Hybrid Decision Model and the Ranking of Experiments*

Edi Karni[†]   Zvi Safra[‡]

July 14, 2021

Abstract

Departing from the reduction of compound lotteries axiom on multi-stage lotteries this paper proposes a new hybrid model to analyze decision trees. Applied to multi-stage decision trees induced by experiments we show that Blackwell's (1953) definition of the relation "more informative" on the set of information structures is equivalent to experiments being more valuable to a class of non-expected utility preferences. This result extends Blackwell's theorem and provides new insights regarding the evaluation of information produced by experiments.

Keyword: Blackwell's theorem; Comparison of experiments; Reduction of compound lotteries; Value of information; Hybrid decision analysis.

JEL classification: D81, D83

1   Introduction

From a decision making point of view, experimentation is valuable because it provides information that helps decision makers choose courses of actions whose payoffs are higher in the states that are more likely to obtain. Blackwell (1953) formalized this perception as a binary relation, "more informative than" on the set of experiments. According to Blackwell one experiment is more informative than another if, *for every set of feasible actions, it yields a richer menu of experiment-wise expected payoffs* (i.e., expected-loss vectors) each of which corresponds to an action taken contingent on the experimental observations. Blackwell characterized this relation by proving that one experiment is more informative than another if and only if the information content of the latter is obtained by garbling the information content of the former. Equivalently, an experiment is more informative if it allows choices that have higher expected utility. We refer to this equivalence as Blackwell's theorem.[1]

---

[1]Cremer (1982) and Leshno and Spector (1992) provide simple proofs of Blackwell's theorem.

Seen in this way, being better informed seems unambiguously beneficial. Thus, the equivalence between ranking experiments by their information content and their ranking by the expected utility criterion seems oddly restrictive. This equivalence is particularly disconcerting in view of experimental evidence suggesting that subjects violate the tenets of expected utility theory–the sure thing principle and the independence axiom–systematically, and the proliferation, over the last 40 years, of non-expected utility models of decision making under risk and under uncertainty.

To grasp the issue, consider a decision maker facing a choice among feasible actions whose consequences depend on the realization of some underlying states. Suppose that the likelihood of the various states to obtain is quantified by a (prior) probability distribution function. Before choosing an action, the decision maker receives a signal (i.e., an observation), produced by an experiment, that informs him about the likely realization of the states. Upon receiving such signal, the decision maker invokes Bayes' rule to update the prior state probabilities, and then proceeds to choose an action from the feasible set. This process may be thought of as two-stage compound lottery. In the first stage, the experiment produces a signal, according to some probability distribution on the set of signals, following which the decision maker chooses an action. In the second stage, a state is selected (according to the posterior distribution), and the decision maker is awarded the prize that corresponds to the image of the selected state under the chosen action.[2] The question is how do decision makers perceive this two-stage lottery.

In this paper we argue that the critical aspect of the expected utility model that is underlying Blackwell's theorem is the way these compound lotteries are handled. The standard way of handling compound lotteries is by applying the reduction of compound lotteries axiom (henceforth, RCLA). This axiom asserts that a multi-stage lottery is reduced to a single-stage lottery by attributing to each ultimate payoff a probability equal to the product of the probabilities on the events that lead to it. In the context of the choice of experiments the RCLA assigns the ultimate outcome the probability of the signal multiplied by the posterior probabilities of the states to which the chosen course of action assigns that outcome. Analysis that treats the two-stage process as equivalent to its one-stage reduction runs the risk of disregarding subtleties that beset the extensive form decision process. Wakker (1988) and Safra and Sulganik (1995) showed that departing from the independence axiom and maintaining the RCLA implies a widespread and robust aversion towards information.

An alternative procedure of handling compound lotteries is the certainty equivalent reduction. According to this procedure, multi-stage lotteries are reduced to single stage lottery by folding back the lottery tree replacing the lotteries along the branches by their certainty equivalents. This procedure was shown by Schlee (1990) and Safra and Sulganik (1995) to imply that, in non-expected utility theories, information is not always valuable.

---

[2] If the prize itself is a lottery ticket then the procedure described above amounts to three-stage lottery in which, in the third and final stage, the lottery corresponding to the image of the selected state under the chosen action is played out to determine the prizes.

In this work we propose a hybrid decision model that invokes the RCLA in the second stage of the decision making process and a procedure analogous to certainty equivalent reduction in the first stage. We contend that the application of the RCLA, which seems compelling when the transition between the stages is automatic, seems less so when the two stages are separated by an intermediate decision. To formalize this idea we propose a new model, dubbed the *hybrid model*, and show that it identifies a class of preferences that unambiguously value one experiment over another if and only if the information content of the latter is obtained by garbling that of the former. The expected utility model is a special case of this class – in fact it is the only hybrid model that is consistent with the RCLA. Moreover, the application of the hybrid model provides new insight into the manner in which information is evaluated and, while admitting the possibility that some decision makers may not consider information to be unambiguously beneficial, it maintains that, in Blackwell's analytical framework, such behavior is unreasonable.

The surprising (difficult) aspect of Blackwell's theorem is that a more informative experiment (that is, experiment that affords better decisions by the expected utility criterion) implies clearer signals. In this paper, informativeness corresponds to an experiment being more valuable in the sense of affording better decisions for a broader set of preferences, including expected utility preferences. Consequently, this direction of the proof relies on Blackwell's Theorem. The novelty of this paper is the observation that the full power of expected utility, in particular, the RCLA, is not needed for Blackwell's result.

