Warwick Statistics

WARWICK
THE UNIVERSITY OF WARWICK

Decision models and real world applications

Julia Brettschneider

Warwick, 3.3.2020

Content

- Traditional Markov decision processes and their limitations
- Joint decisions with multiple perspectives
  - Ownership & influence
  - Utility
  - Awareness
  - Further directions
- Applications
  - UG admissions
  - Cancer treatment
  - Agri-environment

Decision trees

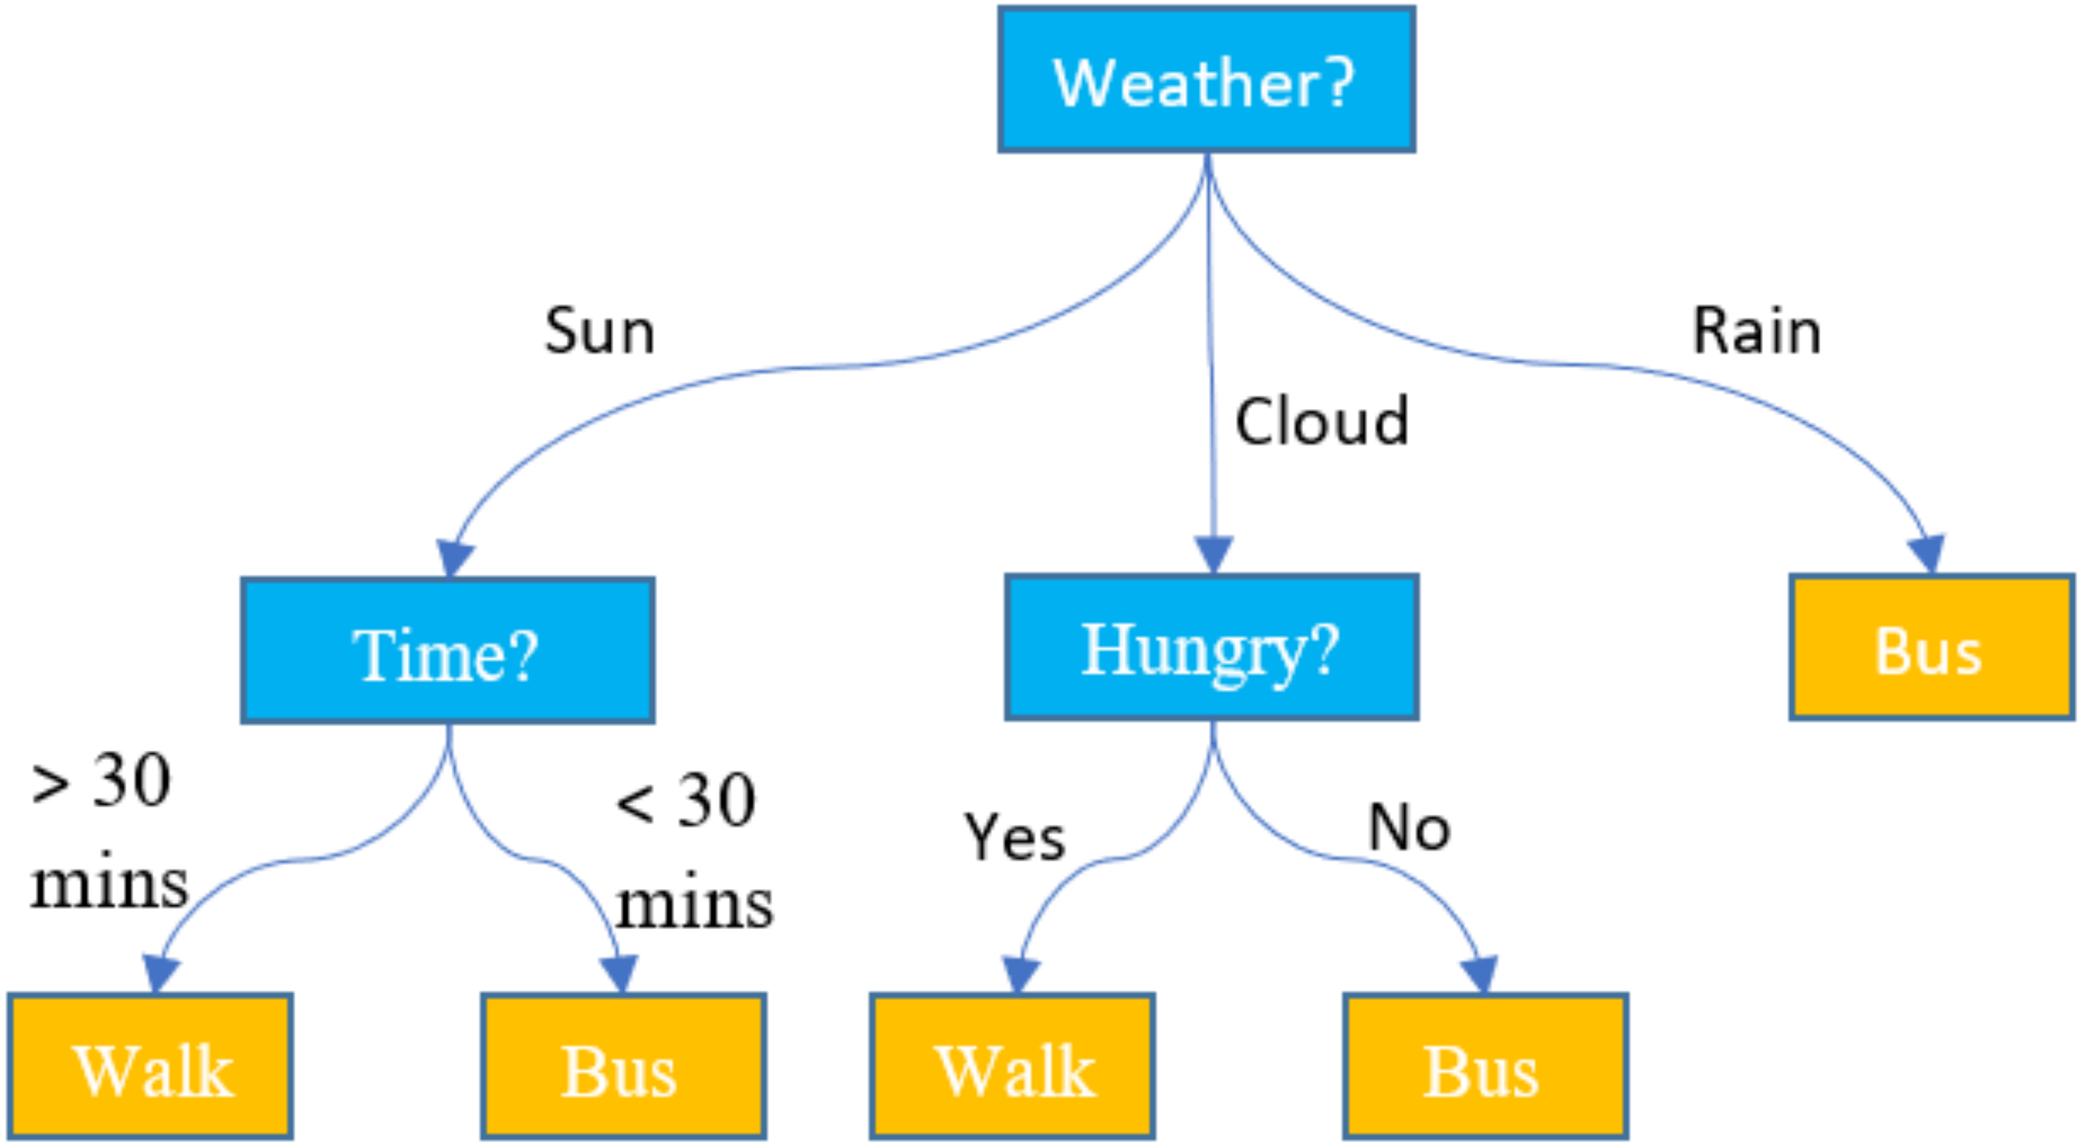

Decision rule?

You are *not* a decision *maker*, you just follow the chart.

