On the value of information in games:
or, some things you might rather not know

from Osborne's *Game Theory*, p. 283

$\frac{1}{2}$ 1 $\frac{1}{2}$

$\frac{1}{2}$ 2 $\frac{1}{2}$

State ω_1

	L	M	R
T	$1, 2\epsilon$	$1, 0$	$1, 3\epsilon$
B	$2, 2$	$0, 0$	$0, 3$

State ω_2

	L	M	R
T	$1, 2\epsilon$	$1, 3\epsilon$	$1, 0$
B	$2, 2$	$0, 3$	$0, 0$

Figure 283.1 The first Bayesian game considered in Section 9.3.1.

9.3.1 More information may hurt

A decision-maker in a single-person decision problem cannot be worse off if she has more information: if she wishes, she can ignore the information. In a game the same is not true: if a player has more information and the other players know that she has more information, then she may be worse off.

Consider, for example, the two-player Bayesian game in Figure 283.1, where $0 < \epsilon < \frac{1}{2}$. In this game there are two states, and neither player knows the state. Player 2's unique best response to each action of player 1 is L. (If player 1 chooses T, L yields 2ϵ, whereas M and R each yield $\frac{3}{2}\epsilon$; if player 1 chooses B, L yields 2, whereas M and R each yield $\frac{3}{2}$.) Further, player 1's unique best response to L is B. Thus (B, L) is the unique Nash equilibrium of the game; it yields each player a payoff of 2. (If you have studied Chapter 4, you will be able to verify that, moreover, the game has no other mixed strategy equilibrium.)

Now consider the variant of this game in which player 2 is informed of the state: player 2's signal function τ_2 satisfies $\tau_2(\omega_1) \neq \tau_2(\omega_2)$. In this game $(T, (R, M))$ is the unique Nash equilibrium. (Each type of player 2 has a strictly dominant action, to which T is player 1's unique best response.)

Player 2's payoff in the unique Nash equilibrium of the original game is 2, whereas her payoff in the unique Nash equilibrium of the game in which she knows the state is 3ϵ in each state. Thus she is worse off when she knows the state than when she does not. To understand this result, notice that among player 2's actions, R is good only in state ω_1, M is good only in state ω_2, and L is a compromise. When she does not know the state she chooses L, inducing player 1 to choose B. When she is fully informed she tailors her action to the state, choosing R in state ω_1 and M in state ω_2, inducing player 1 to choose T. The game has no steady state in which she ignores her information and chooses L because this action leads player 1 to choose B, making R better for player 2 in state ω_1 and M better in state ω_2.

A Basic Theorem of (Bayesian) Expected Utility Theory:

If you can postpone a terminal decision in order to observe, *cost free*, an experiment whose outcome might change your terminal decision, then it is strictly better to postpone the terminal decision in order to acquire the new evidence.

The analysis also provides a value for the new evidence, to answer:

How much are you willing to "pay" for the new information?

An agent faces a current decision:

- **with k terminal options $D = \{d_1, ..., d^*, ..., d_k\}$ (d^* is the best of these)**
- **and one sequential option: first conduct experiment X, with outcomes $\{x_1, ..., x_m\}$ that are observed, then choose from D.**

