

Blackwell's Informativeness Theorem Using Category Theory

Henrique de Oliveira*

March 29, 2016

1 Introduction

An agent faces a *decision problem under uncertainty*, whose payoff $u(a, \omega)$ depends on her action $a \in A$ and the state of the world $\omega \in \Omega$. The agent does not know ω but observes an informative random signal $s \in S$, drawn according to the *information structure* $\sigma : \Omega \rightarrow \Delta(S)$, which specifies the conditional probability $\sigma(s|\omega)$ of observing signal s when the state is ω .

Blackwell (1951, 1953) proved that three different methods to rank information structures generate the same order. The first ranking comes from a notion of “adding noise”: Say that σ' is a *garbling* of σ if an agent who knows σ could replicate σ' by randomly drawing a signal $s' \in S'$ after each observation of $s \in S$, i.e. there exists a function $\gamma : S \rightarrow \Delta(S')$ (the garbling) such that

$$\sigma'(s'|\omega) = \sum_{s \in S} \gamma(s'|s) \sigma(s|\omega).$$

The second ranking comes from a notion of feasibility: Given σ , a mixed strategy $\alpha : S \rightarrow \Delta(A)$ induces a distribution over actions conditional on ω . We can then rank σ and σ' according to which yields the larger set of feasible conditional distributions of actions. The third ranking comes from thinking in terms of expected utility: Say that σ is more informative than σ' if every Bayesian agent, facing any decision problem, can obtain a higher expected utility using σ than by using σ' . Blackwell's theorem can then be stated thus:

Theorem 1. *The following statements are equivalent:*

1. σ' is a garbling of σ ;
2. The set of conditional distributions over actions that are feasible under σ contains those that are feasible under σ' ;

*Pennsylvania State University. Part of the research that led to this paper was conducted while visiting Princeton University. I thank Tomasz Strzalecki for a stimulating conversation that originated this project. I also thank Nageeb Ali, Stephen Morris, Wolfgang Pesendorfer, Philipp Sadowski, Todd Sarver and Asher Wolinsky for their helpful comments and discussions.

3. Every Bayesian agent prefers σ to σ' , for any possible decision problem.

The purpose of this note is to provide a simple proof of this theorem. Blackwell's original proof was rather involved and there have been other attempts to simplify it (Cr mer (1982), Bielinska-Kwapisz (2003), Leshno and Spector (1992)). These papers represent information structures and garblings as matrices and use matrix algebra to derive the result. The proof presented here uses the categorical properties of stochastic maps, avoiding cluttering notation that comes from working with matrices. These categorical properties have been implicitly used by Lehrer et al. (2013) in recent work, but have not been explicitly used before.

The ideas present in the proof can be useful to think about information in other contexts as well. To illustrate this, I show two other applications: to information that arrives over time, reformulating a result of Greenshtein (1996), and to information in strategic environments, deriving the concept of Independent Garbling (Lehrer et al., 2010).

2 The category of stochastic maps

Given a finite set X , let $\Delta(X)$ denote the set of probability distributions over X . A *stochastic map* between two finite sets X and Y is a function $\alpha : X \rightarrow \Delta(Y)$. Given two stochastic maps $\alpha : X \rightarrow \Delta(Y)$ and $\beta : Y \rightarrow \Delta(Z)$, their *composition* is the stochastic map $\beta \circ \alpha : X \rightarrow \Delta(Z)$ defined by

$$\beta \circ \alpha(z|x) = \sum_{y \in Y} \beta(z|y) \alpha(y|x).$$

This composition operation is associative: given another stochastic map $\gamma : Z \rightarrow \Delta(W)$, we can write

$$(\gamma \circ \beta) \circ \alpha = \gamma \circ (\beta \circ \alpha) = \sum_{z \in Z} \sum_{y \in Y} \gamma(w|z) \beta(z|y) \alpha(y|x)$$

and the order of summation does not matter.