Decision trees

Decision trees with *uncertainty*

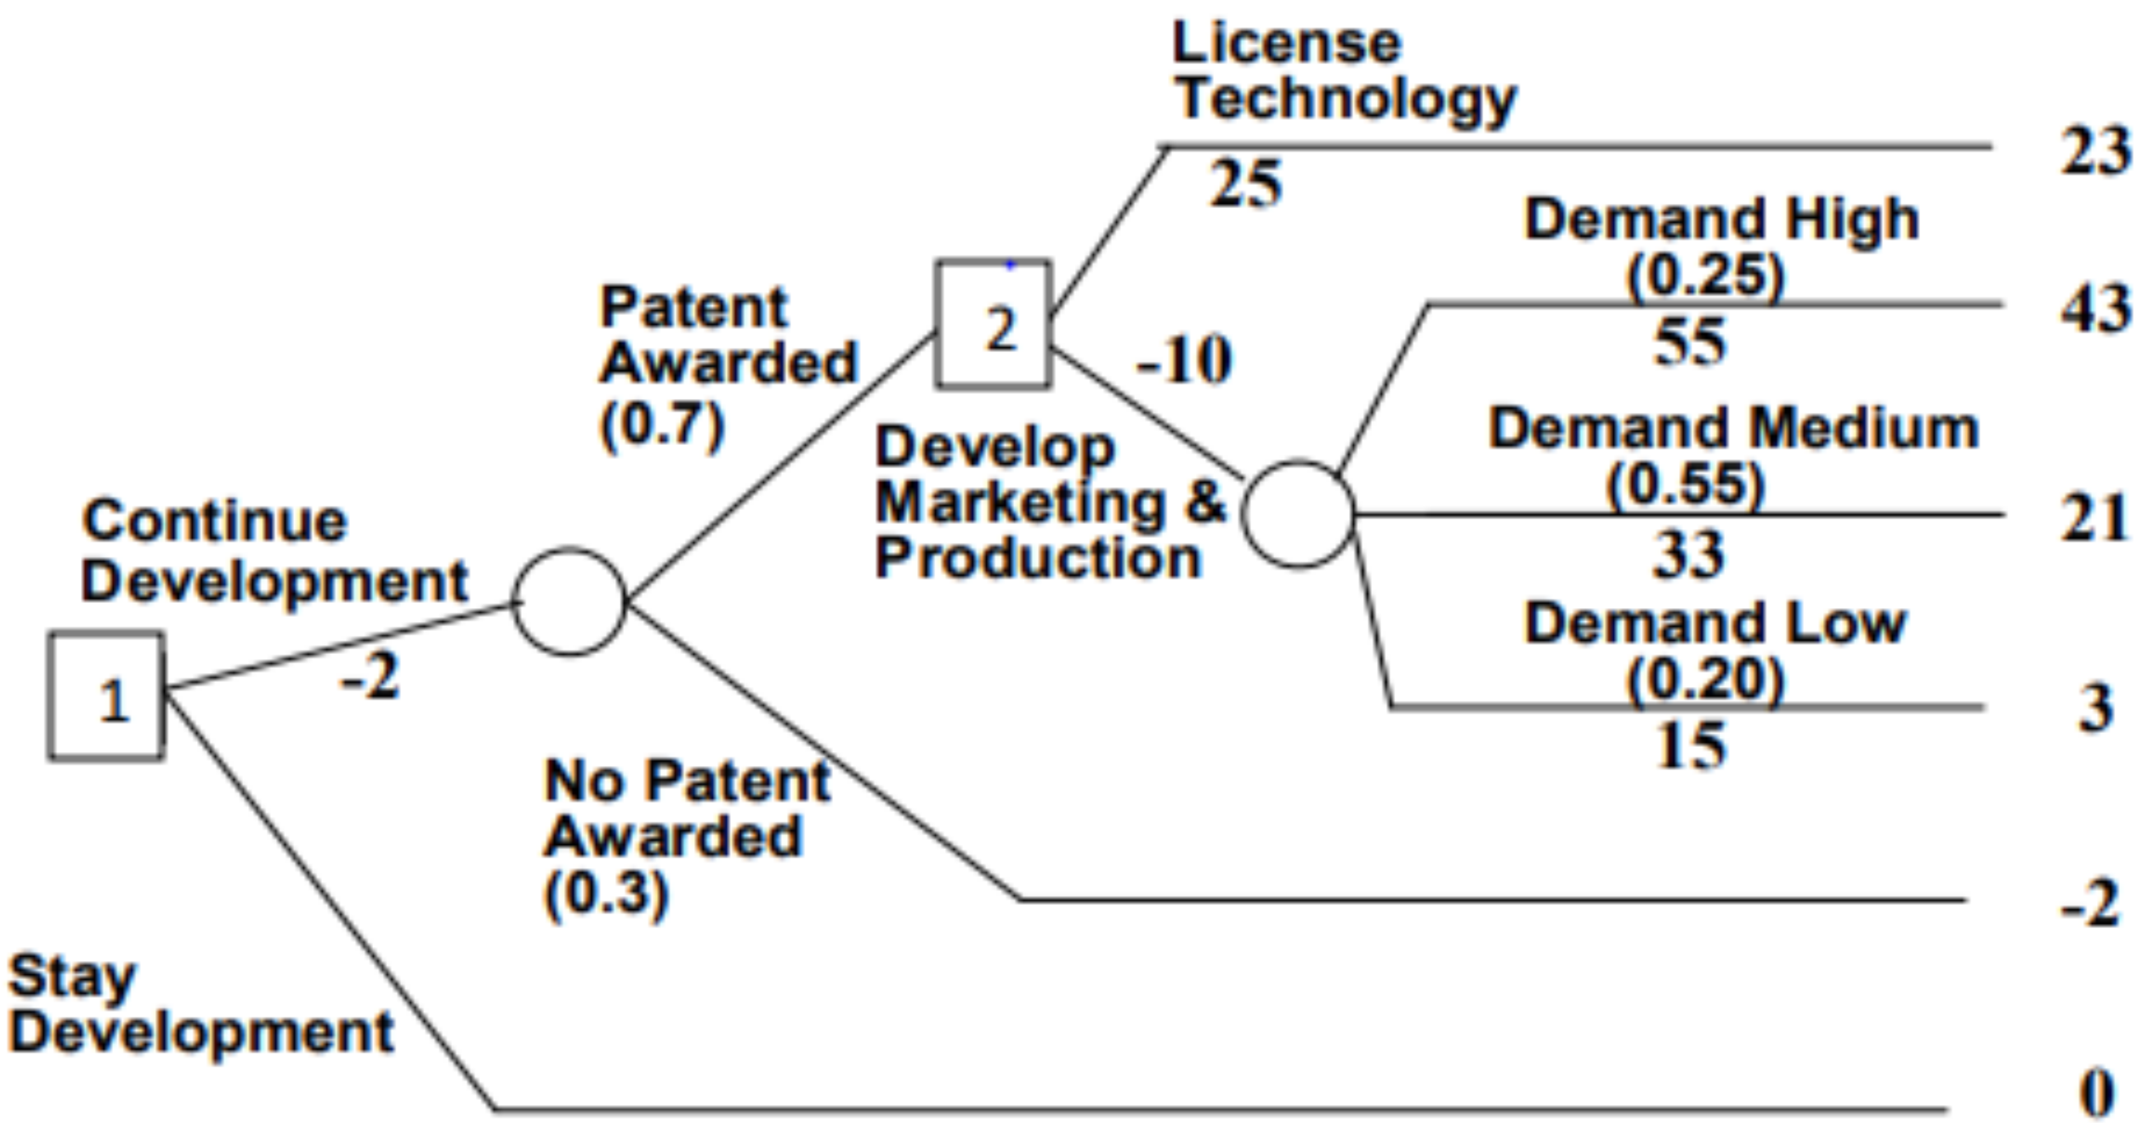

Decision rule:

Choose option that *maximises expected pay-off*

Markov decision processes (MDPs)

$\mathcal{M} = (T, A, \Theta, R)$

Dynamic system under partial control of DM

$\sigma = S_0, \ldots, S_\tau$ **Subsequent states**

$\alpha = a_0, a_1, \ldots, a_\tau$ **Action sequence**

$$P_\tau^{(S,\alpha)}(\sigma) = \prod_{t=0}^{\tau-1} \theta_t(S_t, a_t, S_{t+1})$$

Decision node
Operated by Decision Maker (DM)

Chance node
Operated by prob. distr. (in Economics: "nature")