The rest of the paper is organized as follows. In the next section we provide a brief review of Blackwell's (1953) theorem. In section 3 we review the reduction procedures. In section 4 we introduce and characterize the hybrid decision model. In section 5 we extend Blackwell's theorem. In section 6 we discuss the value of information and review the related literature.

2   The Analytical Framework and Blackwell's Theorem

2.1   The analytical framework

Let $S = \{s_1, ..., s_n\}$ be a finite set of *states* and denote by $\Delta(S)$ the simplex in $\mathbb{R}^n$. Let X be a set of *outcomes* and denote by $\Delta(X)$ the set of simple probability distributions on X referred to as *lotteries*.[3] For all $p, q \in \Delta(X)$ and $\alpha \in [0, 1]$ define $\alpha p + (1 - \alpha) q \in \Delta(X)$ by $(\alpha p + (1 - \alpha) q)(x) = \alpha p(x) + (1 - \alpha) q(x)$, for all $x \in X$. Mappings from S to $\Delta(X)$ are referred to as *acts*. Acts represent potential courses of action. The set of all acts is denoted by $\mathcal{H}$. For all $f, g \in \mathcal{H}$ and $\alpha \in [0, 1]$ define $\alpha f + (1 - \alpha) g \in \mathcal{H}$ by $(\alpha f + (1 - \alpha) g)(s) = \alpha f(s) + (1 - \alpha) g(s)$, for all $s \in S$. Constant acts (i.e., acts that assign the same image to every state) are identified with elements of $\Delta(X)$. Thus, $\Delta(X) \subset \mathcal{H}$. Throughout, we denote by δ_x the distribution function that assigns $x \in X$ the unit probability mass and

---

[3]Simple probability distributions are probability distribution functions with finite supports.

$\Delta\left(\Delta\left(X\right)\right)$ denotes the set of simple probability distributions with supports in $\Delta\left(X\right)$.

An *experiment* is a random variable, $\tilde{y}$, taking values in a finite set $\mathbb{Y}$ whose generic elements, denoted by y, are *signals.* Let $\mathcal{Y}$ denote the set of all experiments. Each experiment is identified with a joint probability distribution μ on $\mathbb{Y}\times S$. Given $\tilde{y}\in\mathcal{Y}$ and $s\in S$, the conditional probability on $\mathbb{Y}$ is depicted by a vector of probabilities $\left(\mu\left(y_1\mid s\right),...,\mu\left(y_{|Y|}\mid s\right)\right)$. Let the *image* of $\tilde{y}$ be the set $Y=\{y\in\mathbb{Y}\mid \exists s\in S\ :\mu\left(y\mid s\right)>0\}$ of all potentially realized signals under $\tilde{y}$. For a given prior $\pi\in\Delta\left(S\right)$, denote $\mu\left(y\right)=\Sigma_{s\in S}\mu\left(y\mid s\right)\pi\left(s\right)$.

2.2 Blackwell's theorem

Consider an expected utility maximizing decision maker characterized by a utility function u in $\mathcal{U}$, the set of all real-valued functions on X. Consider a probability distribution $\pi\in\Delta\left(S\right)$ and a non-empty set $B\subseteq\mathcal{H}$ of feasible acts. Suppose that, before choosing an act from B, the decision maker observes the signal y, generated by an experiment $\tilde{y}\in\mathcal{Y}$, and updates the prior distribution π according to Bayes' rule to obtain the posterior probability distribution $\pi\left(s\mid y\right)=\mu\left(y\mid s\right)\pi\left(s\right)/\mu\left(y\right)$, for all $s\in S$. The decision maker's ex-post problem is:

$$\max_{f\in B}\Sigma_{s\in S}\pi\left(s\mid y\right)\Sigma_{x\in X}u\left(x\right)f\left(s\right)\left(x\right). \tag{1}$$

Denoting the maximal value by $U\left(u,\pi\left(\cdot\mid y\right),B\right)$, the expected utility associated with the experiment $\tilde{y}$ is given by:

$$\hat{U}\left(\tilde{y};u,\pi,B\right):=\Sigma_{y\in Y}\mu\left(y\right)U\left(u,\pi\left(\cdot\mid y\right),B\right). \tag{2}$$

Definition 1: An experiment $\tilde{y}$ is *more informative* than another experiment $\tilde{y}'$ if, for all $(u,\pi,B)\in\mathcal{U}\times\Delta S\times 2^{\mathcal{H}}\backslash\varnothing$,

$$\hat{U}\left(\tilde{y};u,\pi,B\right)\geq\hat{U}\left(\tilde{y}';u,\pi,B\right).$$

An *information structure* is a $|S|\times|Y|$ right-stochastic matrix, $P\left(\tilde{y}\right)$, whose generic element is $\mu\left(y\mid s\right)$. Let $\mathcal{M}$ be the set of $|Y|\times|Y'|$ left-stochastic matrices dubbed *garbling matrices.* A garbling matrix M introduces noise that blurs the information in $P\left(\tilde{y}\right)$. More formally, for each entry, $\mu'\left(y'_r\mid s\right)$, of the less informative structure, $P'\left(\tilde{y}'\right)$, $\mu'\left(y'_r\mid s\right)=\Sigma_{y_j\in Y}\mu\left(y_j\mid s\right)m_{jr}$, where $\Sigma_{\{r|y_r\in Y'\}}m_{jr}=1$ for all $\{j\mid y_j\in Y\}$.

Definition 2: An experiment $\tilde{y}$ is *sufficient* for $\tilde{y}'$ if the corresponding information structures satisfy $P\left(\tilde{y}\right)M=P'\left(\tilde{y}'\right)$, for some $M\in\mathcal{M}$.