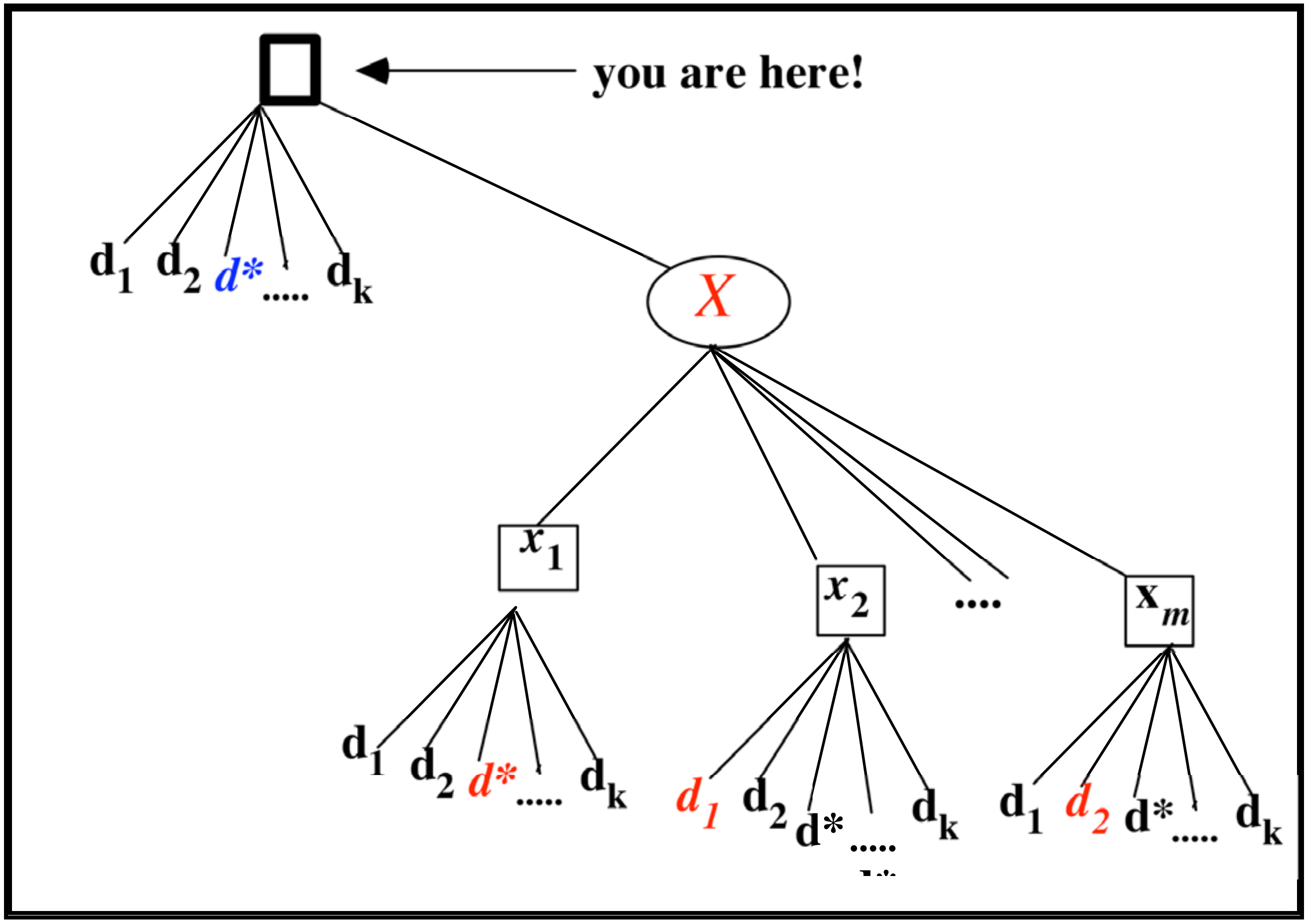

Terminal decisions (acts) as functions from states to outcomes

The canonical decision matrix: **decisions × states**

	s_1	s_2			s_j			s_n
d_1	O_{11}	O_{12}			O_{1j}			O_{1n}
d_2	O_{21}	O_{22}			O_{2j}			O_{2n}
d_m	O_{m1}	O_{m2}			O_{mj}			O_{mn}

$d_i(s_j)$ = outcome $\mathbf{o_{ij}}$.

What are "**outcomes**"? That depends upon which version of expected utility you consider. We will allow arbitrary outcomes, providing that they admit a von Neumann-Morgenstern cardinal utility $U(\bullet)$.