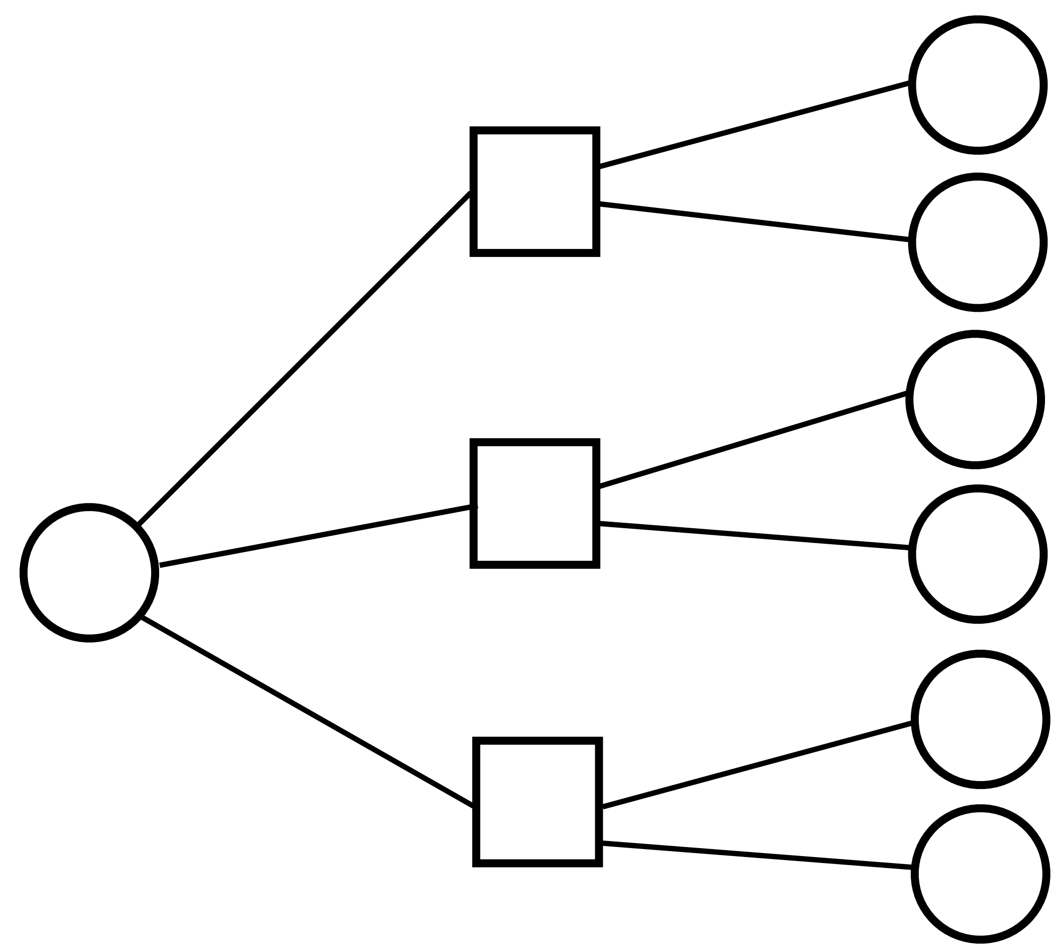

Markov decision processes (MDPs)

$\mathcal{M} = (T, A, \Theta, R)$

Dynamic system under partial control of DM

$\sigma = S_0, \ldots, S_\tau$ **Subsequent states**

$\alpha = a_0, a_1, \ldots, a_\tau$ **Action sequence**

$$P_\tau^{(S,\alpha)}(\sigma) = \prod_{t=0}^{\tau-1} \theta_t(S_t, a_t, S_{t+1})$$

$$h = (S_0, \ldots, S_N, a_0, \ldots, a_N)$$

$$u(h) = \sum_{t=0}^{N} \lambda^t r_t(S_t, a_t)$$ **Utility**

$$P_\tau^{(S,\pi)}(h) = \prod_{t=0}^{\tau-1} \theta_t(S_t, a_t, S_{t+1})$$

Markov decision processes (MDPs)

$\mathcal{M} = (T, A, \Theta, R)$

Dynamic system under partial control of DM

$\sigma = S_0, \ldots, S_\tau$ **Subsequent states**

$\alpha = a_0, a_1, \ldots, a_\tau$ **Action sequence**

$$P_\tau^{(S,\alpha)}(\sigma) = \prod_{t=0}^{\tau-1} \theta_t(S_t, a_t, S_{t+1})$$

$$h = (S_0, \ldots, S_N, a_0, \ldots, a_N)$$

$$u(h) = \sum_{t=0}^{N} \lambda^t r_t(S_t, a_t)$$ **Utility**

$$P_\tau^{(S,\pi)}(h) = \prod_{t=0}^{\tau-1} \theta_t(S_t, a_t, S_{t+1})$$

Decision rule:

Maximise the expected utility

$$E_{P_N^{(S,\pi)}}(u) = \sum_{h \in H_N} u(h) \cdot P_N^{(S,\pi)}(h)$$

Limitations of MDPs

- One decision maker only
- Fixed order of alternating decision and chance nodes
- Fixed utility representing only one perspective
- Utilities numerical (real world outcomes may be incommensurable)
- Utilities not fully multivariate

$$u(h) \quad = \quad \sum_{t=0}^{N} \lambda^t r_t(S_t, a_t)$$

- Probabilities not fully path dependent (Markov)

$$P_{\tau}^{(S,\alpha)}(\sigma) \quad = \quad \prod_{t=0}^{\tau-1} \theta_t(S_t, a_t, S_{t+1})$$

- Knowledge of all probabilities required
- No time varying or state dependent covariates

Example: Genomic Testing

Hiiiiiih! How are you today?

Good? Well, I've got some GREAT news for you. You've got a gene variant that means you've only got a 10% chance for Parkinson's disease before you turn 40 and even after that it only increases by 20% annually.

Sorry for the inconvenience caused, but we thought you'd appreciate to know that beforehand….. See YAH