A special stochastic map is the *identity map* $id_Y : Y \rightarrow \Delta(Y)$, which takes a point $y \in Y$ to the measure $\delta_y \in \Delta(Y)$ that puts probability 1 on y . The identity map has the property that given any maps $\alpha : X \rightarrow \Delta(Y)$ and $\beta : Y \rightarrow \Delta(Z)$, we have $id_Y \circ \alpha = \alpha$ and $\beta \circ id_Y = \beta$. The existence of an identity, together with the associativity property of composition, formally make stochastic maps a category¹ As is typical in category theory, we will often represent stochastic maps as diagrams of arrows between finite sets, and say that such a *diagram*

¹A *category* is a collection of objects (X, Y, Z etc), arrows between objects (α, β, γ etc) and a binary operation \circ , which takes two arrows $\alpha : X \rightarrow Y$ and $\beta : Y \rightarrow Z$ into a new arrow $\beta \circ \alpha : X \rightarrow Z$, satisfying two properties: (1) Every object Y has an *identity arrow* $id_Y : Y \rightarrow Y$ such that for any $\alpha : X \rightarrow Y$ we have $id_Y \circ \alpha = \alpha$ and for any $\beta : Y \rightarrow Z$ we have $\beta \circ id_Y = \beta$ and (2) \circ is associative (Mac Lane, 1978, page 7).

commutes if any two paths with the same beginning and end define the same stochastic map through composition. For example, saying that the diagram

$$\begin{array}{ccc} A & \xrightarrow{\alpha} & B \\ \downarrow \alpha' & & \downarrow \beta \\ B' & \xrightarrow{\beta'} & C \end{array}$$

commutes is to say that $\beta \circ \alpha = \beta' \circ \alpha'$.

3 Garblings and strategies

Information structures, garblings and strategies are all stochastic maps. Thus, Condition (1) in Theorem 1 can be succinctly stated as the existence of a stochastic map γ such that $\sigma' = \gamma \circ \sigma$. In a diagram,

$$\begin{array}{ccc} \Omega & \xrightarrow{\sigma} & S \\ \sigma' \downarrow & \swarrow \gamma & \\ S' & & \end{array}$$

Likewise, the conditional distribution over actions induced by a strategy α is simply $\alpha \circ \sigma$. Condition (2) in the statement of Blackwell's theorem can be stated in a diagram as "for every α' there exists an α such that the following diagram commutes:"

$$\begin{array}{ccc} \Omega & \xrightarrow{\sigma} & S \\ \sigma' \downarrow & & \downarrow \alpha \\ S' & \xrightarrow{\alpha'} & A \end{array}$$

The expected utility of a strategy only depends on the conditional probability over actions, for it can be written as

$$\sum_{\omega} \left(\sum_{a \in A} u(a, \omega) \alpha \circ \sigma(a|\omega) \right) p(\omega),$$

where $u : A \times \Omega \rightarrow \mathbb{R}$ is the utility function and p is the prior.

4 Proof of Blackwell's Theorem

(1) \Rightarrow (2). We can see this by simply completing the diagram that defines a garbling by defining $\alpha = \gamma \circ \alpha'$:

$$\begin{array}{ccc} \Omega & \xrightarrow{\sigma} & S \\ \sigma' \downarrow & \swarrow \gamma & \downarrow \gamma \circ \alpha' \\ S' & \xrightarrow{\alpha'} & A \end{array}$$

This diagram commutes, so we have (2).

(2) \Rightarrow (1). Let α' be the identity map for the signals of σ' , $Id_{S'}$. Condition (2) tells us that there exists an α that makes the following diagram commutes

$$\begin{array}{ccc} \Omega & \xrightarrow{\sigma} & S \\ \sigma' \downarrow & & \downarrow \alpha \\ S' & \xrightarrow{Id_{S'}} & S' \end{array}$$

But note that in that case we have $\sigma' = Id_{S'} \circ \sigma' = \alpha \circ \sigma$, so that α is a garbling from σ to σ' .

(3) \Rightarrow (3). Fix a set of actions A and a utility function $u : A \times \Omega \rightarrow \mathbb{R}$. Let Λ_σ be the set of all conditional probability over actions that are feasible under σ

$$\Lambda_\sigma = \{\lambda : \Omega \rightarrow \Delta(A) \mid \lambda = \alpha \circ \sigma\}.$$

We can write the expected utility of an agent as

$$\max_{\lambda \in \Lambda_\sigma} \sum_{\omega} \left(\sum_{a \in A} u(a, \omega) \lambda(a|\omega) \right) p(\omega).$$

Clearly, if $\Lambda_\sigma \supseteq \Lambda_{\sigma'}$ the agent can obtain higher expected utility using σ than using σ' .

