

Comparison of Experiments in Monotone Problems

Alfredo Di Tillio¹ Marco Ottaviani² Peter N. Sørensen³

June 22, 2022

¹Bocconi

²Bocconi

³Copenhagen

Decisions under Uncertainty and Experiments


- Uncertainty: Unknown state of nature $\theta \in \Theta$, prior probability $\pi(\theta)$
- Experiment: Observe $s \in S$ with probability $P(s|\theta)$

$$\implies \text{Posterior } \pi(\theta|s) = \frac{\pi(\theta)P(s|\theta)}{\sum_{\theta'} \pi(\theta')P(s|\theta')}$$

- Decision: Take action $a \in A$ to maximize


$$\sum_{\theta} \pi(\theta|s)u(a, \theta)$$

Blackwell's Comparison of Experiments (1951)



Experiment P is **more informative** than experiment Q when P allows higher expected utility than Q for **every** decision problem

Blackwell's Comparison of Experiments (1951)



Experiment P is **more informative** than experiment Q when P allows higher expected utility than Q for **every** decision problem



Q is a **garbling** of P

Garbling

Let $P : \Theta \rightarrow \Delta(S)$ and $Q : \Theta \rightarrow \Delta(T)$ be two experiments.

Q is a **garbling** of P if there exists $g : S \rightarrow \Delta(T)$ such that

$$Q(t|\theta) = \sum_s P(s|\theta) g(t|s) \quad \forall \theta$$

i.e. with experiment P you can “imitate” experiment Q ...

Comparison of Experiments in Monotone Problems

We restrict attention to **monotone** decision problems.

To fix ideas: $A = \{a_1, a_2\}$, with $u(a_2, \theta) - u(a_1, \theta)$ increasing in θ .

- Fewer decision problems: more experiments are comparable

Comparison of Experiments in Monotone Problems

We restrict attention to **monotone** decision problems.

To fix ideas: $A = \{a_1, a_2\}$, with $u(a_2, \theta) - u(a_1, \theta)$ increasing in θ .

- Fewer decision problems: more experiments are comparable

What makes good information in monotone problems?

We present new findings

- For the case where the better experiment P has binary support, we provide two new characterizations of dominated experiments Q

Comparison of Experiments in Monotone Problems

We restrict attention to **monotone** decision problems.

To fix ideas: $A = \{a_1, a_2\}$, with $u(a_2, \theta) - u(a_1, \theta)$ increasing in θ .

- Fewer decision problems: more experiments are comparable

What makes good information in monotone problems?

We present new findings

- For the case where the better experiment P has binary support, we provide two new characterizations of dominated experiments Q
- We consider further parts of Blackwell's theorem (not today)

Looking Ahead: Monotone Problems \Leftrightarrow IDO Preferences

States and actions are ordered so that (Quah-Strulovici, 2009)

- For all $\theta' > \theta$ and $a' > a$,

$$u(a', \theta) \geq (>) u(a, \theta) \implies u(a', \theta') \geq (>) u(a, \theta')$$

whenever $u(a', \theta) \geq u(a'', \theta)$ for all a'' with $a \leq a'' \leq a'$

*If action a' is the best action in the interval $[a, a']$ when state is θ ,
 a' remains the best action in $[a, a']$ when state is raised to θ'*

Looking Ahead: Monotone Problems \Leftrightarrow IDO Preferences

States and actions are ordered so that (Quah-Strulovici, 2009)

- For all $\theta' > \theta$ and $a' > a$,

$$u(a', \theta) \geq (>) u(a, \theta) \implies u(a', \theta') \geq (>) u(a, \theta')$$

whenever $u(a', \theta) \geq u(a'', \theta)$ for all a'' with $a \leq a'' \leq a'$

*If action a' is the best action in the interval $[a, a']$ when state is θ ,
 a' remains the best action in $[a, a']$ when state is raised to θ'*

IDO (interval dominance order) class of preferences includes:

- Milgrom-Shannon (1994) single-crossing preferences
- Karlin-Rubin (1956) monotone preferences

Equivalent Comparisons 1:

Less Utility \rightarrow LR Box

Likelihood Ratio Order

In experiment P , signal s' is LR-greater than s ($s' >_{LR} s$) when

$$\frac{P(s'|\theta')}{P(s'|\theta)} > \frac{P(s|\theta')}{P(s|\theta)} \quad \forall \theta' > \theta$$

or equivalently

$$\frac{\pi(\theta'|s')}{\pi(\theta|s')} > \frac{\pi(\theta'|s)}{\pi(\theta|s)} \quad \forall \theta' > \theta$$

Likelihood Ratio Order

In experiment P , signal s' is LR-greater than s ($s' >_{LR} s$) when

$$\frac{P(s'|\theta')}{P(s'|\theta)} > \frac{P(s|\theta')}{P(s|\theta)} \quad \forall \theta' > \theta$$

or equivalently

$$\frac{\pi(\theta'|s')}{\pi(\theta|s')} > \frac{\pi(\theta'|s)}{\pi(\theta|s)} \quad \forall \theta' > \theta$$


An implication of IDO utility:

- Quah-Strulovici, Theorem 2: If $s' >_{LR} s$, the optimal action at s' is no smaller than the optimal action at s

Pictures with Three States

$$\Theta = \{\theta_1, \theta_2, \theta_3\} \text{ with } \theta_1 < \theta_2 < \theta_3$$


Represent beliefs as points in a Marschak triangle:



Binary IDO problem

Start from binary action problems, $A = \{a_1, a_2\}$ with $a_1 < a_2$, then:


Green line: indifference



Binary IDO problem

Start from binary action problems, $A = \{a_1, a_2\}$ with $a_1 < a_2$, then:

Green line: indifference




IDO \Leftrightarrow

1. green line crosses base of triangle &
2. high action a_2 optimal left of green line, low action a_1 to right

Binary Experiment


Binary experiment ($S = \{s_1, s_2\}$) spreads arbitrary prior π to posteriors p, q



Likelihood Ratio Order

Belief p dominates another belief q in LR order ($p >_{LR} q$) when


$$\frac{p_2}{p_1} > \frac{q_2}{q_1} \quad \text{and} \quad \frac{p_3}{p_2} > \frac{q_3}{q_2}$$



Likelihood Ratio Order

Belief p dominates another belief q in LR order ($p >_{LR} q$) when

$$\frac{p_2}{p_1} > \frac{q_2}{q_1} \quad \text{and} \quad \frac{p_3}{p_2} > \frac{q_3}{q_2}$$




$\Leftrightarrow q$ lies below red line and above blue line

Likelihood Ratio Order

Blue area: beliefs **below** p in LR order


Yellow area: beliefs **above** p in LR order



Likelihood Ratio Order

Blue area: beliefs **below** p in LR order


Yellow area: beliefs **above** p in LR order



Thm of QS visible: any IDO indifference line through p is steeper than rays:
 p to right of line \Rightarrow entire blue area to right of line


Likelihood Ratio Order

Similarly for another belief q :




Non-monotone LR experiment

A binary experiment violating monotone LR property:



MLR

A binary experiment satisfying MLR (monotone LR):



Aside: Two states

With only two states, signals of binary experiments can always be arranged to satisfy MLR

Aside: Two states

With only two states, signals of binary experiments can always be arranged to satisfy MLR

In this case, Proposition 1 in Jewitt (2007) implies that our comparison notion (with IDO preferences) is equivalent to Blackwell's (unrestricted)

- Directly, binary action sets suffice when experiment Q is binary
- In any such problem, actions can be ordered so IDO is satisfied

Aside: Two states

With only two states, signals of binary experiments can always be arranged to satisfy MLR

In this case, Proposition 1 in Jewitt (2007) implies that our comparison notion (with IDO preferences) is equivalent to Blackwell's (unrestricted)


- Directly, binary action sets suffice when experiment Q is binary
- In any such problem, actions can be ordered so IDO is satisfied

Three states is the simplest case where monotonicity makes a difference

Likelihood Ratio Box

LR box (or LR interval) from q to $p >_{LR} q$ is the set of beliefs r such that


$$p >_{LR} r >_{LR} q$$



Main Question 1: Better Experiment without MLR

Suppose first that binary experiment (p, q) **violates MLRP**:


Result: (p, q) **dominates (only) its Blackwell garblings**



Main Question 1: Better Experiment without MLR

Suppose first that binary experiment (p, q) **violates MLRP**:

Result: (p, q) **dominates (only) its Blackwell garblings**



Proof:

For every experiment with at least one posterior outside the p - q segment, there is an IDO problem that separates posteriors of that experiment, but does not separate p from q


Main Question 2: Better Experiment with MLR

Now suppose binary experiment (p, q) **satisfies MLRP**.