A central theme of Subjective Expected Utility [*SEU*] is this:

- axiomatize preference $\leq$ over decisions so that

$$d_1 \leq d_2 \ \textbf{iff} \ \Sigma_j \mathbf{P}(s_j)\mathbf{U}(o_{1j}) \leq \Sigma_j \mathbf{P}(s_j)\mathbf{U}(o_{2j}),$$

for **one** subjective (personal) probability $\mathbf{P}(\bullet)$ defined over *states*

and **one** cardinal utility $\mathbf{U}(\bullet)$ defined over *outcomes*.

- Then the decision rule is to choose that (an) option that *maximizes SEU*.

Note: In this version of SEU, which is the one that we will use here:

(1) decisions and states are probabilistically independent, $\mathbf{P}(s_j) = \mathbf{P}(s_j \mid d_i)$.

Reminder: This is sufficient for a fully general *dominance* principle.

(**2**) Utility is state-independent, $\mathbf{U_j}(o_{i,j}) = \mathbf{U_h}(o_{g,h})$, if $o_{i,j} = o_{g,h}$.

Here, $\mathbf{U_j}(o_{\bullet j})$ is the conditional utility for outcomes, given state s_j.

(**3**) (Cardinal) Utility is defined up to positive linear transformations,

$\mathbf{U}'(\bullet) = a\mathbf{U}(\bullet) + b$ (a > 0) is also the same utility function for purposes of *SEU*.

Note: More accurately, under these circumstances with act/state prob. independence, utility is defined up to a similarity transformation: $\mathbf{U_j}'(\bullet) = a\mathbf{U_j}(\bullet) + b_\mathbf{j}$. So, maximizing SEU and Maximizing Subjective Expected Regret-Utility are equivalent decision rules.

Reconsider the value of cost-free evidence when decisions conform to maximizing *SEU*. Recall, the decision maker faces a choice *now* between k-many terminal options $\boldsymbol{D} = \{d_1, ..., d^*, ..., d_k\}$ (d^* maximizes SEU among these k options). There is one sequential option: first conduct experiment X, with sample space $\{x_1, ..., x_m\}$, and then choose from $\boldsymbol{D}$ having observed X. Options in *red* maximize SEU at the choice nodes, using $\mathbf{P}(s_j \mid X = x_i)$.

