Optimal Decision-Making Architecture

H. Moulin[1] A. S. Nesterov[2] D. Tsilevich[2]

[1]Adam Smith Business School
University of Glasgow, UK

[2]International Laboratory of Game Theory and Decision Making
HSE SPB

Introduction

What is common among the three pictures?

Step-by-step structure.

Introduction

We will mainly focus on the justice-system example: independent evaluation at each stage. At each stage, the case can be passed on or closed:

$$\text{Investigation} \longrightarrow \text{Initial hearing} \longrightarrow \text{Appellate court} \longrightarrow \ \ldots$$

But is it always like this?

$$\cancel{\text{Investigation}} \longrightarrow \cancel{\text{Initial hearing}} \longrightarrow \text{Three-Judge Court (Troika - 1936)}$$

Introduction

Research Question

What is the optimal design of a Decision-Making System?

- **High operational costs.** Each stage takes time and time is money.
- **Accuracy-efficiency trade-off**. Reducing costs by limiting reviews results in lower accuracy.

Literature Review

- *CJT*: Uniformity and heterogeneity[Boland, 1989]; Validness [Berend and Paroush, 1998] and correlation between votes[Kaniovski, 2010];
- *Optimal Stopping*: Series of r.v. that lead to a decision based on realization [Siegmund, 1967];
- Judicial Design: Single-stage[Cameron and Kornhauser, 2006]; Appeal as a tool for reducing judicial errors [Shavell, 1995]; Instance that reviews the decisions of others [Siegel and Strulovici, 2023];
- Miscellaneous: Medical Screening[Özekici and Pliska, 1991]; Inspections[Chen et al., 1998].

General Framework: Testers

The goal is to correctly decide Y or N. The decision is made by passing trough stages, that can either **pass on** ($\hat{Y}$) or **drop** ($\hat{N}$) the case. We model stages as **testers**:

Definition 1

Tester T_i is a non-strategic probabilistic element with preassigned levels of sensitivity $Pr(\hat{Y}|Y) = p_i$ and specificity $Pr(\hat{N}|N) = p_i'$, $0.5 < p_i, p_i' < 1$. $T = \{T_1, \ldots, T_K\}$ is a set of available testers.

Depending on the true state of the world ($Pr(Y) = q$; $Pr(N) = 1 - q$), four outcomes are possible: $\hat{Y}|Y$, $\hat{Y}|N$ or *Type I Error*, $\hat{N}|N$, $\hat{N}|Y$ or *Type II Error*.

General Framework: Architectures

From testers, we want to construct a decision-making system. We denote different systems as different **architectures**:

Definition 2

Let **architecture** $\mathcal{V} : (k_1, k_1^*) \longrightarrow \cdots \longrightarrow (k_L, k_L^*)$ be the decision-making system consisting of K testers distributed across L consecutive stages with k_i testers per stage i, s.t. $\sum_{i=1}^{L} k_i = K$. k^* is the threshold on how many testers have to agree on $\hat{Y}$ in order for the decision to be $\hat{Y}$.

General Framework: Examples

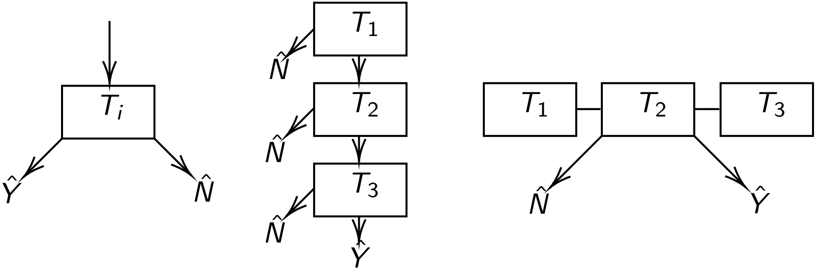

(a) :Tester T_i (b) :Sequential architecture (c) :Parallel architecture

Figure: Tester T_i and examples of trivial architectures.

For $K = 4$, it is possible to create 5 unique structures: fully-sequential, fully-parallel, $\mathcal{V} : (3, k_1^*) \longrightarrow (1, 1)$, $\mathcal{V} : (2, k_1^*) \longrightarrow (1, 1) \longrightarrow (1, 1)$ and $\mathcal{V} : (2, k_1^*) \longrightarrow (2, k_2^*)$ - intermediate.