Main Question 2: Better Experiment with MLR


Now suppose binary experiment (p, q) **satisfies MLRP**.

Take another experiment with **posteriors not all inside box**, as in (r, s)
(going forward, notation abuse, s denotes one posterior)



Comparing Experiments: LR Box

If comparison experiment (r, s) has at least one posterior outside box:



Some IDO problem separates (r, s) but not (p, q) .

$\Rightarrow (p, q)$ **does not dominate any experiment with some posteriors outside box.**

Equivalent Comparisons 2: LR Box \rightarrow Quasi-Garbling

Quasi-Garblings

Q is a **quasi-garbling** of P if for every order on T there exists $g : \Theta \times S \rightarrow \Delta(T)$ such that


- $Q(t|\theta) = \sum_s P(s|\theta)g(t|\theta, s) \quad \forall \theta$
- $g(\cdot|\theta', s) >_{FOSD} g(\cdot|\theta, s) \quad \forall s \text{ and } \forall \theta' > \theta$

LR Box \subset Quasi-Garblings

Consider experiment Q with **posteriors all inside box**.

Result: Q is a quasi-garbling of P .

E.g. binary experiment (r, s)



LR Box \subset Quasi-Garblings

Proof for binary P with $S = \{s_1, s_2\}$ and binary Q with $T = \{t_1, t_2\}$.

Assume P satisfies MLRP when $s_1 > s_2$.

Assume $t_1 > t_2$ (case $t_2 > t_1$ analogous).

Write $p_\theta := P(s_1|\theta)$, $g_\theta := g(t_1|\theta, s_1)$ and $h_\theta := g(t_1|\theta, s_2)$.

Fix any θ and g_θ, h_θ such that $Q(t_1|\theta) = p_\theta g_\theta + (1 - p_\theta)h_\theta$.

Consider higher state $\theta' > \theta$ and equation

$$Q(t_1|\theta') = p_{\theta'} g_{\theta'} + (1 - p_{\theta'}) h_{\theta'}.$$

When P satisfies MLRP, p_θ increasing in θ , so equation has solution $g_{\theta'}, h_{\theta'}$ such that $g_{\theta'} \leq g_\theta$ and $h_{\theta'} \leq h_\theta$.

Equivalent Comparisons 3: Quasi-Garbling \rightarrow Less Utility

Informativeness of Quasi-Garblings

Let p_θ denote optimal probability of choosing a_2 in state θ under P .

If Q is a quasi-garbling of P , probability of choosing a_2 in state θ under Q has form

$$g_\theta p_\theta + h_\theta(1 - p_\theta)$$

where g_θ and h_θ are both decreasing in θ .


By linearity of payoffs, payoff under Q is maximized when each sequence is, in fact, constantly 0 or constantly 1.

When both constantly 0 or both constantly 1, you are not using information under Q .

When g_θ constantly 1 and h_θ constantly 0, you are replicating what you get with P . When g_θ constantly 0 and h_θ constantly 1, you are inverting what you do with P , something you could have done (but chose not to do) with P . Thus, in any case, A does better than B.

Direct Proof that LR Box \rightarrow Less Utility

Normalization WLOG: $u(a_1, \theta) = 0$ for all θ , and $u(a_2, \theta_3) > 0 > u(a_2, \theta_1)$



Picture shows hardest case (when s leads to a_2)

Only interesting IDO problem where different actions are taken at r, s

By IDO and LR-bounds, also a_2 optimal at p and a_1 optimal at q

Proof for binary-binary comparison

Claim: (r, s) can be improved upon unless $s_2/s_1 = p_2/p_1$ or $r_2/r_1 = q_2/q_1$ (or both)

- If both constraints are relaxed, room to reduce $P(s|\theta_1)$
- The change raises s_2/s_1 and reduces r_2/r_1 , staying in LR box
- The change reduces weight on $u(a_2, \theta_1) < 0$ in expected utility

Proof for binary-binary comparison

Claim: (r, s) can be improved upon unless $s_2/s_1 = p_2/p_1$ or $r_2/r_1 = q_2/q_1$ (or both)