(3) \Rightarrow (2). We can prove the contrapositive: If (2) is not true, then there exists a $\lambda' \in \Lambda_{\sigma'}$ such that $\lambda' \notin \Lambda_\sigma$. Seeing Λ_σ as a subset of $\mathbb{R}^{\Omega \times A}$, it is easy to show that it is compact and convex. This means that there must exist a separating hyperplane $v \in \mathbb{R}^{\Omega \times A}$ such that

$$\sum_{\omega, a} v(a, \omega) \lambda(a|\omega) < \sum_{\omega, a} v(a, \omega) \lambda'(a|\omega) \quad \forall \lambda \in \Lambda_\sigma$$

Thus, a Bayesian agent with utility v and uniform prior can obtain strictly higher utility using σ' than using any strategy under σ .

5 Dynamic Informativeness

Now consider an agent whose information arrives over time. Let $t = 1, \dots, T$ denote the time periods. At each period, the agent observes a new signal $s_t \in S_t$ and then takes an action $a_t \in A_t$. A *dynamic information structure* $\sigma : \Omega \rightarrow \Delta(S_1 \times \dots \times S_T)$ specifies the probability of all sequences of signals, given a state of the world. A *dynamic decision problem* is given by the sets A_1, \dots, A_T and a utility function $u : A_1 \times \dots \times A_T \times \Omega \rightarrow \mathbb{R}$.

The agent can condition the choice of $a_t \in A_t$ on the realization of all past signals s_1, \dots, s_t but not on the signals s_{t+1}, \dots, s_T , which have not been revealed yet. Hence a *strategy* of the agent specifies a (possibly random) action to be taken after each history of past signals and past actions, that is, it is

a collection of mappings $(\alpha_t)_{t=1}^T$ where $\alpha_1 : S_1 \rightarrow \Delta(A_1)$, $\alpha_2 : S_1 \times S_2 \times A_1 \rightarrow \Delta(A_2)$ etc. For our purposes, it will be more convenient to represent strategies in an alternative way, which we explain below. To simplify notation, let $S^t = S_1 \times \dots \times S_t$, $A^t = A_1 \times \dots \times A_t$ etc for all $t = 1, \dots, T$.

Definition 2. A mapping $\alpha : S^T \rightarrow \Delta(A^T)$ is *adapted* if its marginal distribution on A^t depends only on S^t , that is, if there exist maps $\alpha^t : S^t \rightarrow \Delta(A^t)$ such that

$$\alpha^t(a_1, \dots, a_t | s_1, \dots, s_t) = \sum_{a_{t+1}, \dots, a_T} \alpha(a_1, \dots, a_t, a_{t+1}, \dots, a_T | s_1, \dots, s_T).$$

Lemma 3. *Every strategy induces a unique adapted mapping. Every adapted mapping is induced by some strategy.*

Proof. Given a strategy $(\alpha_t)_{t=1}^T$, let $\alpha^1 = \alpha_1$ and define inductively

$$\alpha^t(a_1, \dots, a_t | s_1, \dots, s_t) = \alpha^{t-1}(a_1, \dots, a_{t-1} | s_1, \dots, s_{t-1}) \alpha_t(a_t | s_1, \dots, s_t, a_1, \dots, a_{t-1}).$$

We can prove by induction that α^t is adapted for every t . Indeed, suppose that α^{t-1} is adapted. The marginal of α^t on A^{t-1} is simply α^{t-1} , which depends only on s_1, \dots, s_{t-1} . The marginal of α^t on A^r for $r < t-1$ is the same as the marginal of α^{t-1} on A^r and therefore depends only on S^r , since α^{t-1} is adapted. Therefore α^t is adapted.

