to change his action from D to U, and this is what hurts the team. If agent 1 did not know that the signal was informative to agent 2, then agent 1 would have kept his action at D. Then the team would actually be better off with the signal than without it. So in general, if an agent is making a mistake, then his knowledge that the other agent's view will change may be harmful to him (and the team) when his own view of the possibilities does not change. Notice, that this phenomenon only matters when the signal is uninformative, since when it is informative, it reduces, and hence eliminates, his error at that state.

One subtle feature of the analysis is that the reason that 1 thinks that the signal will harm agent 2 in this example is that 1 thinks that the signal will cause 2 to take seriously a possibility ω_1, which from 1's perspective, is not really possible. The error which 1 thinks that 2 is making is the opposite of the type of error which I have assumed agents to make in this paper. That is, in my model, agents make errors be ruling out states that they shouldn't, whereas 1 thinks that 2 makes an error by allowing a state that he shouldn't. This mistaken attribution of error to agent 2 causes agent 1 to take an action which is harmful to the team. One could repeat the analysis of this paper, allowing for more general forms of error from the objective perspective by measuring welfare by $p^*(\cdot|E)$ for some event E rather than by p^*.

7. Conclusion

This paper studies the interaction of error and information both in a single-person setting and in an interactive setting. The perspective of this paper is that while a lot of information is beneficial, a little information can be harmful. I believe that the consequences of this view are significant, both within the field of economics and elsewhere. To take just one example, consider an organization in which information is dispersed, and members of the organization may have opposing views. Suppose there is mediator within the organization who has the authority to settle disagreements. The mediator will gather possibly contradictory information from members of the organization, and will be forced to discard at least some information to maintain consistency. If the mediator can verify all of the information that he is given, then there will be no incentive for anyone to withhold information, since this means that the mediator will only act on the basis of true information, which is in everyone's interests. On the other hand, if there is even a small amount of information that the mediator cannot verify, then this may corrupt true information in the manner described above, and create incentives for members of the organization to withhold information. In this case, as in may other social interactions, the role of error and disagreement is crucial.

References

1. C. Alchourrón, P. Gärdenfors, and D. Makinson, *On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision*, Journal of Symbolic Logic **50** (1985), 510–530.
2. R. Aumann, *Agreeing to Disagree*, Annals of Statistics **4** (1976), 1236–1239.
3. B. Bassan, M. Scarsini, and S. Zamir, *"I Don't Want to Know!": Can it be Rational?*, Discussion Paper 158, Center for the Study of Rationality, Hebrew University, 1997.
4. Bruno Bassan, *Positive Value of Information in Games*, International Journal of Game Theory **32** (2003), 17–31.
5. P. Battigalli and M. Siniscalchi, *Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games*, Journal of Economic Theory **88** (1999), 188–230.
6. ______, *Strong Belief and Forward-Induction Reasoning*, Journal of Economic Theory **106** (2002), 356–391.
7. D. Blackwell, *Comparison of Experiments*, Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability (J. Neyman, ed.), University of California Press, Berkeley, CA, 1951, pp. 93–102.
8. ______, *Equivalent Comparisons of Experiments*, The Annals of Mathematical Statistics **24** (1953), 265–272.
9. O. Board, *Algorithmic Characterization of Rationalizablity in Extensive Form Games*, Working Paper, Department of Economics, Oxford (2002).
10. ______, *Dynamic Interactive Epistemology*, Games and Economic Behavior **49** (2003).
11. G. Bonnano, *A Simple Modal Logic for Belief Revision*, Working Paper, University of California, Davis (2004).
12. J. Geanakoplos, *Game Theory Without Partitions, and Applications to Speculation and Consensus*, Cowles Foundation Discussion Paper No. 914, 1989.
13. O. Gossner, *Comparison of Information Structures*, Games and Economic Behavior **30** (2000), 44–63.
14. S. Morris, *The Logic of Belief and Belief Change; a Decision-Theoretic Approach*, Journal of Economic Theory **69** (1996), 1–23.
15. R. Myerson, *Multistage Games with Communication*, Econometrica **54** (1986), 328–358.
16. A. Neyman, *The Positive Value of Information*, Games and Economic Behavior **3** (1991), 350–355.
17. F. Ramsey, *Weight and the Value of Knowledge*, British Journal for the Philosophy of Science **41** (1990), 1–4.
18. A. Renyi, *On a New Axiomatic Theory of Probability*, Acta Mathematica, Academiae Scientiarum Hungaricae **6** (1955), 285–335.

19. Z. Safra and E. Sulganik, *On the Non-Existence of Blackwell's Theorem-Type Results with General Preference Relations.*, Journal of Risk and Uncertainty **10** (1995), 187–201.
20. E. Schlee, *The Value of Perfect Information in Nonlinear Utility Theory*, Theory and Decision **30** (1991), 127–131.
21. I. Sher, *Mediating Disagreement: A Theory of Purely Informational Conflict*, Working Paper, 2005.
22. W. Spohn, *The Representation of Popper Measures*, Topoi **5** (1986), 69–74.
23. R. Stalnaker, *Knowledge, Belief and Counterfactual Reasoning in Games*, Economics and Philosophy **12** (1996), 133–163.
24. ———, *Belief Revision in Games: Forward and Backward Induction*, Mathematical Social Sciences **36** (1998), 31–56.
25. J-M Tallon, J-C Vergnaud, and S. Zamir, *Contradicting Beliefs and Communication*, Discussion Paper 311, Center for the Study of Rationality, Hebrew University, 2003.
26. P. Wakker, *Non-Expected Utility as Aversion to Information*, Journal of Behavioral Decision Making **1** (1988), 169–175.