General Framework: Costs I

We propose a constraint on how many testers can be involved per case **on average**: at any stage there is an opportunity for a case to be dropped, so the probability of triggering the next stage k_i is $\neq 1$. For a sequential-system:

$$Pr(k_i) = \begin{cases} 1, i = 1 \\ q \prod_{j=1}^{i-1} p_j + (1-q) \prod_{j=1}^{i-1} (1-p_j), i > 1 \end{cases}$$

For a general system:

$$Pr(k_i) = q \prod_{j=1}^{i-1} \sum_{n=k^*}^{k_j} \binom{k_j}{n} p^n (1-p)^{k_j - n} + (1-q) \prod_{j=1}^{i-1} \sum_{n=k^*}^{k_j} \binom{k_j}{n} p^{k_j - n} (1-p)^n \tag{1}$$

General Framework Costs II

Having established the probability of reaching the i-th stage of the decision-making process, we can estimate the expected number of testers per stage, $E(k_i) = k_i \times E[Pr(k_i)]$. Assuming $E(q) = E(1-q) = 0.5$, i.e. equal probabilities of the states:

$$E(k_i) = 0.5k_i[\prod_{j=1}^{i-1} \sum_{n=k^*}^{k_j} \binom{k_j}{n} p^n (1-p)^{k_j - n} + \prod_{j=1}^{i-1} \sum_{n=k^*}^{k_j} \binom{k_j}{n} p^{k_j - n} (1-p)^n]$$

Then,

$$E(K|\mathcal{V}, T, q) = \sum_{i=1}^{L} E(k_i) = k_1 + \sum_{i=2}^{L} E(k_i), L \geq 2. \quad (2)$$

General Framework: Problem

In this paper we define the objective function $\mathcal{E}$ as the weighted sum of Type I and Type II Errors - the goal of the system is to minimize the probability of making a mistake.

$$\mathcal{E} = \lambda Err_1 + (1 - \lambda) Err_2 \tag{3}$$

We want to find optimal architecture $\mathcal{V}^*$:

$$\mathcal{V}^* = \arg\min_{\mathcal{V}^* \in V} \mathcal{E}, \tag{4}$$

where V is set of all architectures $\mathcal{V}$ s.t. $E(K|\mathcal{V}, T, q) \leq \overline{c}$. Here, $\overline{c}$ is upper limit on how many testers we are willing to place **on average**.

Assumptions

1. *Symmetry* Tester T_i is *symmetric* in accuracy if $TPR_i = TNR_i$:

$$Pr(\hat{Y}|Y) = Pr(\hat{N}|N) = p_i$$

2. *Uniformity* Testers in set $T = \{T_1, \dots, T_K\}$ are *uniform* in terms of their respective accuracy if sensitivity and specificity is the same across all testers:

$$p_1 = p_2 = \cdots = p_K = p$$

$$p'_1 = p'_2 = \cdots = p'_K = p'$$

3. *Independence* Any two adjacent testers T_i and T_{i-1} in a system are *independent* if the decision of the first tester does not influence the decision of the second tester:

$$Pr(\hat{Y}_i \cap \hat{Y}_{i-1}|Y) = Pr(\hat{Y}_i|Y) \times Pr(\hat{Y}_{Y-1}|Y) = p_i \times p_{i-1}$$

Sequential Case

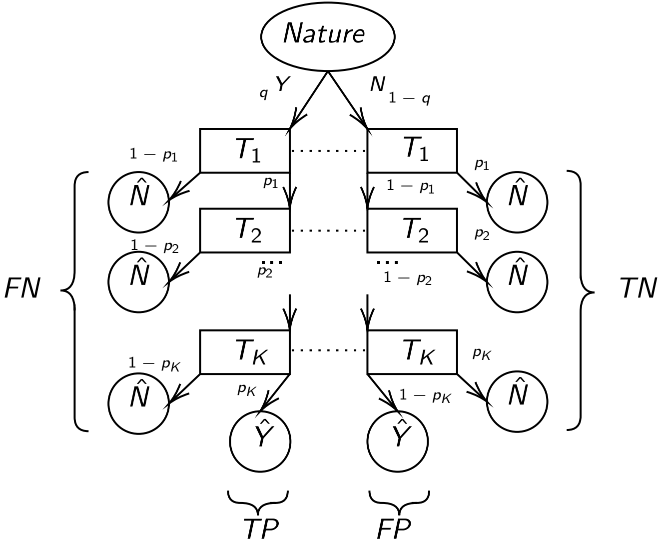

Figure: The Decision Tree: Sequential case

Sequential Case

Lemma 1

Let $T = \{T_1, T_2, \ldots, T_K\}$ be the set of symmetric, uniform, independent tester. Then, the optimal length L^* of the sequential decision-making tree can be calculated with the following formula:

$$L^* = \begin{cases} 1 & , \lambda = 0 \\ min[max(\lfloor \frac{\ln(\frac{\lambda}{1-\lambda} \; \frac{\ln(1-p)}{\ln p})}{\ln \frac{p}{1-p}} \rfloor, 1), K] & , \lambda \in (0; 1) \\ K & , \lambda = 1 \end{cases} \tag{5}$$

Parallel Case

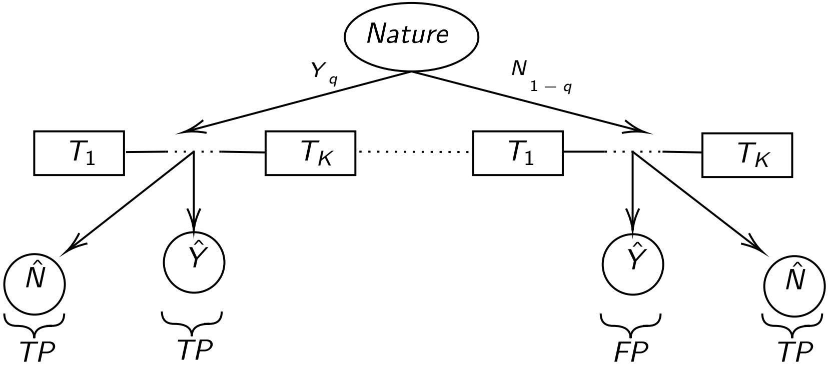

Figure: The Decision Tree: Parallel case

Parallel Case

Lemma 3

For an arbitrary parallel architecture $\mathcal{V} : (K, k^*)$, the number $k^*$of symmetric, uniform and independent testers k^* that have to agree on $\hat{Y}$ in order for the final verdict $\hat{Y}$ to be optimal:

$$k^* = \begin{cases} 1 & , \lambda = 0 \\ \lfloor \frac{K}{2} + 1 - \frac{\ln \frac{1-\lambda}{\lambda}}{2 \ln \frac{p}{1-p}} \rfloor & , \lambda \in (0, 1) \\ K & , \lambda = 1 \end{cases} \tag{6}$$

General Case: $\overline{c} = 2$

Case for $\overline{c}$=2 is minimally-viable, since it produces all types of architecture: Parallel $\mathcal{V} : (2, 2), \mathcal{V} : (2, 1)$; Sequential $\mathcal{V} : (1, 1) \longrightarrow \ldots$; Pyramid $\mathcal{V} : (1, 1) \longrightarrow (2, 1), \mathcal{V} : (1, 1) \longrightarrow (2, 2)$

Lemma 2

For a constraint $\overline{c} = 2$, the maximal length K of a Sequential-Decision Making process is such that:

$$p^{K+1} + (1 - p)^{K+1} = (2p + 1)^2, p \in (0.5; 1) \tag{7}$$

Then,

$$\mathcal{K} = \frac{2\ln(2p - 1) + \ln(p)}{\ln(p)}, \ |K - \mathcal{K}| < 1 \tag{8}$$

For $p \in [0.55; 1), \lfloor K \rfloor = \lfloor \mathcal{K} \rfloor$;

As $p \longrightarrow 0.5, K\&\mathcal{K} \longrightarrow \infty$; $p \longrightarrow 1, K\&\mathcal{K} \longrightarrow 3$