Now suppose that we have an adapted mapping $\alpha : S^T \rightarrow \Delta(A^T)$. Given the mappings corresponding to the marginals $\alpha^t : S^t \rightarrow \Delta(A^t)$, let $\alpha_t : S^t \times A^{t-1} \rightarrow \Delta(A_t)$ give a conditional probability of a_t given a_1, \dots, a_{t-1} and s_1, \dots, s_t .² Then we can show that $\alpha^t = \alpha^{t-1} \alpha_t$ as before, so that the strategy $(\alpha_t)_{t=1}^T$ induces the adapted mapping α^T . \square

Adapted mappings can be conveniently represented using a commutative diagram. Let $\pi_t : S^{t+1} \rightarrow S^t$ and $\rho_t : A^{t+1} \rightarrow A^t$ be the projection mappings (these can be regarded as stochastic mappings that put probability one on a single value). Given a stochastic mapping $\alpha^{t+1} : S^{t+1} \rightarrow \Delta(A^{t+1})$, the composition $\rho_t \circ \alpha^{t+1} : S^{t+1} \rightarrow \Delta(A^t)$ gives precisely the marginal distribution of α^{t+1} on A^t . If α^{t+1} is adapted, then that marginal depends only on S^t : there exists a mapping $\alpha^t : S^t \rightarrow \Delta(A^t)$ such that $\alpha^t \circ \pi_t = \rho_t \circ \alpha^{t+1}$. Thus a mapping $\alpha : S^T \rightarrow \Delta(A^T)$ is adapted if there exist mappings α^t that make the following diagram commute:

$$\begin{array}{ccccccc} S^1 & \xleftarrow{\pi_1} & S^2 & \xleftarrow{\pi_2} & \dots & \xleftarrow{\pi_T} & S^T \\ \alpha^1 \downarrow & & \alpha^2 \downarrow & & & & \alpha \downarrow \\ A^1 & \xleftarrow{\rho_1} & A^2 & \xleftarrow{\rho_2} & \dots & \xleftarrow{\rho_T} & A^T \end{array}$$

²Bayes rule must be respected when the probability of a_1, \dots, a_{t-1} given s_1, \dots, s_t is greater than zero. When it is zero, α_t may be defined arbitrarily.

Lemma 4. *The adapted mappings form a category under the operation of composition of stochastic mappings.*

Proof. Let $\alpha : S^T \rightarrow \Delta(A^T)$ and $\beta : A^T \rightarrow \Delta(B^T)$ be adapted mappings. The only nontrivial step here is to show that their composition $\beta \circ \alpha : S^T \rightarrow \Delta(B^T)$ is also an adapted mapping. To see this, we can write the diagram that represents their composition:

$$\begin{array}{ccccccc}
 S^1 & \xleftarrow{\pi_1} & S^2 & \xleftarrow{\pi_2} & \dots & \xleftarrow{\pi_T} & S^T \\
 \alpha^1 \downarrow & & \alpha^2 \downarrow & & & & \alpha^T \downarrow \\
 A^1 & \xleftarrow{\rho_1} & A^2 & \xleftarrow{\rho_2} & \dots & \xleftarrow{\rho_T} & A^T \\
 \beta^1 \downarrow & & \beta^2 \downarrow & & & & \beta \downarrow \\
 B^1 & \xleftarrow{\lambda_1} & B^2 & \xleftarrow{\lambda_2} & \dots & \xleftarrow{\lambda_T} & B^T
 \end{array}$$

It is easy to see that this diagram commutes. For example,

$$\lambda_1 \circ \beta_2 \circ \alpha_2 = \beta_1 \circ \rho_1 \circ \alpha_2 = \beta_1 \circ \alpha_1 \circ \pi_1.$$

This means that $\beta \circ \alpha$ is an adapted mapping, with the mappings $\beta_t \circ \alpha_t$ corresponding to its marginals. \square

In the proof of Theorem 1, a key property used was that garblings and strategies were the same kind of mathematical objects. There, they were stochastic mappings; here, strategies must be adapted mappings. This suggests that the right notion of garbling for the present dynamic setting is one where the garbling is required to be an adapted mapping. The following theorem confirms this suggestion.