23andMe®

Genomic Testing: Scenarios

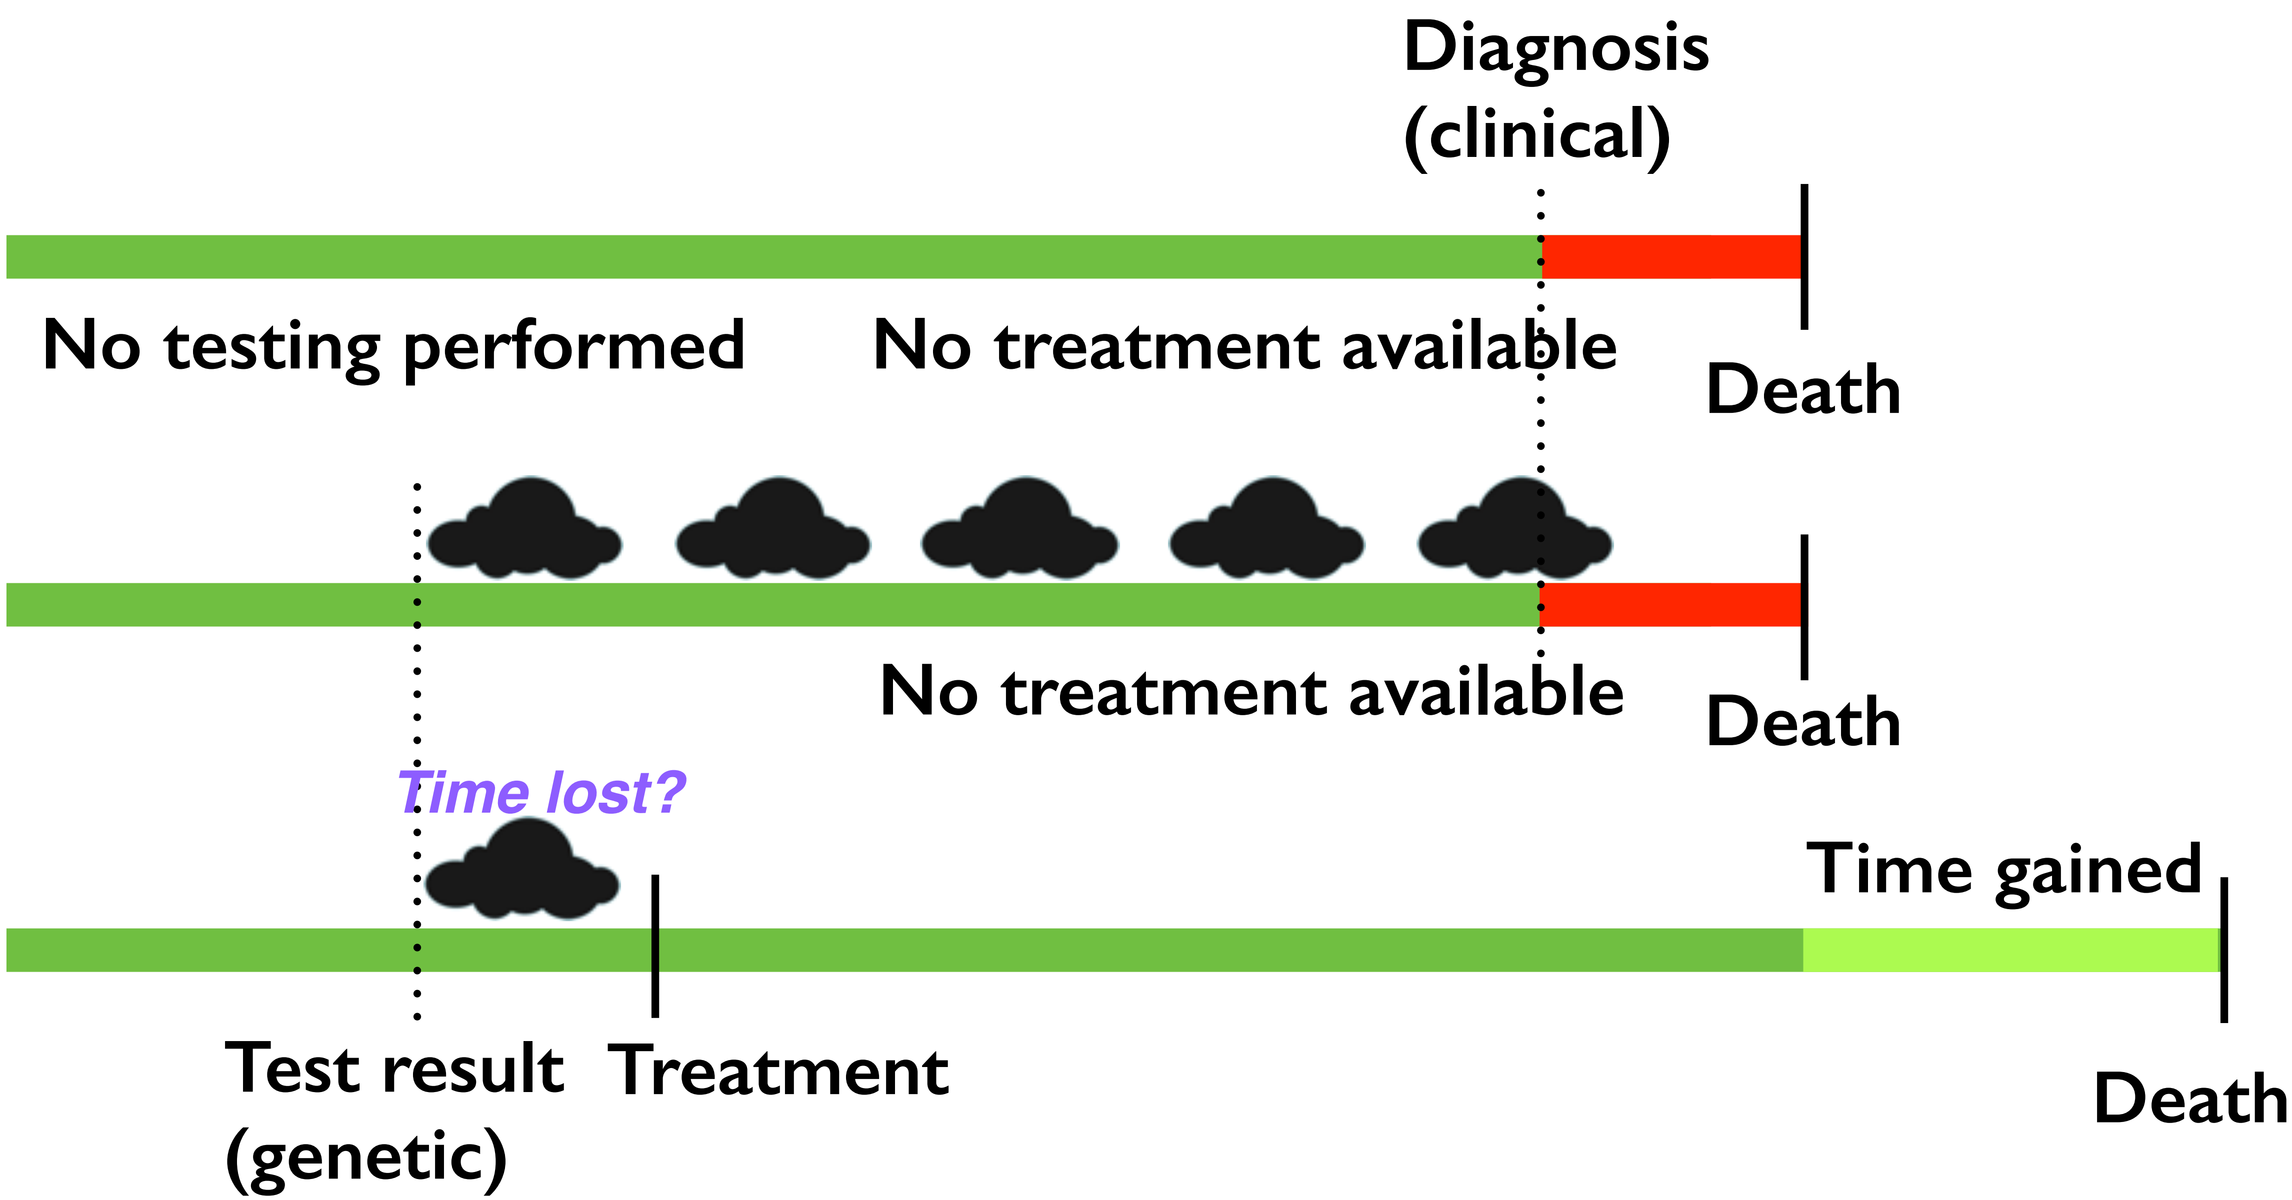

- Test result may cause anxiety and apathy during lead time (could be 50 years!)
- Test result may be wrong (e.g. immature research, multiple testing)
- If correct and if treatment is available testing may increase survival time.

Decision to take a test

- How do we compare outcomes? Cost of lost years?
- What is the loss for living with bad prospects?
- Consider probability weighting (Tversky/Kaneman's Prospect theory)
- Consider cost for others (relatives) who may ***not*** have asked for the information

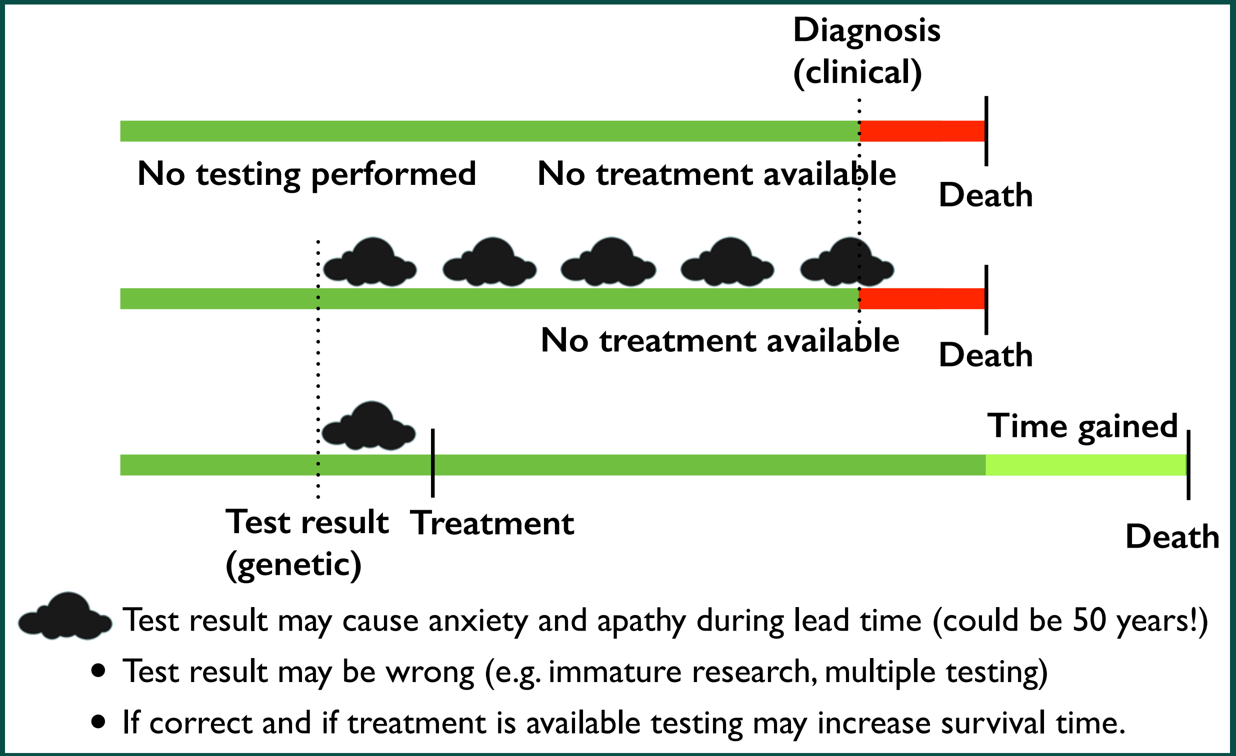

Example: Breast Cancer prognosis

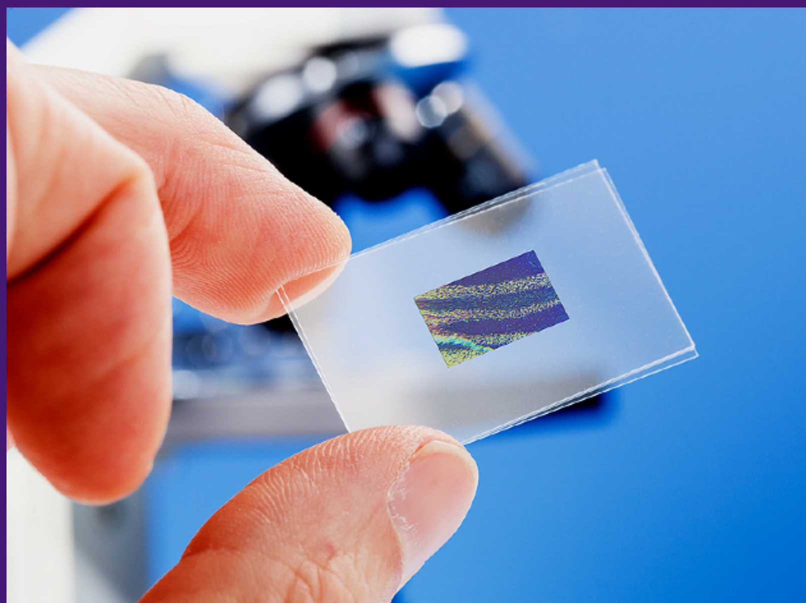

Oncotype DX®:
multigene diagnostic test that determines the individual risk of cancer recurrence in early-stage invasive breast cancer