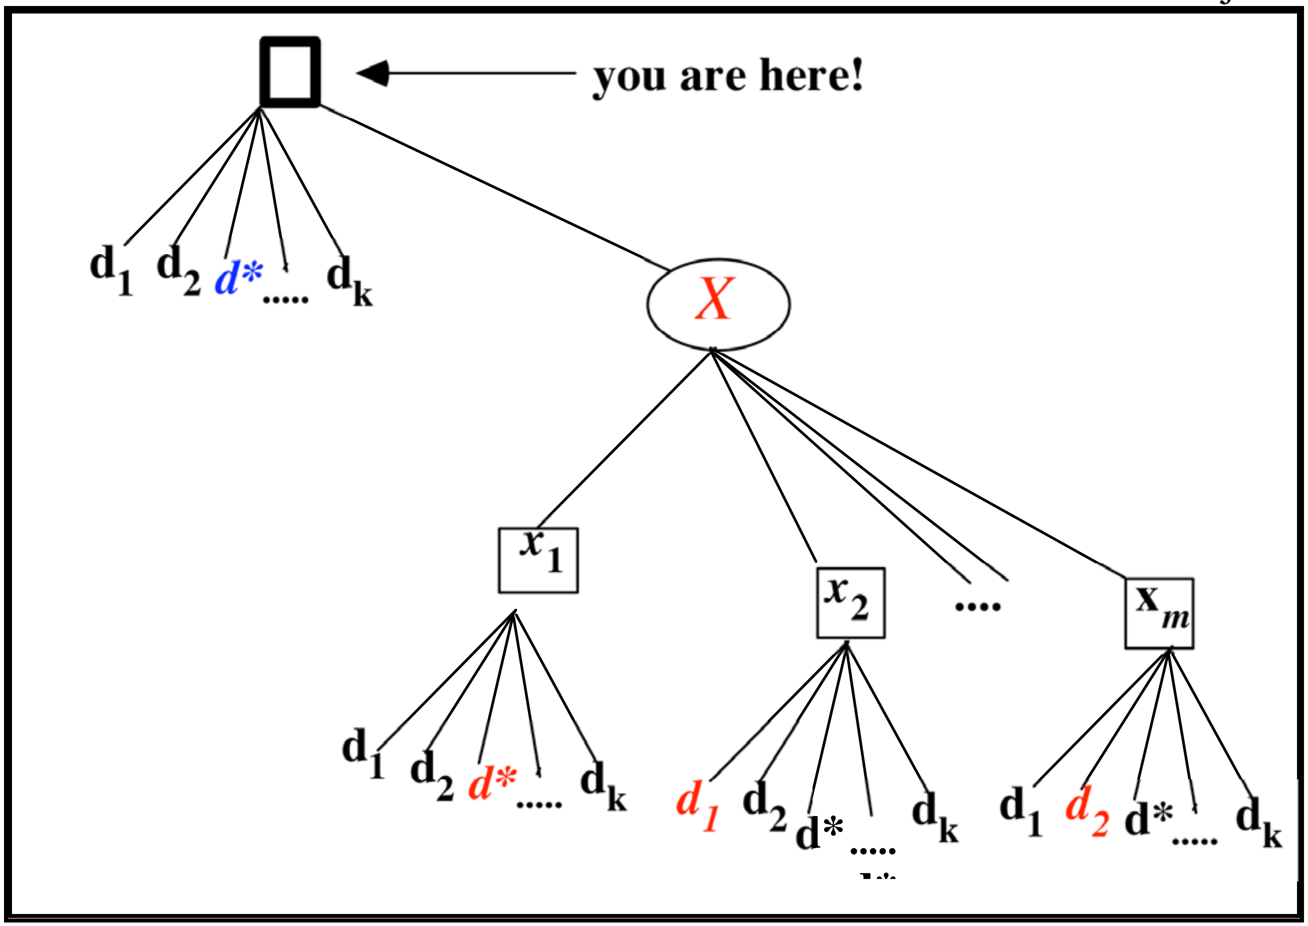

By the law of conditional expectations: $E(Y) = E(E[Y \mid X])$**.**

With Y **the Utility of an option** $U(d)$**, and** X **the outcome of the experiment,**

$$\mathrm{Max}_{d \in D}\, E(U(d)) = E(U(d^*))$$

$$= E(E(U(d^*) \mid X))$$ **("ignoring** X**" when choosing)**

$$\leq E(\mathrm{Max}_{d \in D}\, E(U(d) \mid X))$$

$$= U(\text{sequential option}).$$

- **Hence, the academician's** ***first-principle*****:**
  Never decide today what you might postpone until tomorrow
  in order to learn something new.

- $U(d^*) = U(\text{sequential option})$ **if and only if the new evidence** X **never leads you to a different terminal option.**

- $U(\text{sequential option}) - E(U(d^*))$ **is the** ***value of the experiment*****: what you will pay (at most) in order to conduct the experiment prior to making a terminal decision.**

4 cases under which the central result does not obtain, and you are willing to pay in order not to learn prior to making a terminal decision.

The first 3 cases are discussed in *Is Ignorance Bliss?*, Kadane, Schervish & Seidenfeld, *J.Phil*. 2008.

1. **The decision rule is not Expected Utility Maximization with a single probability distribution. For example, represent uncertainty of an event using a (convex) <u>set of probabilities</u>, $\mathcal{P}$. Let the decision rule be Γ-Maximin – choose an act whose min expected utility is max w.r.t. set $\mathcal{P}$. Then the value of (cost free) information may be negative.**
   - **This is the fate of inference with *pivotal variables* in statistical inference.**

2. **The Law of Conditional Expectations fails: $E(Y) \neq E(\,E\,[Y \mid X]\,)$. For example, if expectations are based on a finitely, but not countably additive probability – corresponding to an "improper" prior in Bayesian statistical inference – then the value of (cost free) information may be negative.**