- If both constraints are relaxed, room to reduce $P(s|\theta_1)$
- The change raises s_2/s_1 and reduces r_2/r_1 , staying in LR box
- The change reduces weight on $u(a_2, \theta_1) < 0$ in expected utility

Same claim holds for s_3/s_2 or r_3/r_2

- Otherwise room to raise $P(s|\theta_3)$

Proof for binary-binary comparison

Claim: (r, s) can be improved upon unless $s_2/s_1 = p_2/p_1$ or $r_2/r_1 = q_2/q_1$ (or both)

- If both constraints are relaxed, room to reduce $P(s|\theta_1)$
- The change raises s_2/s_1 and reduces r_2/r_1 , staying in LR box
- The change reduces weight on $u(a_2, \theta_1) < 0$ in expected utility

Same claim holds for s_3/s_2 or r_3/r_2

- Otherwise room to raise $P(s|\theta_3)$

For more than three states, it works inductively

Proof for binary-binary comparison

Claim: $s \sim_{LR} p$ or $r \sim_{LR} q$ (or both)

- Relies on previous claim and $p >_{LR} q$
- If $s_2/s_1 = p_2/p_1$ and $r_2/r_1 \geq q_2/q_1$ then $P(s|\theta_1) \leq P(p|\theta_1)$ and $P(s|\theta_2) \leq P(p|\theta_2)$ (algebra shows, next slide)
- Then also $s_3/s_2 = p_3/p_2$ binds (algebra shows, next slide)
- And so on with more states, inductively
- Likewise for the case with r

Proof for binary-binary comparison

Claim: $s \sim_{LR} p$ or $r \sim_{LR} q$ (or both)

- Relies on previous claim and $p >_{LR} q$
- If $s_2/s_1 = p_2/p_1$ and $r_2/r_1 \geq q_2/q_1$ then $P(s|\theta_1) \leq P(p|\theta_1)$ and $P(s|\theta_2) \leq P(p|\theta_2)$ (algebra shows, next slide)
- Then also $s_3/s_2 = p_3/p_2$ binds (algebra shows, next slide)
- And so on with more states, inductively
- Likewise for the case with r

Finally, when $s \sim_{LR} p$

- then $P(s|\theta_k) \leq P(p|\theta_k)$ in all states (and proportional)
- Expected utility under (r, s) is proportional (less than one-for-one) to that under (p, q) which is positive

Beyond Binary Dominated Experiments Q


Argument with three states

Suppose dominated experiment has **three posteriors** r, s, t

Argument with three states

Suppose dominated experiment has **three posteriors** r, s, t


Say s lies on one side of p - q segment while r, t on the opposite side:



Argument with three states

Suppose dominated experiment has **three posteriors** r, s, t

Say s lies on one side of p - q segment while r, t on the opposite side:



View experiment (r, s, t) as **composite**:

Starting from prior π , belief is spread

1. either to π_1 and then to s or r ,
2. or to π_2 and then to s or t .

Beyond Binary Dominated Experiment

Let λ be such that $\pi^1 = \frac{P(r)r + \lambda P(s)s}{P(r) + \lambda P(s)}$.

Let $P^1 = P(r) + \lambda P(s)$ and $P^2 = 1 - P^1 = P(t) + (1 - \lambda) P(s)$

Beyond Binary Dominated Experiment

Let λ be such that $\pi^1 = \frac{P(r)r + \lambda P(s)s}{P(r) + \lambda P(s)}$.

Let $P^1 = P(r) + \lambda P(s)$ and $P^2 = 1 - P^1 = P(t) + (1 - \lambda) P(s)$

Define $\mu^i \in (0, 1)$ such that $\pi^i = \mu^i p + (1 - \mu^i) q$.

Then $P(p) = \mu^1 P^1 + \mu^2 P^2$, since $\pi = P(p)p + P(q)q = P^1 \pi^1 + P^2 \pi^2$.

Beyond Binary Dominated Experiment

Let λ be such that $\pi^1 = \frac{P(r)r + \lambda P(s)s}{P(r) + \lambda P(s)}$.

Let $P^1 = P(r) + \lambda P(s)$ and $P^2 = 1 - P^1 = P(t) + (1 - \lambda) P(s)$

Define $\mu^i \in (0, 1)$ such that $\pi^i = \mu^i p + (1 - \mu^i) q$.