Reveals the underlying tumour biology on the molecular level to help guide treatment decisions (adjuvant chemotherapy or not)

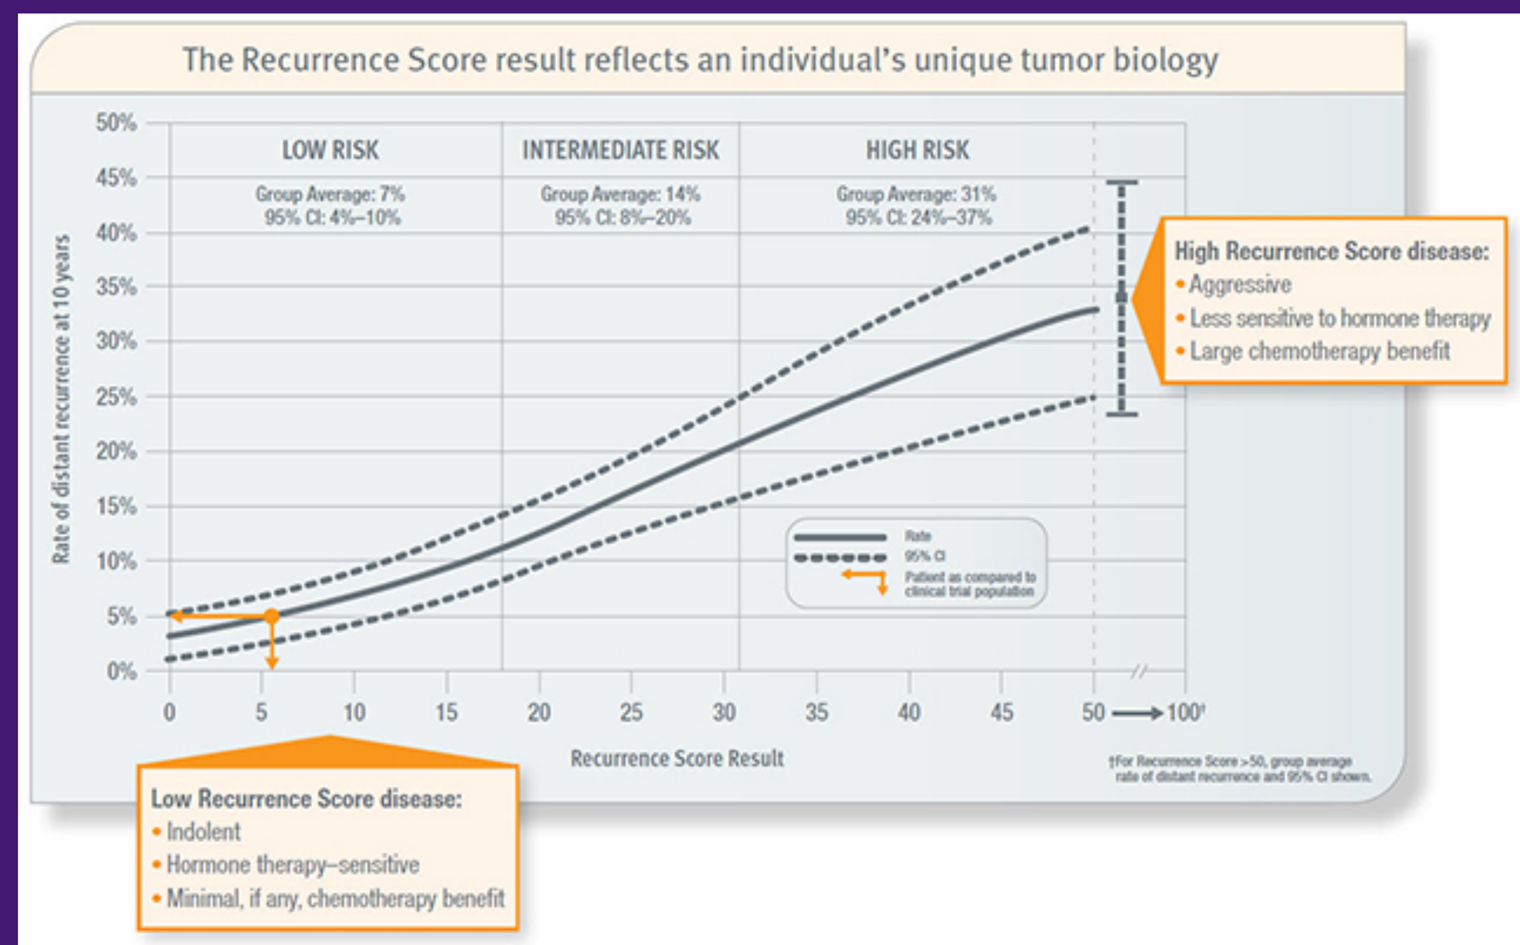

Medical Treatment Decision

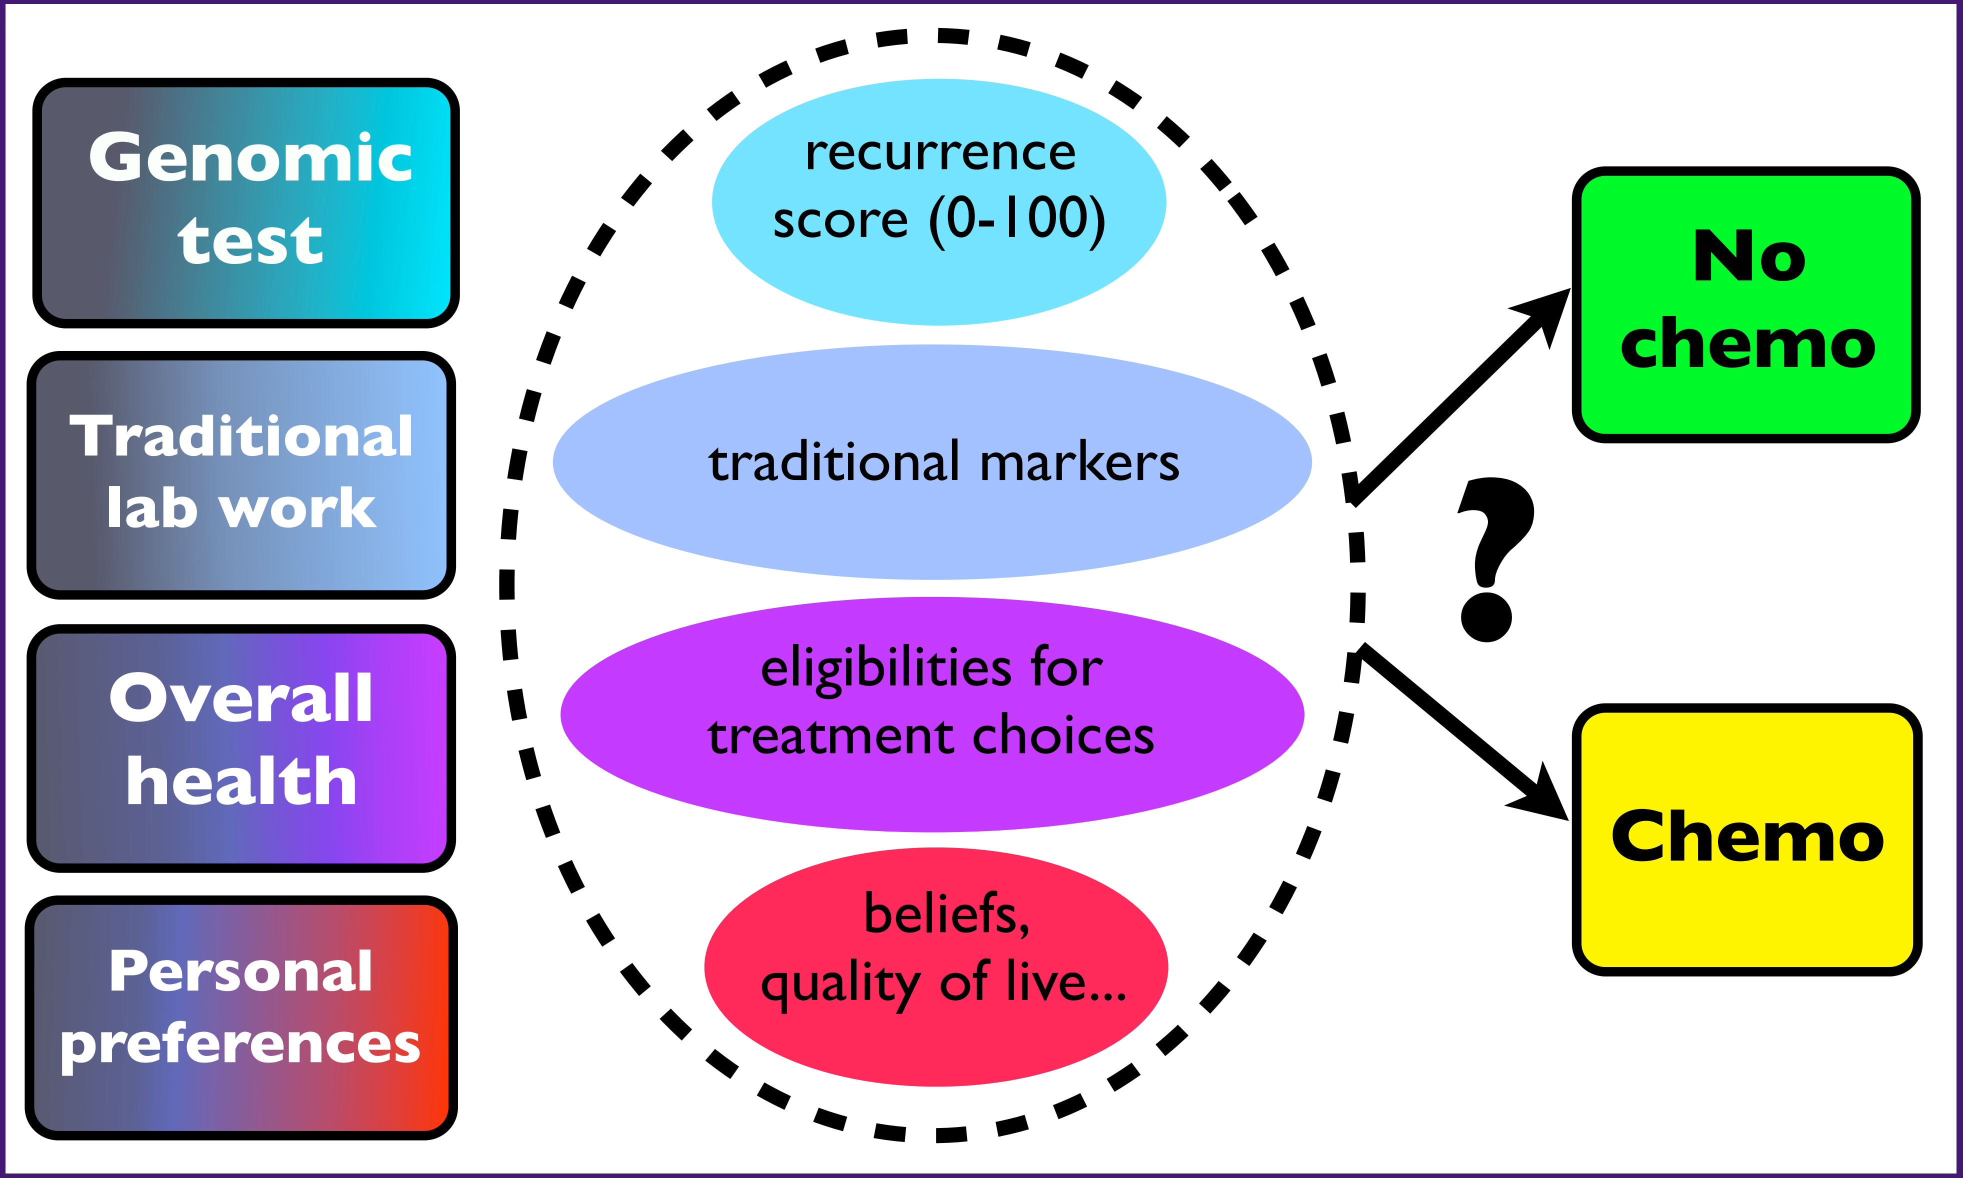

- Complex information with uncertainty (Oncotype DX)
- Emotions interfering with judgement