General Case, $\overline{c} = 2$

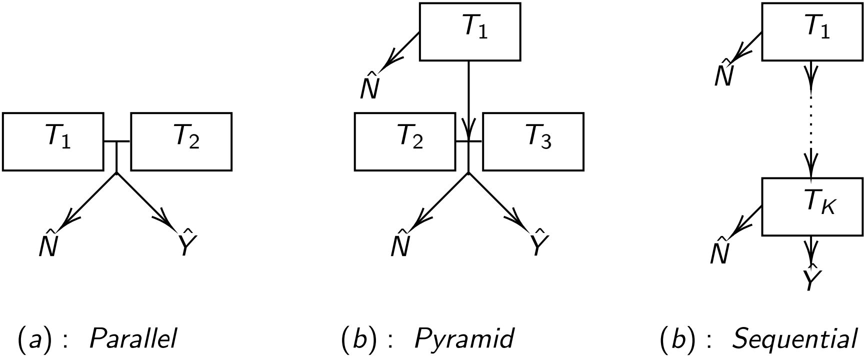

Figure: Available Architectures

General Case, $\overline{c} = 2$

Theorem 1

For a non-trivial decision-making system ($\lambda \in (0;1)$) comprised of symmetric, uniform, independent testers of the expected length of utmost two, the optimal architecture $\mathcal{V}^*$ is:

$$\mathcal{V}^* = \begin{cases} \mathcal{V}: (2,1) & , \lambda \leq \frac{2-p}{3} \\ \mathcal{V}: (1,1) \longrightarrow (2,1) & , \frac{2-p}{3} < \lambda < p \\ \mathcal{V}: (1,1) \underrightarrow{L^*} \dots & , \lambda \geq p \end{cases} \tag{9}$$

where

$$L^* = min(\lfloor \frac{\ln(\frac{\lambda}{1-\lambda} \frac{\ln(1-p)}{\ln p})}{\ln \frac{p}{1-p}} \rfloor, \lfloor \frac{2\ln(2p-1) + \ln(p)}{\ln(p)} \rfloor)$$

General Case, $\overline{c} = 2$

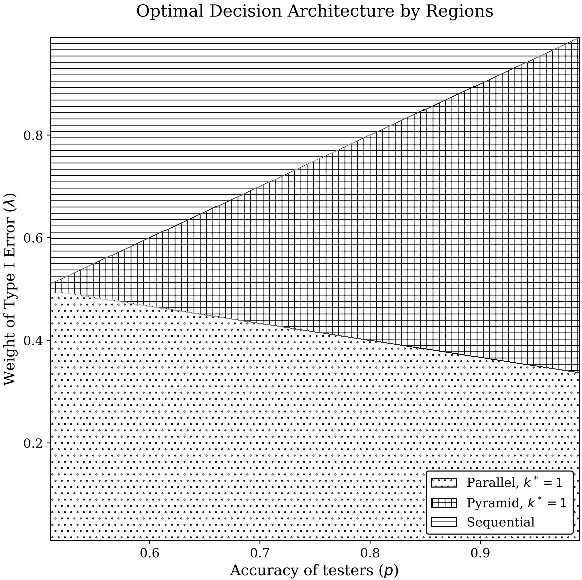

Conclusion

1. **Each type** (sequential, parallel, intermediate) **is used** for some combination of "softness" and "accuracy".
2. Even in the case of $\bar{c} = 2$, Sequential Architecture can be (very close to) infinite in length - **but is often not**.
3. If we value the outcomes equally, **pyramids** are the way to go **for any accuracy**.
4. Sequential Juries - **same in accuracy, but cheaper**.

Berend, D. and Paroush, J. (1998).
When is condorcet's jury theorem valid?
Social Choice and Welfare, 15(4):481–488.

Boland, P. J. (1989).
Majority systems and the condorcet jury theorem.
Journal of the Royal Statistical Society Series D: The Statistician, 38(3):181–189.

Cameron, C. M. and Kornhauser, L. A. (2006).
Appeals mechanisms, litigant selection, and the structure of judicial hierarchies.
Institutional Games and the U.S. Supreme Court, 173:173–204.

Chen, J., Yao, D. D., and Zheng, S. (1998).
Quality control for products supplied with warranty.
Operations Research, 46(1):107–115.

Kaniovski, S. (2010).
Aggregation of correlated votes and condorcet's jury theorem.
Theory and Decision, 69(3):453–468.

Özekici, S. and Pliska, S. R. (1991).
Optimal scheduling of inspections: A delayed markov model with false positives and negatives.
Operations Research, 39(2):261–273.

Shavell, S. (1995).
The appeals process as a means of error correction.
The Journal of Legal Studies, 24(2):379–426.

Siegel, R. and Strulovici, B. (2023).
Judicial mechanism design.
American Economic Journal: Microeconomics, 15(3):243–70.

Siegmund, D. O. (1967).
Some problems in the theory of optimal stopping rules.
The Annals of Mathematical Statistics, pages 1627–1640.