Inspections for Decision Makers
(or: you may fool me, but not hurt me)

Federico Echenique and Eran Shmaya

California Institute of Technology

November 10, 2007

Existing result: Inspections are manipulable

- Outcome $\omega \in \Omega$ governed by a prob. distribution μ.
  ($\omega = (z_0, z_1, \ldots)$, a seq. of outcomes)
- A putative *expert* claims the distribution is ν.

Existing result: Inspections are manipulable

- Outcome $\omega \in \Omega$ governed by a prob. distribution μ.
  ($\omega = (z_0, z_1, \ldots)$, a seq. of outcomes)
- A putative *expert* claims the distribution is ν.
- An *inspector* says she believes the expert
  iff ω is in test set $T_\nu \subseteq \Omega$.
- (T_ν) is s.t. true expert passes the test w/prob. 1.

Existing result: Inspections are manipulable

- Outcome $\omega \in \Omega$ governed by a prob. distribution μ.
  ($\omega = (z_0, z_1, \ldots)$, a seq. of outcomes)
- A putative *expert* claims the distribution is ν.
- An *inspector* says she believes the expert
  iff ω is in test set $T_\nu \subseteq \Omega$.
- (T_ν) is s.t. true expert passes the test w/prob. 1.

Result: A false expert can always manipulate the test,
and pass for all ω.
(Foster & Vohra, Lehrer, Olszewski & Sandroni, Shmaya)

We study case where inspector cares about ν because she has to make a decision.

Inspector is a DM

- A putative *expert* claims the distribution is ν.
- DM cares about ω: she makes a *decision* $a \in A$; her payoff depends on (ω, a).

Inspector is a DM

- A putative *expert* claims the distribution is ν.
- DM cares about ω: she makes a *decision* $a \in A$;
  her payoff depends on (ω, a).

$$\text{DM} \begin{cases} \text{believe } \nu & \longrightarrow \text{ use } a^*_\nu, \text{ optimal action for } \nu \\ \text{reject } \nu & \longrightarrow \text{ use } a^*_\pi, \text{ optimal action for } \pi. \end{cases}$$

π is the DM's existing belief about Ω

We compare

$$\text{Payoff}(\omega, a^*_\nu) - \text{Payoff}(\omega, a^*_\pi)$$

under two criteria: ν and π.

Our test

We show: there is a test (T_ν) s.t.

- true expert passes the test w/prob. 1.

Our test

We show: there is a test (T_ν) s.t.

- true expert passes the test w/prob. 1.
- DM can follow expert recommendation (i.e. choose a^*_ν). *regardless* of whether the expert is true.

Our test

We show: there is a test (T_ν) s.t.

- true expert passes the test w/prob. 1.
- DM can follow expert recommendation (i.e. choose a^*_ν).
  regardless of whether the expert is true.
  - if DM thinks expert is true, she should choose a^*_ν

Our test

We show: there is a test (T_ν) s.t.

- true expert passes the test w/prob. 1.
- DM can follow expert recommendation (i.e. choose a^*_ν). *regardless* of whether the expert is true.
  - if DM thinks expert is true, she should choose a^*_ν
  - if DM thinks expert is false, under beliefs π, ν will not pass the test when a^*_ν is bad for DM.

(...if DM is patient enough)

Our test

We show: there is a test (T_ν) s.t.

- true expert passes the test w/prob. 1.
- DM can follow expert recommendation (i.e. choose a^*_ν). *regardless* of whether the expert is true.
  - if DM thinks expert is true, she should choose a^*_ν
  - if DM thinks expert is false, under beliefs π, ν will not pass the test when a^*_ν is bad for DM.

(...if DM is patient enough)

Let me say it again.

Model

Model: (Z, A, r, λ, π).

- Z finite or countable set.
- Ω: infinite sequences $z_1, z_2, \ldots$ in Z.
- A : set of actions.
- $r : Z \times A \to [0, 1]$: payoff function.
- $\lambda \in (0, 1)$: discount factor.
- $\pi \in \Delta(\Omega)$: beliefs.

Model

A *test* is a function $T : \Delta(\Omega) \to$ subsets of Ω.

A test T is *type-I error free* if $\nu(T(\nu)) = 1 \ \forall \nu \in \Delta(\Omega)$.

Model

A *strategy* for DM is $f : Z^{<\mathbb{N}} \to A$. $f(z_0, \ldots, z_{n-1})$ is the action taken by the DM after observing $(z_0, \ldots, z_{n-1})$.

Model

A *strategy* for DM is $f : Z^{<\mathbb{N}} \to A$. $f(z_0, \ldots, z_{n-1})$ is the action taken by the DM after observing $(z_0, \ldots, z_{n-1})$.

Payoff:

$$R_\lambda(\omega, f) = (1 - \lambda) \sum_{n \in \mathbb{N}} \lambda^n r\left(z_n, f(z_0, \ldots, z_{n-1})\right)$$

from strategy f and outcome $\omega = (z_0, z_1, \ldots)$.

f is ν-optimal iff

$$f \in \operatorname{argmax} \int R_\lambda(x, g)\nu(\mathrm{d}x).$$

Theorem

There exists a type-I error free test T s.t.

$$\lim_{\lambda\to 1}\int_{T(\nu)} \left(R_\lambda(\omega,g) - R_\lambda(\omega,f)\right)\ \pi(d\omega) \leq 0$$

for every $\nu \in \Delta(X)$ and every ν-optimal strategy f and π-optimal strategy g.

Merging

$\pi, \nu \in \Delta(\Omega)$.
ν *merges with* π if

$$\lim_{n\to\infty} d(\pi_{\omega|_n}, \nu_{\omega|_n}) = 0,$$

with π-prob. 1.

Merging

$\pi, \nu \in \Delta(\Omega)$.
ν *merges with* π if

$$\lim_{n\to\infty} d(\pi_{\omega|_n}, \nu_{\omega|_n}) = 0,$$

with π-prob. 1.

π is abs. cont. w.r.t. ν if $\nu(A) = 0 \Rightarrow \pi(A) = 0$.

Proposition (Blackwell-Dubins Theorem)

If π is abs. cont. w.r.t. ν, then ν merges with π

Our test

There is T_ν s.t. $\nu(T_\nu) = 1$,
and $(\nu(A) = 0 \Rightarrow \pi(A) = 0)$ on T_ν.
exists by application of Lebesgue's Decomposition Theorem.

Then, on T_ν, ν merges with π.
If λ is large enough, payoffs under π are close.

Our test

Turns out:

$$T_\nu = \left\{ \omega : \limsup_{n\to\infty} \frac{\pi(z_1, \dots z_n)}{\nu(z_1, \dots z_n)} < \infty \right\}$$

(a "likelihood ratio" test).

Related work.

- Olszewski & Sandroni
- Al-Najjar & Weinstein
- Feinberg & Stewart