With these definitions in mind Blackwell's theorem is stated as follows:

Blackwell's Theorem: *An experiment $\tilde{y}$ is more informative than another experiment $\tilde{y}'$ if and only if $\tilde{y}$ is sufficient for $\tilde{y}'$.*

That is, the relation "being sufficient" is equivalent to being ranked higher by all preference relations that admit expected utility representations.

3 Informative Signals and Reduction Procedures

3.1 Signals

According to Blackwell's theorem more informative experiments produce clearer signals in the following sense: When comparing two experiments, every posterior distribution produced by the less informative experiment is a weighted average of the posterior distributions produced by the more informative experiment (i.e., the sufficient experiment). Formally, let $\tilde{y}, \tilde{y}' \in \mathcal{Y}$ with images Y and Y', respectively. Suppose that $\tilde{y}$ is sufficient for $\tilde{y}'$, hence there exists $M \in \mathcal{M}$ such that $P(\tilde{y}) M = P'(\tilde{y}')$. Observe that the equalities

$$\mu'(y_r' \mid s) = \Sigma_{\{j \mid y_j \in Y\}} \mu(y_j \mid s) \, m_{jr}, \tag{3}$$

for all $y_r' \in Y'$, $y_j \in Y$ and $s \in S$, imply

$$\mu'(y_r') = \Sigma_{\{j \mid y_j \in Y\}} \mu(y_j) \, m_{jr}. \tag{4}$$

Therefore, by Bayes' rule,

$$
\begin{aligned}
\pi(s \mid y_r') &= \frac{\pi(s)\,\mu'(y_r' \mid s)}{\Sigma_{s' \in S} \pi(s')\,\mu'(y_r' \mid s')} = \frac{\pi(s)\,\Sigma_{\{j \mid y_j \in Y\}} \mu(y_j \mid s)\, m_{jr}}{\mu'(y_r')} \\[2mm]
&= \frac{\Sigma_{\{j \mid y_j \in Y\}} \pi(s)\,\mu(y_j \mid s)\, m_{jr}}{\mu'(y_r')} = \frac{\Sigma_{\{j \mid y_j \in Y\}} \frac{\mu(y_j)}{\mu(y_j)} \pi(s)\,\mu(y_j \mid s)\, m_{jr}}{\mu'(y_r')} \\[2mm]
&= \frac{\Sigma_{\{j \mid y_j \in Y\}} \mu(y_j)\,\pi(s \mid y_j) m_{jr}}{\mu'(y_r')} = \Sigma_{\{j \mid y_j \in Y\}} \frac{\mu(y_j)\, m_{jr}}{\mu'(y_r')} \pi(s \mid y_j).
\end{aligned} \tag{5}
$$

and the posterior distributions satisfy,

$$\pi(\cdot \mid y_r') = \Sigma_{\{j \mid y_j \in Y\}} \frac{\mu(y_j)\, m_{jr}}{\mu'(y_r')} \pi(\cdot \mid y_j).$$

Consequently, for each act-posterior probability pair $(f, \pi(\cdot \mid y_r')) \in \mathcal{H} \times \Delta(S)$ that is feasible under the less informative experiment, corresponds a set $\{(f, \pi(\cdot \mid y_j)) \mid y_j \in Y\} \subset \mathcal{H} \times \Delta(S)$ of act-posterior probability pairs of the more informative experiment. Thus, from the ex-ante viewpoint, the more informative experiment offers a richer set of opportunities to match feasible acts to the perceived likelihood of the states depicted by their posterior probabilities.

3.2 Reduction procedures

Consider a pair $(f, \pi) \in \mathcal{H} \times \Delta(S)$ of an act and a probability distribution on S. The pair may be regarded as a two-stage lottery in which, in the first stage, a state $s \in S$ is drawn

at random according to the distribution π and, in the second stage, an outcome $x \in X$ is determined by the lottery $f(s) \in \Delta(X)$.

The most common way to reduce such a two-stage lottery to a one-stage lottery in $\Delta(X)$ is RCLA, the *reduction of compound lotteries axiom*. This reduction procedure identifies (f, π) with the one-stage mixture lottery $\Sigma_{s \in S} \pi(s) f(s) \in \Delta(X)$. The axiomatic structure underlying expected utility theory implicitly utilizes RCLA.

Another way to reduce two-stage lotteries is the *certainty equivalent reduction*.[4] This reduction procedure assumes that the decision maker possesses a preference relation over $\Delta(X)$ and that, for every lottery $f(s) \in \Delta(X)$, there exists an element $c(f(s)) \in X$, the *certainty equivalent* of $f(s)$, such that the decision maker is indifferent between $f(s)$ and $\delta_{c(f(s))}$.[5] The certainty equivalent reduction identifies (f, π) with the one-stage mixture lottery $\Sigma_{s \in S} \pi(s) \delta_{c(f(s))} \in \Delta(X)$. Like RCLA, this reduction is implicit in expected utility theory.

In general the resulting one-stage lotteries $\Sigma_{s \in S} \pi(s) f(s)$ and $\Sigma_{s \in S} \pi(s) \delta_{c(f(s))}$ are not equivalent. However, under expected utility they are closely related in the sense that the decision maker is always indifferent between them.