Then $P(p) = \mu^1 P^1 + \mu^2 P^2$, since $\pi = P(p)p + P(q)q = P^1 \pi^1 + P^2 \pi^2$.

We demonstrate that $U \geq V$ in three steps:

(i) $U = P^1 U^1 + P^2 U^2$ where U^i is utility from (p, q) with prior π^i ,

Beyond Binary Dominated Experiment

Let λ be such that $\pi^1 = \frac{P(r)r + \lambda P(s)s}{P(r) + \lambda P(s)}$.

Let $P^1 = P(r) + \lambda P(s)$ and $P^2 = 1 - P^1 = P(t) + (1 - \lambda) P(s)$

Define $\mu^i \in (0, 1)$ such that $\pi^i = \mu^i p + (1 - \mu^i) q$.

Then $P(p) = \mu^1 P^1 + \mu^2 P^2$, since $\pi = P(p)p + P(q)q = P^1 \pi^1 + P^2 \pi^2$.

We demonstrate that $U \geq V$ in three steps:

- (i) $U = P^1 U^1 + P^2 U^2$ where U^i is utility from (p, q) with prior π^i ,
- (ii) $V = P^1 V^1 + P^2 V^2$ where V^1 is utility from (r, s) with prior π^1 and V^2 is utility from (s, t) with prior π^2 , and

Beyond Binary Dominated Experiment

Let λ be such that $\pi^1 = \frac{P(r)r + \lambda P(s)s}{P(r) + \lambda P(s)}$.

Let $P^1 = P(r) + \lambda P(s)$ and $P^2 = 1 - P^1 = P(t) + (1 - \lambda) P(s)$

Define $\mu^i \in (0, 1)$ such that $\pi^i = \mu^i p + (1 - \mu^i) q$.

Then $P(p) = \mu^1 P^1 + \mu^2 P^2$, since $\pi = P(p)p + P(q)q = P^1 \pi^1 + P^2 \pi^2$.

We demonstrate that $U \geq V$ in three steps:

- (i) $U = P^1 U^1 + P^2 U^2$ where U^i is utility from (p, q) with prior π^i ,
- (ii) $V = P^1 V^1 + P^2 V^2$ where V^1 is utility from (r, s) with prior π^1 and V^2 is utility from (s, t) with prior π^2 , and
- (iii) $U^i \geq V^i$.

Steps

(i) Experiment (p, q) with prior π_i has utility $U^i = \mu^i u_p + (1 - \mu^i) u_q$.
 $\Rightarrow U = P^1 U^1 + P^2 U^2$ then follows from $P(p) = P^1 \mu^1 + P^2 \mu^2$.

Steps

(i) Experiment (p, q) with prior π_i has utility $U^i = \mu^i u_p + (1 - \mu^i) u_q$.
 $\Rightarrow U = P^1 U^1 + P^2 U^2$ then follows from $P(p) = P^1 \mu^1 + P^2 \mu^2$.

(ii) Experiment (r, s) with prior π^1 has utility $V^1 = \frac{P(r)u_r + \lambda P(s)u_s}{P^1}$.

Also, $V^2 = \frac{P(t)u_t + (1-\lambda)P(s)u_s}{P^2}$.

$\Rightarrow V = P^1 V^1 + P^2 V^2$ follows from definition of V .

Steps

(i) Experiment (p, q) with prior π_i has utility $U^i = \mu^i u_p + (1 - \mu^i) u_q$.
 $\Rightarrow U = P^1 U^1 + P^2 U^2$ then follows from $P(p) = P^1 \mu^1 + P^2 \mu^2$.

(ii) Experiment (r, s) with prior π^1 has utility $V^1 = \frac{P(r)u_r + \lambda P(s)u_s}{P^1}$.

Also, $V^2 = \frac{P(t)u_t + (1-\lambda)P(s)u_s}{P^2}$.

$\Rightarrow V = P^1 V^1 + P^2 V^2$ follows from definition of V .

(iii) Binary experiment (r, s) with prior π^i delivers utility V^i .

By earlier claim for binary experiments, (p, q) with prior π^i yields

$$U^i \geq V^i.$$

Steps

(i) Experiment (p, q) with prior π_i has utility $U^i = \mu^i u_p + (1 - \mu^i) u_q$.
 $\Rightarrow U = P^1 U^1 + P^2 U^2$ then follows from $P(p) = P^1 \mu^1 + P^2 \mu^2$.