3. **The new information may fail to be "cost free." A familiar setting is where sampling carries an explicit cost. A less familiar setting for costly information is where your utility of an outcome includes your own state of ignorance about that outcome.**
   - **The Taxi-Driver (our Example 12) illustrates this: I hold a ticket to a mystery play and take a taxi to the theater. The taxi-driver knows who-done-it and offers (threatens?) to inform me unless the tip is sufficiently large. It may be sensible to pay to avoid learning this information!**

4. A case that we do not discuss at length is where act/state dependence obtains in the choice whether or not to learn new information prior to making a terminal decision.

Recall: With act/state dependence even simple dominance is no longer valid!

	ω_1	ω_2
Act_1	**3**	**1**
Act_2	**4**	**2**

Regardless that Act_2 dominates Act_1, if $P(\omega_i \mid Act_i) > 3/4$ then Act_1 has greater (conditional) expected utility than Act_2.

The typical model for act/state dependence is the presence of Moral Hazard, (e.g., in insurance) where the states of uncertainty for the decision maker involve the actions of another (rational) agent – as in a game!

However, regarding the principal result about the value of cost-free information, it is a side-issue whether the act/state dependence involves the actions of another decision-maker, or not.

A Toy Example of act/state dependence without Moral Hazard where new (cost free) information has negative value.

<u>Binary Terminal Decision</u>

	ω_1	ω_2
d_1	1	0
d_2	0	1

Suppose $P(\omega_1) = .75$. Without added information $d^* = d_1$, and $U(d^*) = .75$.

Let $X = \{0,1\}$ be an irrelevant binary variable with likelihood,

$$P(X=0 \mid \omega_1) = P(X=0 \mid \omega_2) = .80.$$

So, X is irrelevant to Ω.

However, suppose that the decision to *<u>observe</u>* X alters the "prior" probability over Ω so that, $P(\omega_1 \mid \textit{observe } X) = .60 < .75$. Then $U(\textit{observe } X) = .60 < 75$.

In this case, because of act/state dependence, the decision maker strictly prefers not to observe (cost free) X prior to making the terminal decision $D = \{d_1, d_2\}$.

> **Aside: The example allows for an unconditional probability over X, but not over Ω, because it makes no sense to put a (non-trival) probability distribution over the space of one's own current options!**

Now, let's reconsider Osborne's example where cost-free information carries negative value in a game.

9.3 Two examples concerning information 283

Player 1: $\frac{1}{2}$, $\frac{1}{2}$

Player 2: $\frac{1}{2}$, $\frac{1}{2}$

State ω_1

	L	M	R
T	$1, 2\epsilon$	$1, 0$	$1, 3\epsilon$
B	$2, 2$	$0, 0$	$0, 3$

State ω_2

	L	M	R
T	$1, 2\epsilon$	$1, 3\epsilon$	$1, 0$
B	$2, 2$	$0, 3$	$0, 0$

Figure 283.1 The first Bayesian game considered in Section 9.3.1.

9.3.1 More information may hurt

A decision-maker in a single-person decision problem cannot be worse off if she has more information: if she wishes, she can ignore the information. In a game the same is not true: if a player has more information and the other players know that she has more information, then she may be worse off.

Consider, for example, the two-player Bayesian game in Figure 283.1, where $0 < \epsilon < \frac{1}{2}$. In this game there are two states, and neither player knows the state.

	P(T & τ_1) = $\alpha/2$	P(B & τ_1) = $(1-\alpha)/2$	P(T & τ_2) = $\alpha/2$	$P(B$ & $\tau_2)$ = $(1-\alpha)/2$
L	2ε	2	2ε	2
M	0	0	3ε	3
R	3ε	3	0	0

with $0 < \varepsilon < 1/2$

<u>Column-player's probability assumptions</u>

P(type = τ_1) = ½. P(Top) = α. P(*Row* & type) = P(*Row*)/2 – these are independent factors. Moreover, since play is simultaneous between players:

P(*Row* & type | *Column's act*) = P(*Row* & type).