To see how these reductions affect the evaluation of experiments, let a probability distribution $\pi \in \Delta(S)$ and a nonempty set $B \subseteq \mathcal{H}$ of feasible acts be given and consider a decision maker facing an experiment $\tilde{y} \in \mathcal{Y}$. For the decision maker, this decision problem requires a plan of choosing acts in B contingent on the realization of signals produced by $\tilde{y}$. This problem can be described as a three-stage compound lottery. In the first stage, a signal $y \in Y$ is drawn according to the distribution μ. Contingent on the signal and the corresponding posterior distribution $\pi(\cdot \mid y)$, an act $f^*(\pi(\cdot \mid y)) \in B$ is chosen. In the second stage a state s is selected according to the posterior distribution $\pi(\cdot \mid y)$ and the lottery $f^*(\pi(\cdot \mid y))(s) \in \Delta(X)$ is awarded as a prize. In the third and final stage, the lottery $f^*(\pi(\cdot \mid y))(s)$ determines the final outcome $x \in X$. Using the previous notation, for each y the final two stages are described by the pair $(f^*(\pi(\cdot \mid y)), \pi(\cdot \mid y)) \in \mathcal{H} \times \Delta(S)$.

Applying RCLA to the second stage, the pairs $(f^*(\pi(\cdot \mid y)), \pi(\cdot \mid y))$ are reduced to the one-stage lotteries $\Sigma_{s \in S} \pi(s \mid y) f^*(\pi(\cdot \mid y))(s) \in \Delta(X)$. Applying RCLA again, this time to the first stage, further simplifies the three-stage compound lottery and reduces it to the equivalent one-stage mixture lottery

$$\sum_{y \in Y} \mu(y) \sum_{s \in S} \pi(s \mid y) f^*(\pi(\cdot \mid y))(s). \tag{6}$$

The process under which RCLA is applied twice is called the RCLA *procedure* and is commonly used to analyze behavior of non-expected utility decision makers. This procedure implies a widespread and robust aversion to information. The reason for this aversion is that, under RCLA, non-expected utility decision makers may be dynamically inconsistent.

---

[4]For more detailed discusssion and application of certainty equivalent reduction see Segal (1987).

[5]The existence of certainty equivalents depends on the richness of the set of outcomes X.

That is, evaluating the optimal acts $f^*\left(\pi\left(\cdot\mid y\right)\right)$ from an *ex-ante* point of view, by looking at the one-stage lottery (6), the decision maker may find out that he prefers to replace some of these *ex-post* optimal acts by others. However, when he finds himself at the second stage, optimality dictates choosing the acts $f^*\left(\pi\left(\cdot\mid y\right)\right)$.

Alternatively, the certainty equivalent reduction can be applied twice to yield the one-stage mixture lottery

$$\sum_{y\in Y}\mu\left(y\right)\delta_{c\left(\Sigma_{s\in S}\pi(s|y)\delta_{c(f^*(\pi(\cdot|y))(s))}\right)}. \tag{7}$$

This is called the certainty equivalent reduction *procedure* and under it information is not always valuable for non-expected utility decision makers. Unlike RCLA, however, the reason here is not dynamic inconsistency (it is easy to verify that, viewing the one-stage lottery (7) *ex-ante*, the decision maker is always in agreement with his planned *ex-post* optimal acts $f^*\left(\pi\left(\cdot\mid y\right)\right)$). Rather, the reason for this aversion is that the mixture and certainty equivalent operators are not commutative for non-expected utility decision makers. Thus, having larger sets of options under a sufficient experiment does not necessarily translate to being in a preferable situation.

An alternative procedure, which we propose and study in this paper, applies the RCLA to the second stage and a reduction analogous to the certainty equivalent to the first stage.[6] Under this *hybrid procedure*, which is characterized in the next section, for each signal y, the second-stage reduction induces a set of one-stage lotteries $\{\Sigma_{s\in S}\pi\left(s\mid y\right)f\left(s\right)\mid f\in B\}$. Assuming that these lotteries are evaluated by some utility functional v (see next section), let f_y^* be a maximizer of v over this set (that is, $f_y^*\in\arg\max_{f\in B}v\left(\Sigma_{s\in S}\pi\left(s\mid y\right)f\left(s\right)\right)$). Then, the value of the experiment $\tilde{y}$ is given by

$$\sum_{y\in Y}\mu\left(y\right)v\left(f_y^*\right).$$

The justification for applying distinct procedures to the different stages is the nature of the uncertainties involved. In the second stage, given the act and the (updated) state probabilities, the outcome is selected "algorithmically" without interference by the decision maker. By contrast, after the first stage, corresponding to each signal there is an interim step at which the decision maker interferes by updating the state probabilities and choosing an act. This aspect of the dynamic process suggests that decision makers may regard the first stage as qualitatively distinct from the later stages and, consequently, treat them differently. Specifically, according to the hybrid procedure, assessing the value of experiments, decision makers envision the acts that they would choose contingent on the signals, assign these acts utility values and take the mean utility values as the value of the experiment.

---

[6]The analogous reduction does not require the existece of certainty equivalents.

4   Characterization of the Hybrid Representation

Assume that X is a compact topological space and consider a probability distribution $\pi \in \Delta(S)$, a nonempty set $B \subseteq \mathcal{H}$ and an experiment $\tilde{y} \in \mathcal{Y}$. As above, the experiment can be identified with a two-stage decision tree in which, in the first stage, a signal in $y \in Y$ is realized and, in the second stage, by choosing an act $f \in B$ to maximize his ex-post preferences at the decision node associated with y, the decision maker is faced with the reduced lottery $f_y^* \in \Delta(X)$. From an ex-ante viewpoint, the ex-post choices are seen as the $|Y|+1$ tuple $(\mu, (f_y^*)_{y \in Y}) \in \Delta(Y) \times \Delta(X)^{|Y|}$ and the whole experiment can be identified with a subset of all $|Y|+1$ tuples $(\mu, (p_1, ..., p_{|Y|})) \in \Delta(Y) \times \Delta(X)^{|Y|}$. To analyze all possible experiments, assume that the cardinality of $\mathbb{Y}$ (the set of all signals) is N $(< \infty)$ and consider the set of all $N+1$ tuples $(\mu, (p(y))_{y \in \mathbb{Y}}) \in \Delta(\mathbb{Y}) \times \Delta(X)^N$. A decision maker is characterized by two complete and transitive, preference relations: (i) an ex-ante preference relation $\succcurlyeq$ on $\Delta(\mathbb{Y}) \times \Delta(X)^N$, that ranks decision trees, and (ii) an ex-post preference relation $\succcurlyeq^*$ on $\Delta(X)$, that is used to choose optimal acts at the realized decision nodes.