Theorem 5. *The following statements are equivalent:*

1. σ' is an adapted garbling of σ ;
2. The set of conditional distributions over actions that are feasible under σ contains those that are feasible under σ' ;
3. Every Bayesian agent prefers σ to σ' , for any possible dynamic decision problem.

The proof follows exactly the same lines as the proof of Theorem 1 so it is omitted.

6 Strategic Informativeness

Now consider two agents, each receiving a signal about a common state. We can represent such a *joint information structure* as a stochastic mapping $\sigma : \Omega \rightarrow \Delta(S_1 \times S_2)$, where $s_i \in S_i$ is the signal observed by agent i . After seeing

their signal, each agent takes an action from a set A_i . Thus, a strategy for agent i is a stochastic mapping $\alpha_i : S_i \rightarrow \Delta(A_i)$; the *strategy profile* is a stochastic mapping $\alpha \equiv \alpha_1 \otimes \alpha_2 : S_1 \times S_2 \rightarrow \Delta(A_1 \times A_2)$, obtained by taking the independent product of each player’s strategy. The independence can be understood as an assumption that all the information the players have available to them is expressed in the joint information structure σ .

Together, a joint information structure and a strategy profile induce a conditional distribution over actions $\alpha \circ \sigma : \Omega \rightarrow \Delta(A_1 \times A_2)$. When is it that σ permits more conditional distributions over actions than an alternative joint information structure σ' ? Again, the key property here is that garblings and strategies be the same kind of mathematical objects. In this case, we will say that $g : S_1 \times S_2 \rightarrow \Delta(S'_1 \times S'_2)$ is an *independent garbling* if there exist $g_1 : S_1 \rightarrow \Delta(S'_1)$ and $g_2 : S_2 \rightarrow \Delta(S'_2)$ such that g is the independent product of g_1 and g_2 , i.e.

$$g(s'_1, s'_2 | s_1, s_2) = g_1(s'_1 | s_1) g_2(s'_2 | s_2).$$

That this is the right notion of garbling is confirmed by the following result.

Theorem 6. *The following statements are equivalent*

1. σ' is an independent garbling of σ ;
2. The set of conditional distributions over actions that are feasible under σ contains those that are feasible under σ' .

The proof again follows easily the first half of the proof of Theorem 1 once we note that independent garblings form a category.

Notice that this theorem, unlike Theorems 1 and 5, makes no mention of preferences. In order to do that, we would have to specify a game-theoretic solution concept under which players would prefer σ to σ' . This comes with some difficulties, e.g. multiplicity of equilibria, but most importantly it may be that having more feasible actions actually hurt players (players may value commitment). Some proposed solutions involve restricting the class of games under consideration (Gossner, 2000, Lehrer et al., 2010, Peşki, 2008), restricting the class of information structures, or adopting new solution concepts Bergemann and Morris (2015).

References

- Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures. *Theoretical Economics*, 2015.
- Agnieszka Bielinska-Kwapisz. Sufficiency in blackwell’s theorem. *Mathematical Social Sciences*, 46(1):21–25, 2003.
- D. Blackwell. Comparison of experiments. In *Second Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 93–102, 1951.

- D. Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, pages 265–272, 1953.
- Jacques Crémer. A simple proof of blackwell's "comparison of experiments" theorem. *Journal of Economic Theory*, 27(2):439–443, 1982.
- O. Gossner. Comparison of information structures. *Games and Economic Behavior*, 30(1):44–63, 2000.
- Eitan Greenshtein. Comparison of sequential experiments. *The Annals of Statistics*, 24(1):436–448, 1996.
- E. Lehrer, D. Rosenberg, and E. Shmaya. Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2):670–682, 2010.
- Ehud Lehrer, Dinah Rosenberg, and Eran Shmaya. Garbling of signals and outcome equivalence. *Games and Economic Behavior*, 81:179–191, 2013.
- Moshe Leshno and Yishay Spector. An elementary proof of blackwell's theorem. *Mathematical Social Sciences*, 25(1):95–98, 1992.
- Saunders Mac Lane. *Categories for the working mathematician*, volume 5. Springer Science & Business Media, 1978.
- Marcin Peški. Comparison of information structures in zero-sum games. *Games and Economic Behavior*, 62(2):732–735, 2008.