- Multiple decision makers interacting (physicians, patients, family/friends)

UG Admissions (unique to UK!)

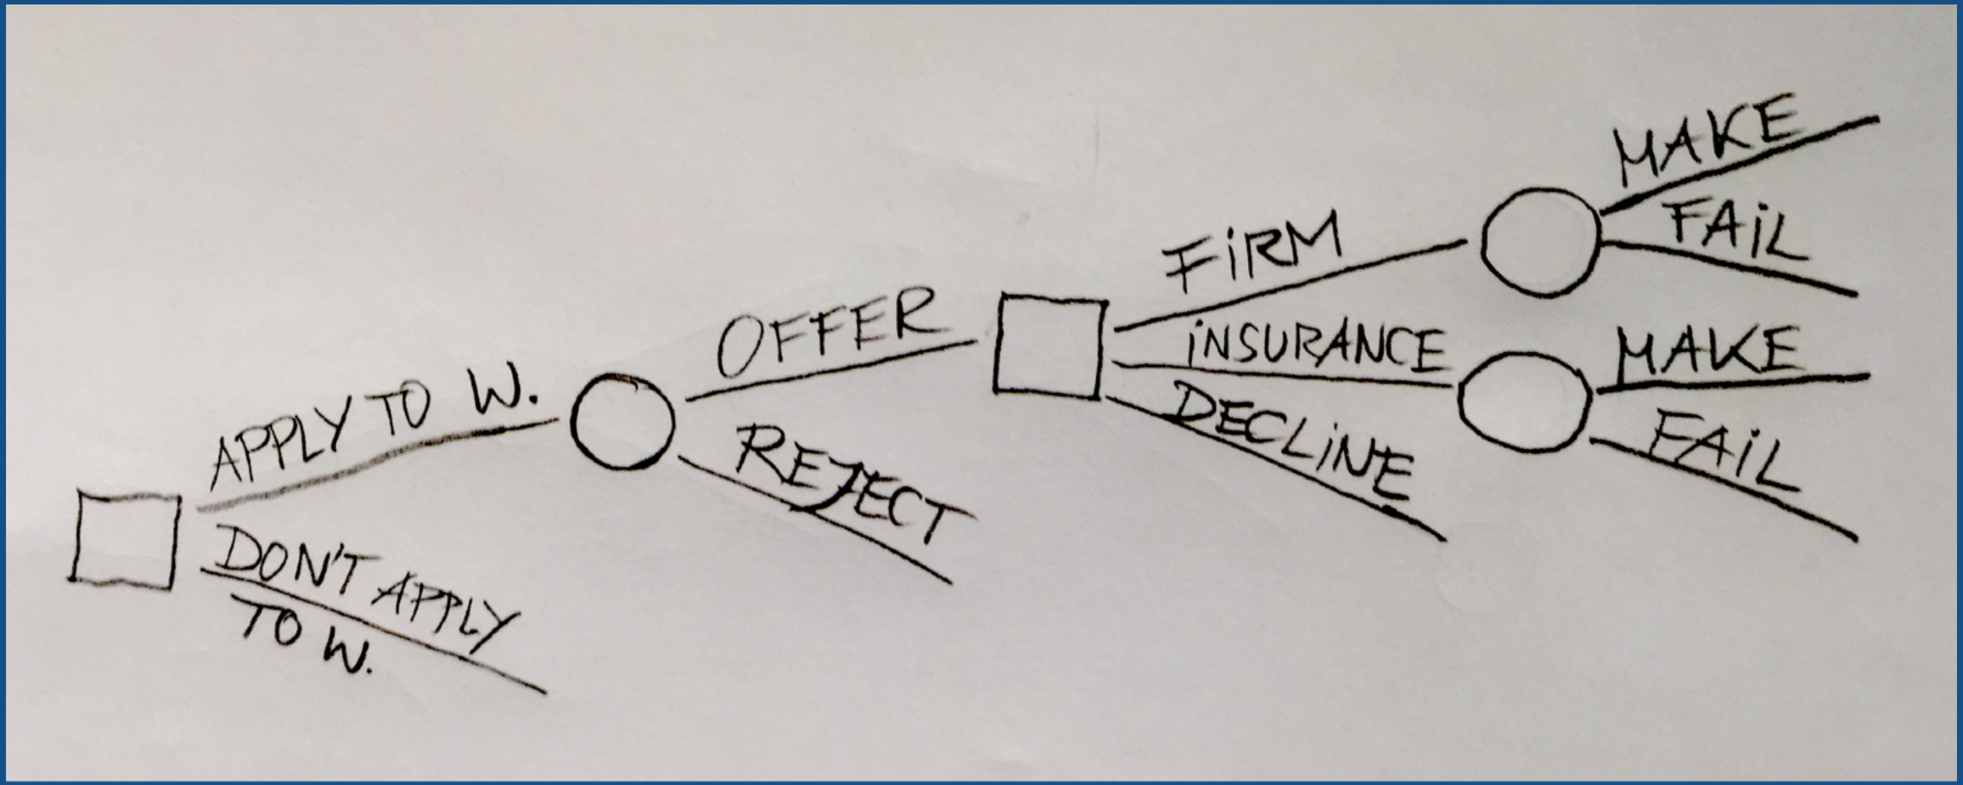

Students receive offers conditional on their A-level results. Decision who gets a conditional offer is based on predicted A-level results, previous marks, recommendation letters, etc.