Continuity - (of $\succcurlyeq$ and $\succcurlyeq^*$ with respect to the corresponding Euclidean topology) The sets $\{(\mu', (p(y))_{y \in \mathbb{Y}} \in \Delta(\mathbb{Y}) \times \Delta(X)^N \mid (\mu, (p(y))_{y \in \mathbb{Y}} \succcurlyeq (\mu', (p'(y))_{y \in \mathbb{Y}})\}$ and $\{(\mu', (p(y))_{y \in \mathbb{Y}} \in \Delta(\mathbb{Y}) \times \Delta(X)^N \mid (\mu', (p'(y))_{y \in \mathbb{Y}} \succcurlyeq (\mu, (p(y))_{y \in \mathbb{Y}})\}$ are closed for all $\left(\mu, (p(y))_{y \in \mathbb{Y}}\right) \in \Delta(\mathbb{Y}) \times \Delta(X)^N$ and the sets $\{p \in \Delta(X) \mid q \succcurlyeq^* p\}$, $\{p \in \Delta(X) \mid p \succcurlyeq^* q\}$ are closed for all $q \in \Delta(X)$.

Henceforth we assume that both $\succcurlyeq$ and $\succcurlyeq^*$, are continuous binary relations. To capture consequentialism we assume that the ex-post preference relation $\succcurlyeq^*$ is independent of the realized signal y and of the probability distribution μ.

The next axiom asserts that if a signal y is certain to obtain then the ex-ante ranking of any two lotteries agrees with their ex-post ranking. Formally, define $\mathbf{p} = (p(y))_{y \in \mathbb{Y}} \in \Delta(X)^N$, then

Consistency - *For all* $\mathbf{p}, \mathbf{q} \in \Delta(X)^N$ *and* $y \in \mathbb{Y}$, $(\delta_y, \mathbf{p}) \succcurlyeq (\delta_y, \mathbf{q})$ *if and only if* $p(y) \succcurlyeq^* q(y)$.

The next axiom asserts that independence applies to the first stage of the decision making process. To formalize this idea we introduce the partial mixture operation on $\Delta(\mathbb{Y}) \times \Delta(X)^N$: For all $(\mu, \mathbf{p})$, $(\nu, \mathbf{p}) \in \Delta(\mathbb{Y}) \times \Delta(X)^N$ and $\lambda \in [0, 1]$, define $\lambda(\mu, \mathbf{p}) + (1-\lambda)(\nu, \mathbf{p}) = (\lambda\mu + (1-\lambda)\nu, \mathbf{p})$. Consider a decision maker who is indifferent between the alternatives $(\mu, \mathbf{p})$ and $(\mu', \mathbf{q})$ and prefers the alternative $(\nu, \mathbf{p})$ over $(\nu', \mathbf{q})$. Suppose that, facing a choice between the decision trees $A = (\lambda\mu + (1-\lambda)\nu, \mathbf{p})$ and $B =$

$(\lambda\mu' + (1-\lambda)\nu', \mathbf{q})$ he reasons that if the event whose probability is λ obtains then if he has chosen A he is faced with the alternative $(\mu, \mathbf{p})$ and if he has chosen B he faces the alternative $(\mu', \mathbf{q})$. Conditional on the realization of this event, he is indifferent between A and B. By the same logic he would prefer A over B conditional on the realization of the complemetary event whose probability is $1-\lambda$. Consequently, he prefers A over B unconditionally.[7] Formally,

First-Stage Independence - *For all $(\mu, \mathbf{p})$, $(\nu, \mathbf{p})$, $(\mu', \mathbf{q})$, $(\nu', \mathbf{q})$ in $\Delta(\mathbb{Y}) \times \Delta(X)^N$ and $\lambda \in [0,1]$, if $(\mu, \mathbf{p}) \sim (\mu', \mathbf{q})$ then $(\nu, \mathbf{p}) \succcurlyeq (\nu', \mathbf{q})$ if and only if $(\lambda\mu + (1-\lambda)\nu, \mathbf{p}) \succcurlyeq (\lambda\mu' + (1-\lambda)\nu', \mathbf{q})$.*

The next theorem provides the hybrid representation and shows that an ex-ante preference relation satisfying consistency (with respect to an ex-post preference relation) and first-stage independence is representable as a weighted sum of the ex-post utilities.

Theorem 1: *Let $\succcurlyeq$ and $\succcurlyeq^*$ binary relations on $\Delta(\mathbb{Y}) \times \Delta(X)^N$ and $\Delta(X)$, respectively, such that $\succ^* \neq \emptyset$. Then the following conditions are equivalent:*

(a) $\succcurlyeq$ and $\succcurlyeq^*$ are complete, transitive, and continuous, jointly they satisfy consistency and $\succcurlyeq$ satisfies first-stage independence.