Here we have act/state <u>*independence*</u> in the game with simultaneous play.

<u>Column Player's Expected Utilities for the three options</u>

U[*L*] = $2(1-\alpha(1-\varepsilon)) >$

U[*M*] = $(3/2)(1-\alpha(1-\varepsilon))$ = U[*R*] = $(3/2)(1-\alpha(1-\varepsilon))$

So, Column-player chooses *L*, regardless the value of α.

This is known to Row-player, who then chooses *B* to maximize her/his utility.

That choice also is known to Column player; hence, <u>$\alpha = 0$</u>.

Then Column's U[*L*] = 2. Likewise, 2 is the sure payoff for Row's choice *B*.

Version 2*a* – Column-player learns her/his type prior to choosing a terminal option, and Row-player knows only that fact.

***Contingent play given Column-player's type*.**

If Column-player may choose among {*L*, *M*, *R*} <u>contingent on his/her type</u>, τ_i (*i* =1,2) then *R* dominates both *M* and *L*, given type = τ_1
and *M* dominates both *L* and *R*, given type = τ_2.
So the dominant contingent strategy for Column player is (*R* if τ_1, M if τ_2).
Since play is simultaneous between players, act/state independence obtains.
So the dominant play for Column has "prior" (ex ante) expected utility,

$$U[R \text{ if } \tau_1; M \text{ if } \tau_2] = 3(1 - \alpha'(1\text{-}\varepsilon)),$$

where α' is Column player's "prior" for Row choosing *Top* in Version 2a.

Recall in Version 1, Column's P(Top) = α. If $\alpha = \alpha'$, then

$$U[R \text{ if } \tau_1; M \text{ if } \tau_2] = 3(1 - \alpha(1\text{-}\varepsilon)) > U[L] = 2(1 - \alpha(1\text{-}\varepsilon))$$

and Column player has <u>positive value</u> for the information of her/his type.

HOWEVER, in the second version of the problem, since Row-player also knows these calculations on behalf of Column-player, and as Row-player's option *T* dominates option *B* given either *M* or R – with payoffs 1 *vs* 0 -- then, Row-player chooses *T*, and Column player knows this too.

$$\text{So, } \underline{\alpha' = 1} \neq \alpha = 0.$$

In Version 2*b* of the game both players learn Column's type prior to making a terminal decision. The upshot is the same.

By dominance, Column plays: *R* if τ_1; *M* if τ_2. Knowing this Row plays *T*, etc.

From Column's perspective,

in Version 2*a* (or 2*b*), U[*R* if τ_1; *M* if τ_2)] = 3ε < 2 = Version 1's U[*L*].

Column prefers the first version of the game. Similarly for Row player!

So, if the *__initial choice__* (for either player to make) is whether to play Version 1, or instead to play Version 2a of the game, the initial choice is to play Version 1 of the game. Likewise in a choice between Version 1 and Version 2b.

However, in this sequential problem, in choosing first between Version 1 and Version 2 of the game, and then playing the version chosen, there is act/state dependence from either player's perspective, probabilistic dependence between the player's choice of Version 1 vs. Version 2 of the game and her/his probability for how the other player chooses.

From Column-player's perspective, the mere choice of version fixes the value of α – Column player's probability that Row player chooses Top, *T*.

Likewise, in choosing between Version 1 and Version 2 of the game, Row player faces act/state dependence in her/his probability for Column's behavior.

Thus, the familiar result about the non-negative value of cost-free information does not apply in this sequential game. Each player prefers Version 1 over Version 2. Each player prefers playing the game with less information rather than more.

But this is also what happens when there is only one decision maker and she/he faces a problem with act/state dependence in probabilities. The opportunity to postpone a cost-free decision may have negative value (with or without the Moral Hazard of *another* decision maker's choice) provided that there is act/state dependence in personal probabilities.