Perspectives in UG Admissions

Student Perspective

University perspective

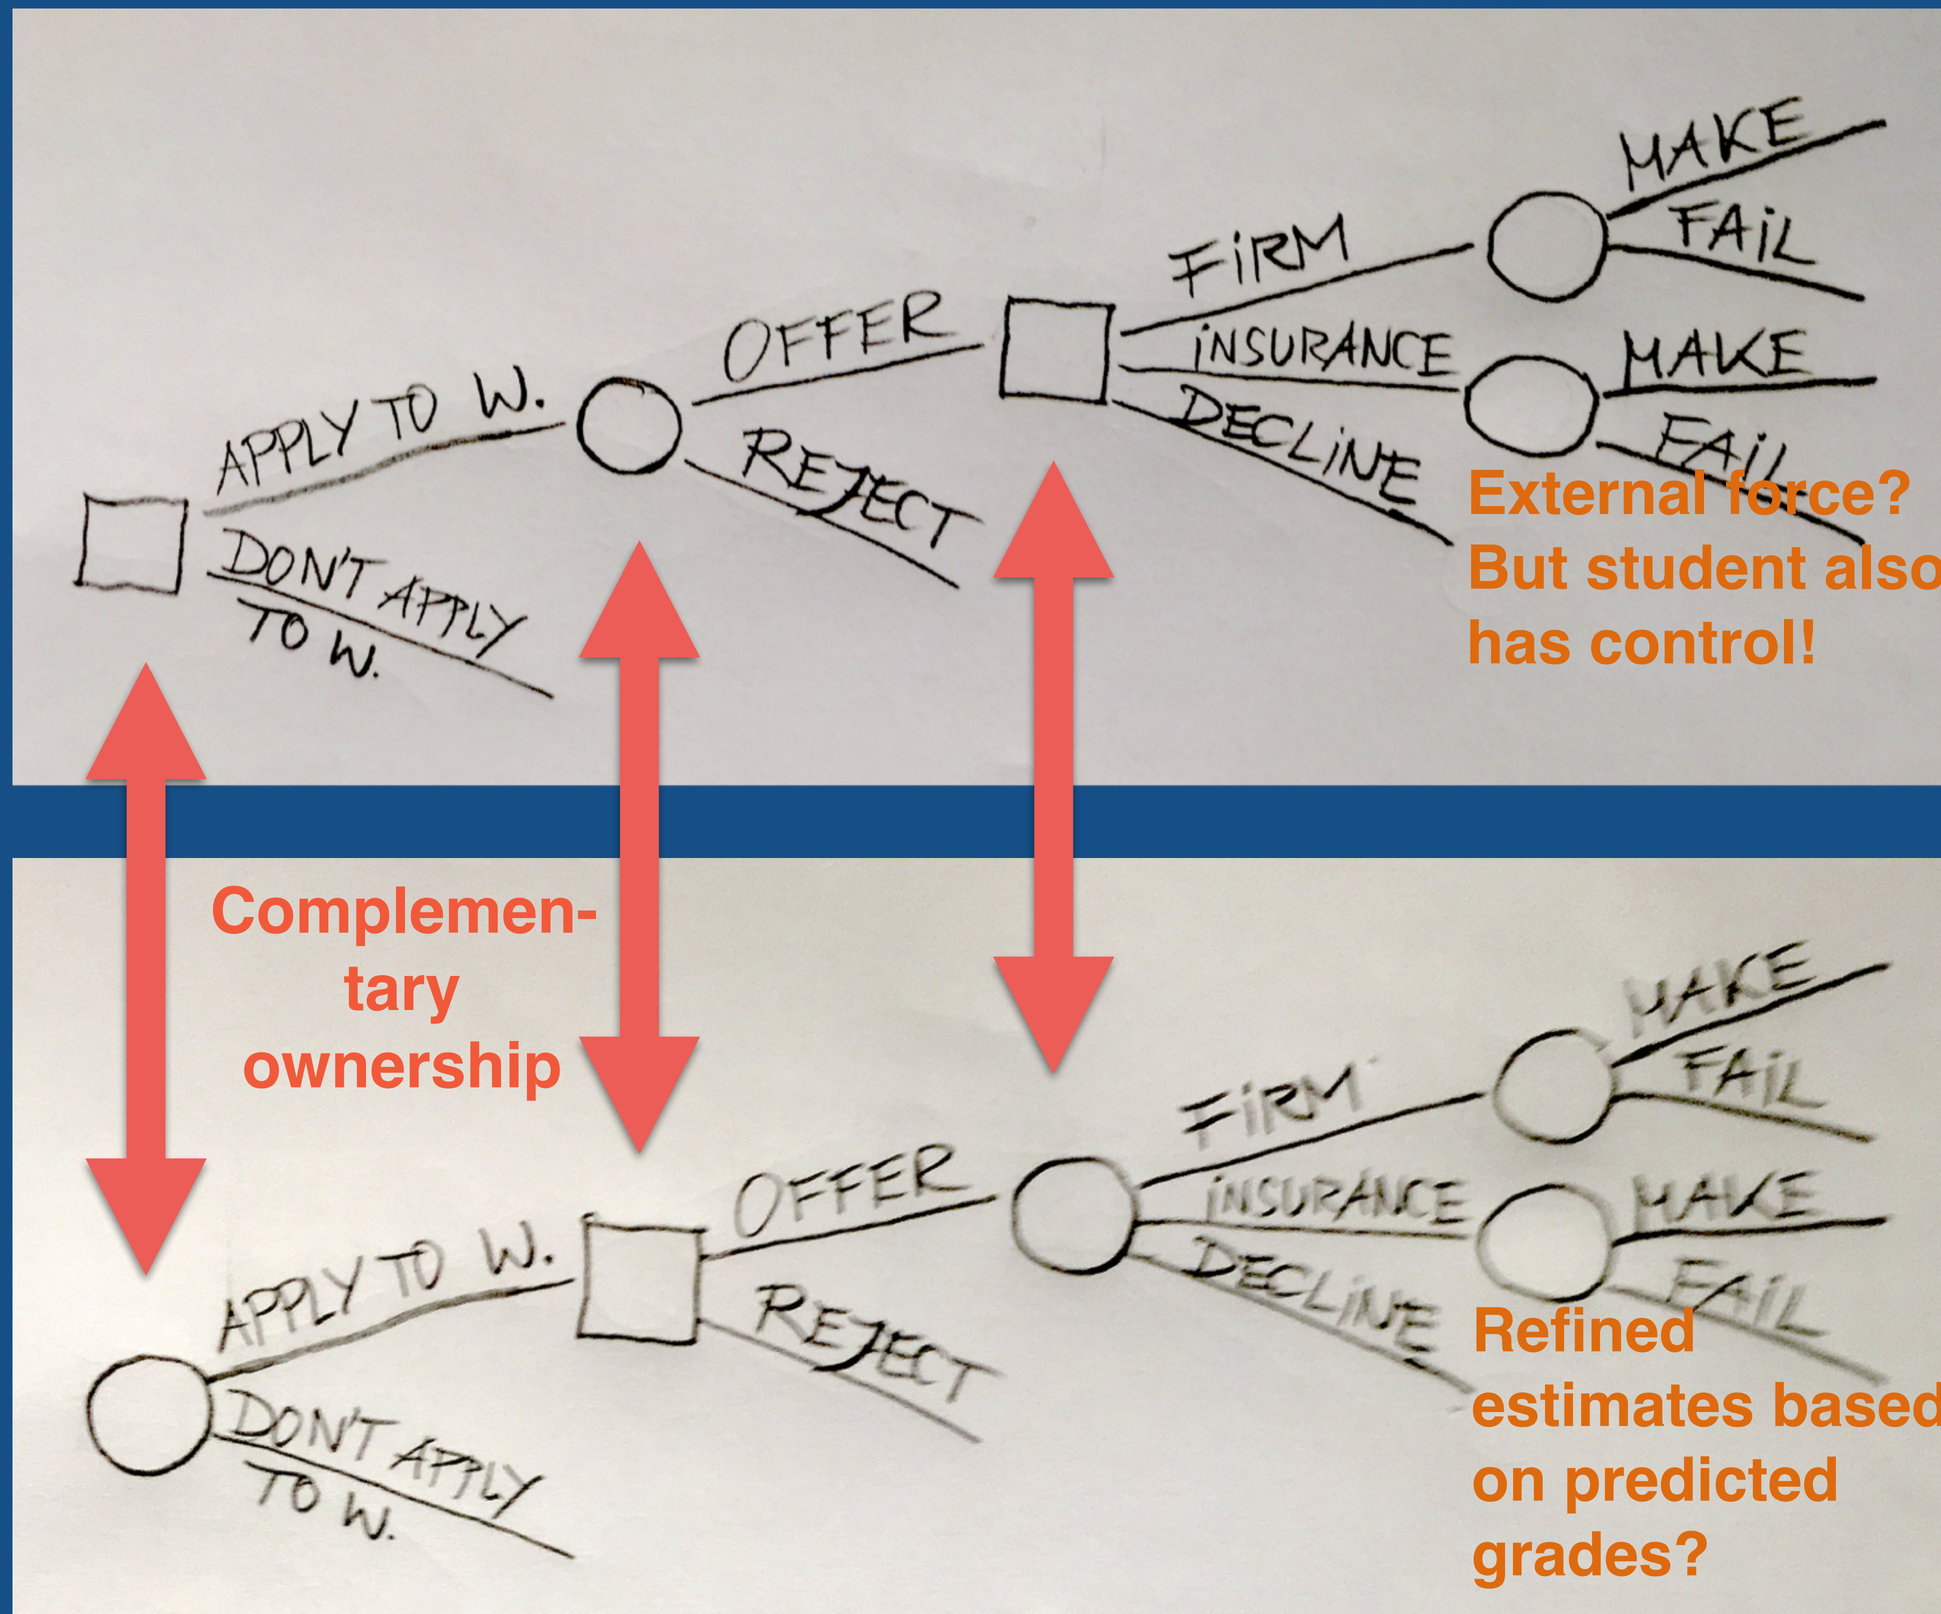

Generalised decision trees

- Multiple decision maker who act together/coordinate choice
- Flexible order of nodes
- Flexible and path-dependent utility
- Non-numerical outcomes, rank-based utilities
- Path dependencies (not Markov)
- Not all information about probabilities (states) available
- Multiple perspectives
- Multiple times scales for evaluation (decision rules)

Game tree

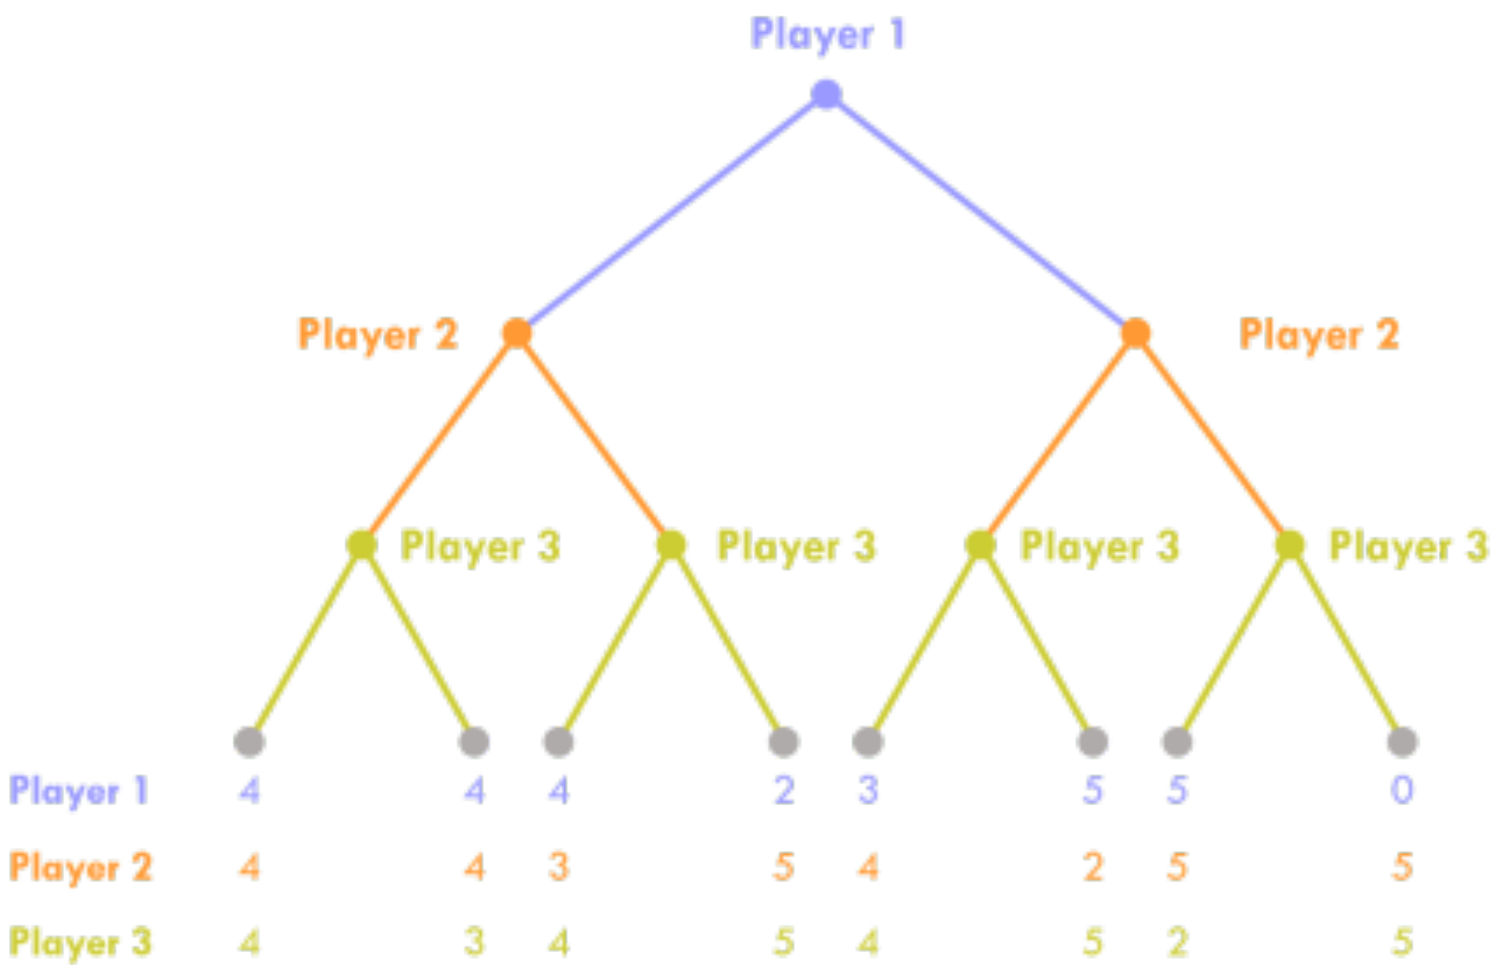

Simultaneous move pay-off matrix

		Player 2	
		Left	Right
Player 1	Up	4 , 3	-1 , -1
	Down	0 , 0	3 , 4

- Moves (choices) of each players at each stage, part of strategy
- Outcomes, pay-offs
- Simultaneous or sequential moves

Tree notation

Mathematical notation for the set $\mathcal{T}$,
i.e. the connected rooted graphs w/o cycles (aka trees):

$$\mathbb{T} = \rho \cup \bigcup_{n \in \mathbb{N}} \mathbb{N}^n \text{ all possible individuals of such trees}$$

$$g_n(T) = \{x \in T \,|\, g(x) = n\} \text{ be the } n\text{th generation of the tree}$$

$$\textit{mother map } m : g_n(\mathbb{T}) \longmapsto g_{n-1}(\mathbb{T})$$

$$C(x) = \{y \in T \,|\, m(y) = x\} \text{ is the set of children of } x$$

$$\text{vertex } x \in T \text{ is called } \textit{leaf} \text{ if } |C(x)| = 0$$

$$T' = T \setminus L(T) = \{x \in T \,|\, |C(x)| > 0\}$$

Decision Owner