(b) There exist continuous non-constant functions $V : \Delta(\mathbb{Y}) \times \Delta(X)^N \to \mathbb{R}$ and $v : \Delta(X) \to \mathbb{R}$, such that V represents $\succcurlyeq$, v represents $\succcurlyeq^*$, and, for all $(\mu, \mathbf{p}) \in \Delta(\mathbb{Y}) \times \Delta(X)^N$,

$$V(\mu, \mathbf{p}) = \sum_{y \in \mathbb{Y}} \mu(y) \, v(p(y)). \tag{8}$$

Moreover, if there are functions $\bar{V} : \Delta(\mathbb{Y}) \times \Delta(X)^N \to \mathbb{R}$ and $\bar{v} : \Delta(X) \to \mathbb{R}$ such that $\bar{V}$ represents $\succcurlyeq$, $\bar{v}$ represents $\succcurlyeq^*$ and $\bar{V}(\mu, \mathbf{p}) = \sum_{y \in \mathbb{Y}} \mu(y) \, \bar{v}(p(y))$ then $\bar{v} = bv + a$, $b > 0$.

Proof: The proof is omitted since it is identical to that of Theorem 1 of Karni and Safra (2000) with the additional assumption that the ex-post preference relation $\succcurlyeq^*$ is signal-independent.

Remark: The behavior of a decision maker with such a pair of ex-ante and ex-post preference relations displays *dynamic consistency.* Formally, if $\mathbf{p}$ and $\mathbf{q}$ differ only when signal y occurs and $\mu(y) > 0$, $p(y)$ is preferred to $q(y)$ ex-post if and only if $\mathbf{p}$ is preferred to $\mathbf{q}$ ex-ante.

---

[7] This axiom is analogous to the constrained independence axiom defined in Karni and Safra (2000).

The case where the ex-post preferences satisfy the independence axiom is of special interest. Let $U(p) = \Sigma_{x \in X} u(x) p(x)$ be an expected utility representation of $\succcurlyeq^*$. Then, by Theorem 1, the ex-ante preferences $\succcurlyeq$ are represented by

$$V(\mu, \mathbf{p}) = \sum_{y \in \mathbb{Y}} \mu(y) w(U(p(y))) \tag{9}$$

where w is strictly increasing. This representation is more general than expected utility (for $\succcurlyeq$) and is reduced to it when w is affine.[8] It resembles the class of non-expected utility models that have representations known as smooth ambiguity attitudes.[9] Note, however, that ambiguity aversion requires that the function w be concave. In the present context, concave w implies aversion to the spread of signal-contingent payoffs that more informative experiments afford. See further discussion of this at the end of the next section.

5 Blackwell's Theorem Extended

Consider all pairs $(\succcurlyeq, \succcurlyeq^*)$ of ex-ante and ex-post preferences satisfying the requirements of Theorem 1(a). Let $\mathcal{V}$ be the set of all pairs (V, v) satisfying (8) and let $\mathcal{V}_{cx} \subset \mathcal{V}$ consist of all pairs in which v is convex (the need for convexity will be clarified in the next theorem). As above, for given $\pi \in \Delta(S)$, $B \subseteq 2^{\mathcal{H}} \backslash \varnothing$ and $\tilde{y} \in \mathcal{Y}$, let

$$f_y^* \in \arg\max_{f \in B} v(\Sigma_{s \in S} \pi(s \mid y) f(s)) \tag{10}$$

be the one-stage reduced lottery that maximizes the ex-post utility v at the decision node associated with y. As the ex-ante preferences $\succcurlyeq$ admit the hybrid representation and are dynamically consistent, the ex-ante value of $\tilde{y}$ is given by $V(\mu, (f_y^*)_{y \in Y}) = \sum_{y \in Y} \mu(y) v(f_y^*)$. Following this, for every $((V, v), \pi, B) \in \mathcal{V}_{cx} \times \Delta(S) \times 2^{\mathcal{H}} \backslash \varnothing$, the value of an experiment $\tilde{y} \in \mathcal{Y}$ with image Y is defined by:

$$\hat{V}(\tilde{y}; (V, v), \pi, B) := \sum_{y \in Y} \mu(y) v(f_y^*). \tag{11}$$

Definition 3: An experiment $\tilde{y}$ is *more hybrid-valuable* than another experiment $\tilde{y}'$ if, for all $((V, v), \pi, B) \in \mathcal{V}_{cx} \times \Delta(S) \times 2^{\mathcal{H}} \backslash \varnothing$,

$$\hat{V}(\tilde{y}; (V, v), \pi, B) \geq \hat{V}(\tilde{y}'; (V, v), \pi, B).$$

The next theorem extends Blackwell's (1953) theorem.

---

[8] In the social choice framework, an axiomatization of such preferences is given by Grant et. al (2010).

[9] See Klibanoff, Marinacci and Mukerji, (2005) and Seo (2009).

Theorem 2: *An experiment $\tilde{y}$ is more hybrid-valuable than another experiment $\tilde{y}'$ if and only if $\tilde{y}$ is sufficient for $\tilde{y}'$.*

The proof below rests on two properties of Blackwell's result: First, that the expected utility under the sufficient experiment is larger. Second, the set of expected utility payoffs generated by the sufficient experiment constitutes a larger "spread" of expected utility values. With v being monotonic increasing and convex these properties make the expected value of v larger for the sufficient experiment.