(ii) Experiment (r, s) with prior π^1 has utility $V^1 = \frac{P(r)u_r + \lambda P(s)u_s}{P^1}$.

Also, $V^2 = \frac{P(t)u_t + (1-\lambda)P(s)u_s}{P^2}$.

$\Rightarrow V = P^1 V^1 + P^2 V^2$ follows from definition of V .

(iii) Binary experiment (r, s) with prior π^i delivers utility V^i .

By earlier claim for binary experiments, (p, q) with prior π^i yields

$$U^i \geq V^i.$$

Argument generalizes for dominated experiment with n **posteriors in box**

Literature

Closest literature

- Lehmann (1988) imposed MLR on experiments (improved by QS)
 - Nothing known about case where better (p, q) violates MLR
 - Gets only part of our box (MLR (r, s)) when (p, q) has MLR
- Kim (2021) proposes another comparison
 - Characterization (monotone quasi-garbling) a priori restricts how signals in non-MLR (r, s) should map to actions
 - Kim's notion gives more than our LR box (we have more DM problems)
 - If s must map to high action, it only gives $p >_{LR} s$ and $r >_{LR} q$
 - Far more when also dominating (p, q) violates MLR:
 r LR-above some point on $p - q$ line, s LR-below
 - To recover our result: monotone quasi-garble for any order on signals
- Athey and Levin (2018) likewise: monotone signal-to-action map
 - They generalize Lehmann by permitting other orders than MLR
 - We let experiment's monotonicity depend on decision problem

Why Drop MLR?

Why Drop MLR?

1. The MLR is strong: possible to order states and signals such that $P(s'|\theta)/P(s|\theta)$ rises in θ for every pair $s' > s$
2. Even when the MLR thus holds, some decision makers may have IDO preferences when Θ is differently ordered
 - For example, best state for investment-action when a moderate political party wins
 - But opinion polls are MLR when states are left-right ordered

General Dominating Experiments (in progress)

General Dominating Experiment: In Progress

We are working toward a general characterization

- If P has a least and greatest signal in the LR order, any less informative Q must have its support in the LR-interval between those
- Let us illustrate the issue in belief space


General Dominating Experiment: In Progress

We are working toward a general characterization

- If P has a least and greatest signal in the LR order, any less informative Q must have its support in the LR-interval between those
- Let us illustrate the issue in belief space
- The appropriate generalization of the LR box may lie in the space of $\Pr(a_j|\theta_k)$ profiles rather than beliefs
- Under Blackwell, the set of $\Pr(a_j|\theta_k)$ profiles expands for better experiments
- IDO allows a weaker experiment to have profiles in a slightly larger set, akin to our LR-box

General Dominating Experiment

- Experiment P has support p^1, p^2, p^3 , while Q is supported on the other six points $s^1, s^2, s^3, r^1, r^2, r^3$
- Mass at p^2 can be combined with mass from p^1 to dominate outcomes from Q in the smaller, red box (i.e., r^1, s^1)



Possible Value Functions

Another part of Blackwell's characterization theorem:

- Experiment P is better than experiment Q if and only if, for every prior π , P gives higher expected value than Q for any convex value function

Possible Value Functions

Another part of Blackwell's characterization theorem:

- Experiment P is better than experiment Q if and only if, for every prior π , P gives higher expected value than Q for any convex value function
- For any decision problem, value function is convex and continuous in posterior beliefs
 - conversely, any convex continuous function is a value function for some decision problem

Possible Value Functions

Another part of Blackwell's characterization theorem:

- Experiment P is better than experiment Q if and only if, for every prior π , P gives higher expected value than Q for any convex value function
- For any decision problem, value function is convex and continuous in posterior beliefs
 - conversely, any convex continuous function is a value function for some decision problem
- When restricting attention to monotone problems
 - what other properties do value functions have in addition to convexity?

Conclusion

We explore Blackwell's program to compare experiments

- restricted to IDO decision problems

When the better experiment P is binary

- If P satisfies MLR, it dominates all experiments supported in its LR-interval \Leftrightarrow all its quasi-garblings
- Otherwise, P dominates all Blackwell-garbled experiments \Leftrightarrow all its quasi-garblings

More generally

- Fruitful to explore the feasible profiles of (random) state-action maps
- Possible to characterize IDO-based value functions