Learning From The Data: A Theory Without Guessing

Karl H. Schlag

August 25, 2025

1 Using Information When Making Choices

Setting:
- Uncertainty about the true enviornment.
- Information provided in the form of a data set about past outcomes.

We show:
Generically, Bayesians can perform on average worse with the data than without it.
We say that they can be caught guessing.

Why? Because they treat the true environment as a random draw from their prior.

We call a rule non-guessing if in the true enviornment it performs on average better when conditioning on the data than when chosing a random action.

E.g. choose between a, a'. Payoffs in $[0, 10]$. Data set: $a : 1, 3, 5$, $a' : 4, 10$.

2 Non-Guessing Methodology

Choice of an action from A, here $A = \{a, a'\}$.
Fix payoff distribution (unknown) of each with support in $[0, 10]$.

A randomized choice $q \in \Delta A$ serves as benchmark (called pseudo-prior).
eg choose each action equally likely.
eg choose according to frequency of recommendations by outsiders.

Add data set consisting of some past independent realizations of each action.

Take a rule that specifies choice given pseudo-prior and data set.

Rule is non-guessing if exp payoffs under rule $\geq$ exp payoffs of pseudo-prior.

Note: To choose q regardless of data set is trivially non-guessing.
Note: Sequential evaluation of non-guessing rules is non-guessing.

3 The Math

$A = \{a, a'\}$
action $a'' \in A$ yields payoff drawn from unknown distribution $P_{a''} \in \mathcal{P}$ ($\mathcal{P}$ known).

Let $\mu_{a''} = E_{P_{a''}}(z)$ mean payoff of action $a'' \in A$

Fix $m_a, m_{a'} \in \mathbb{N}$

Data set $x = \left\{(a, y_i)_{i=1}^{m_a}\right\} \cup \left\{\left(a', z_j\right)_{j=1}^{m_{a'}}\right\}$
where y_i and z_j indep drawn from P_a and $P_{a'}$ respectively

Rule $f : R^{m_a+m_{a'}} \rightarrow \Delta A$

$E_P(f) = \int f(x)\, dP_a(y_1) \cdot ... \cdot dP_a(y_{m_a}) \cdot dP_{a'}(z_1) \cdot ... \cdot dP_{a'}\left(z_{m_{a'}}\right) \in \Delta A$

Rule f is non-guessing given $q \in \Delta A$ if $E_P(f) \cdot \mu \geq q \cdot \mu \ \forall\ P_a, P_{a'} \in \mathcal{P}$.

4 Initial Insight

Lemma: Any non-guessing rule is continuous in the payoffs in the data set, so its choice is randomized in many data sets.

Corollary: Both empiricist and generic Bayesian can be caught guessing.

5 Our Finding

Focus on linear rules (simple, easy to compute performance).

Theorem: We find a linear non-guessing rule that dominates any other linear non-guessing rule.

Here is the rule for a balanced data set:

(i) If only payoffs $\{0, 10\}$ in data set then choose action that yielded more 10's.

(ii) When payoffs in $(0, 10)$:
Transform each payoff $x \in (0, 10)$ into 10 and 0 with probability $\frac{x}{10}$ and $1 - \frac{x}{10}$, and then go back to (i).

Rule for non-balanced data set: replace (i) by
(i'): choose action a if sufficiently many 10's, otherwise choose a'.

6 Back to our Example

Two actions a and a'.

Payoffs drawn from $[0, 10]$.

Data set: action a payoffs $1, 3, 5$, action a' payoffs $4, 10$.

Choose a and a' with probabilities 0.196 and 0.804 respectively.

7 Other Settings

- Many actions and balanced data set.
- One uncertain action.

Same qualitative findings as under two actions.

8 Good at Learning Which Action is Best?

Def: Rule finds the best action in large data sets if it chooses a best action with arbitrarily high probability provided each action is sampled sufficiently often.

Theorem: Both Empiricist and dominant linear non-guessing rule find the best action in large data sets.

Not clear whether true for Bayesian.

Proposition: In any balanced sample, dominant linear non-guessing rule under uniform pseudo-prior maximizes minimal probability of choosing best action for a given minimal difference between the two means.

Moreover, it is best rule with this property!!!!

9 Possible Applications

decision making using noisy data:
- treatment choice (eg whom to give mosquito nets to given a RCT)
- A/B testing (eg improving web page design)
- which worker to promote
- choosing a price given market data
- choosing which firm has highest revenue

decision or game theory based on noisy signals: (not this paper)

10 Related Literature

model:
treatment choice (Manski, 2004), A/B testing, bandit setting
$\neq$ inference about means, $\neq$ dynamic learning from own choices

related concepts:
- absolute expediency = exp payoffs must increase when learning from own previous choice (Lakshmivarahan & Thathachar, 1973, Börgers etal, 2004)
- improving = exp payoffs must increase when learning from others (Schlag, 1998)
- UMPU test as uncovered in this paper

explicit algorithms:
- minimax regret for two actions (approximate if sample unbalanced)
- hypothesis testing
- others?

11 Conclusion

... compared to Bayesian

- choices get simpler as you know less (one universal rule)
- performance is easier to compute
- no loss even if data set large

... compared to empiricist and classical hypothesis testing

- choices are good in any sample size

... compared to minimax regret

- math easier
- methodology allows for bias
- methodology allows for learning

I shot the Bayesian ... but I did not shoot the deputy