Proof: (Sufficiency) Suppose that $\tilde{y}$ is sufficient for $\tilde{y}'$. Let Y and Y' denote the images of $\tilde{y}$ and $\tilde{y}'$, respectively. Fix $((V,v),\pi,B) \in \mathcal{V}_{cx} \times \Delta(S) \times 2^{\mathcal{H}}\backslash\varnothing$. Then,

$$
\begin{aligned}
\hat{V}\left(\tilde{y}';(V,v),\pi,B\right) &= \sum_{\{r|y_r'\in Y'\}} \mu'\left(y_r'\right) v\left(\sum_{s\in S} f_{y_r'}^*\left(s\right)\pi\left(s\mid y_r'\right)\right) \qquad (12) \\
&= \sum_{\{r|y_r'\in Y'\}} \mu'\left(y_r'\right) v\left(\sum_{s\in S} f_{y_r'}^*\left(s\right)\left[\sum_{\{j|y_j\in Y\}} \frac{\mu\left(y_j\right)m_{jr}}{\mu'\left(y_r'\right)}\pi\left(s\mid y_j\right)\right]\right) \\
&= \sum_{\{r|y_r'\in Y'\}} \mu'\left(y_r'\right) v\left(\sum_{\{j|y_j\in Y\}} \frac{\mu\left(y_j\right)m_{jr}}{\mu'\left(y_r'\right)}\left[\sum_{s\in S} f_{y_r'}^*\left(s\right)\pi\left(s\mid y_j\right)\right]\right) \\
&\leq \sum_{\{r|y_r'\in Y'\}} \mu'\left(y_r'\right) \sum_{\{j|y_j\in Y\}} \frac{\mu\left(y_j\right)m_{jr}}{\mu'\left(y_r'\right)} v\left(\sum_{s\in S} f_{y_r'}^*\left(s\right)\pi\left(s\mid y_j\right)\right) \\
&= \sum_{\{r|y_r'\in Y'\}}\sum_{\{j|y_j\in Y\}} \mu\left(y_j\right)m_{jr} v\left(\sum_{s\in S} f_{y_r'}^*\left(s\right)\pi\left(s\mid y_j\right)\right) \\
&\leq \sum_{\{r|y_r'\in Y'\}}\sum_{\{j|y_j\in Y\}} \mu\left(y_j\right)m_{jr} v\left(\sum_{s\in S} f_{y_j}^*\left(s\right)\pi\left(s\mid y_j\right)\right) \\
&= \sum_{\{j|y_j\in Y\}} \mu\left(y_j\right) v\left(\sum_{s\in S} f_{y_j}^*\left(s\right)\pi\left(s\mid y_j\right)\right) \sum_{\{r|y_r'\in Y'\}} m_{jr} \\
&= \hat{V}\left(\tilde{y};(V,v),\pi,B\right)
\end{aligned}
$$

where the second equality follows from (5), the first inequality holds since v is convex, the second inequality holds since, by definition, $v(\Sigma_{s\in S} f_{y_r'}^*\left(s\right)\pi\left(s\mid y_j\right)) \leq v(\Sigma_{s\in S} f_{y_j}^*\left(s\right)\pi\left(s\mid y_j\right))$, and last equality holds since $\Sigma_{\{r|y_r\in Y'\}} m_{jr} = 1$.

(Necessity) If v is affine then the expected utility representations are a subset of the set of preference relations that have hybrid representations with $(V,v) \in \mathcal{V}_{cx}$. Hence, necessity is implied by the necessity part of Blackwell's theorem. ♠

11

The ingredients of the representation of the hybrid model for affine v in (9) are similar to those of the class of non-expected utility models known as smooth ambiguity models. Smooth ambiguity aversion requires that the analogous of the function w be concave. By contrast our extension of Blackwell's theorem requires that the decision makers exhibit *information proclivity*, an attitude that is captured by the convexity of the function w. This difference between the two models may be explained as follows: More informative experiments offer larger spreads of expected utility payoffs to choose from. Consequently, in addition to the unambiguous advantage of allowing decision makers to choose acts whose payoffs are higher in the states that are more likely to obtain, information affects the decision making process through the decision makers' attitudes towards wider spread of expected payoffs. We showed that the value of information increases if the decision maker displays information proclivity (that is, prefer a wider spread of expected payoffs). Information may be less valuable if the decision maker displays payoff-spread aversion, which is captured by concave w. The case of expected utility maximizing behavior (i.e., the function w is linear) is exhibits neutrality towards the payoffs spread.

Ambiguity aversion in the smooth ambiguity models seems justifiable since the decision makers are confronted with a spread of the priors that introduce payoffs variation that they cannot exploit. By the same token, information proclivity is justifiable because the decision makers value the opportunity to choose from, and thereby exploit, the wider spread of expected payoffs afforded by the more informative experiments.[10]

To further explain this point, let $S = \{s_1, s_2\}$, $Y = \{y_1, y_2\}$, $\pi = \left(\frac{1}{2}, \frac{1}{2}\right)$ and assume that v is strictly convex. Choose $p, q \in \Delta(X)$ such that $v\left(\frac{1}{2}p + \frac{1}{2}q\right) < \frac{1}{2}v(p) + \frac{1}{2}v(q)$, and consider $B = \{f\}$ where $f(s_1) = p$ and $f(s_2) = q$. Let $\tilde{y}$ be the fully informative experiment (that is, $\mu(y_j \mid s_j) = 1$, $j = 1, 2$) and $\tilde{y}'$ be the totally uninformative experiment $(\pi(\cdot \mid y_j) = \pi(\cdot)$, $j = 1, 2)$. Clearly, $\tilde{y}$ sufficient for $\tilde{y}'$. Then,

$$\hat{V}\left(\tilde{y}; (V, v), \pi, B\right) = \frac{1}{2}v(p) + \frac{1}{2}v(q) > v\left(\frac{1}{2}p + \frac{1}{2}q\right) = \hat{V}\left(\tilde{y}'; (V, v), \pi, B\right)$$

This illustrates the fact that, in addition to affording better choice of acts information has intrinsic value in the sense that, ex-ante, the decision-maker prefers a larger spread of expected payoffs, even if their expectations is the same.,

This example can be used to demonstrate the necessity of the convexity of v. Assume that v is not convex and choose $p, q \in \Delta(X)$ such that $v\left(\frac{1}{2}p + \frac{1}{2}q\right) > \frac{1}{2}v(p) + \frac{1}{2}v(q)$. Then

$$\hat{V}\left(\tilde{y}; (V, v), \pi, B\right) = \frac{1}{2}v(p) + \frac{1}{2}v(q) < v\left(\frac{1}{2}p + \frac{1}{2}q\right) = \hat{V}\left(\tilde{y}'; (V, v), \pi, B\right)$$

and $\tilde{y}$, the fully informative experiment, becomes less hybrid-valuable than the totally uninformative experiment $\tilde{y}'$.

---

[10] Gensbittel, Renou and Tomala (2015) analysed the amiguity case and provided conditions under which information is desiralbe within the Maxmin model. Our main result does the same for the smooth ambiguity model. Note that in that paper, similarly to ours, the decision maker cannot commit ex-ante to executing ex-post choices. An analysis with commitment (under ambiguity) was carried out by Li and Zhou (2016).'

6   Concluding Remarks

6.1   The value of information

Lurking in the background of Blackwell's theorem are two tacit properties of expected utility theory – *consequentialism* and *reduction of compound lotteries*. The former maintains that, facing sequential decisions involving risky choices, decision makers are "forward looking" in the sense that, at every decision node, their preferences are unaffected by outcomes that did not materialize, or "roads not taken," along the decision making path. The latter asserts that decision makers evaluate acts solely by the ultimate probability distributions they induce on outcomes regardless of whether the outcome is drawn, in a single step, from a known distribution or is arrived at by more convoluted trajectory that includes chance and decision nodes.

Wakker (1988) showed that departing from the independence axiom while maintaining consequentialism and reduction of compound lotteries (at all stages) and assuming convexity, necessarily result in situations in which decision makers refuse free information related to finer partitions. Safra and Sulganik (1995) demonstrated that, maintaining consequentialism and reduction of compound lotteries throughout, violations of the relation of being more informative experiment à la Blackwell are quite robust. More specifically, for almost all pairs of experiments $\tilde{y}$ and $\tilde{y}'$ satisfying that $\tilde{y}$ is sufficient for $\tilde{y}'$, there exists a non-expected utility preferences that strictly prefer the latter.

Replacing the reduction of compound lotteries assumption by the certainty equivalent reduction, used at all stages of the decision problem, is not sufficient to derive an analogue of Blackwell's theorem (see Schlee, [1990], and Safra and Sulganik, [1995]). This and the previous observations clarify the need for hybrid models.

6.2   Related literature

Segal (1990) was the first to propose a model of decision making under risk in which he replaced the reduction of compound lotteries with certainty equivalent reduction, to characterized rank-dependent utility model.[11] Seo (2009) obtains smooth ambiguity averse representations of choice under uncertainty that departs from the reduction of compound lotteries axiom. Halevy (2007) presented experimental evidence suggesting that subjects whose behavior violates reduction of compound lotteries under risk are more likely to exhibit ambiguity aversion when facing decision making under uncertainty.

The present work argues that departing from the independence axiom of expected utility theory while maintaining dynamic consistency and the intuitive presumption that being better informed is unambiguously desirable, justifies the departure from the RCLA in sequential decision situations in which, at the interim stages, information may be exploited

---

[11]Since the independence axiom of expected utility theory implies the RCLA, departure from the latter property implies the departure from the former axiom.

by choosing acts that better match the underlying data.

References

[1] Blackwell, D. (1953) "Equivalent Comparisons of Experiments," *Annals of Mathematics and Statistics*, 74, 265-272.

[2] Cremer, J. (1982) "A Simple Proof of Blackwell's "Comparison of Experiments" Theorem." *Journal of Economic Theory*, 27, 439-443.

[3] Gensbittel, F., Renou, L. and Tomala, T. (2015) "Comparisons of Ambiguous Experiments," Working Paper.

[4] Grant, S. Kajii, A. Polak, B. and Safra Z. (2010) "Generalized Utilitarianism and Harsanyi's Impartial Observer Theorem," *Econometrica,* 78(5):1939–1971.

[5] Halevy, Y. (2007) "Ellsberg Revisited: An Experimental Study," *Econometrica*, 75, 503–536.

[6] Karni, E. and Safra, Z. (2000) "An extension of a theorem of von Neumann and Morgenstern with an application to social choice theory," *Journal of Mathematical Economics*, 34, 315-327.

[7] Klibanoff, P., Marinacci, M. and Mukerji, S. (2005) "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73 , 1849-1892

[8] Leshno, M. and Spector, Y. (1992) "An Elementary Proof of Blackwell' Theorem." *Mathematical Social Sciences*, 25, 95-98.

[9] Li, J. and Zhou, J. (2016) "Blackwell's Informativeness Ranking with Uncertainty Averse Preferences," *Games and Economic Behavior,* 96, 18-29.

[10] Schlee, E. (1990) "The Value of Information in Anticipated Utility Theory," *Journal of Risk and Uncertainty,* 3, 83-92.

[11] Segal, U. (1990) "Two Stage Lotteries without the Reduction Axiom." *Econometrica*, 58, 349–377.

[12] Seo, K. (2009) "Ambiguity and Second-Order Beliefs," *Econometrica*, 77, 1575–1605.

[13] Wakker, P. (1988) "Nonexpected Utility as Aversion to Information," *Journal of Behavioral Decision Making* 1, 169-175.