Definition Decision owner and control tree. *Let* $T \in \mathcal{T}$ *be a decision tree and* $\beta : T' \longrightarrow B$ *a map assigning each decision point* x *a decision owner* $b(x)$. $\beta^T := \beta(T')$ *is called* ownership tree *generated by* β.

Traditional examples

- full control
- MDPs
- sequential games

More general examples

- DM may depend on path, not only on step

Decision Owner: Example

Definition _. Decision owner and control tree. *Let* $T \in \mathcal{T}$ *be a decision tree and* $\beta : T' \longrightarrow B$ *a map assigning each decision point* x *a decision owner* $b(x)$. $\beta^T := \beta(T')$ *is called* ownership tree *generated by* β.

Example Full control. *If there is a decision maker* $b \in B$ *such that* $\beta(x) = b$ *for all* $x \in T'$ *then* b *fully owns or has full control over the decision process.*

Decision Owner: Example

Definition **Decision owner and control tree.** *Let* $T \in \mathcal{T}$ *be a decision tree and* $\beta : T' \longrightarrow B$ *a map assigning each decision point* x *a decision owner* $b(x)$. $\beta^T := \beta(T')$ *is called* ownership tree *generated by* β.

Example **Traditional 2-person sequential game.** N_1 *are be odd numbers and* N_2 *are the even numbers and Let* $B = \{b_1, b_2\}$ *and for* $x \in T'$ *let*

$$\beta(x) = \begin{cases} b_1 & \textit{if } x \textit{ is odd,} \\ b_2 & \textit{if } x \textit{ is even.} \end{cases}$$

Sequential Control

Definition Control of a step. *Let* $n \in \mathbb{N}_0$*. If there is a* $b \in B$ *such that* $\beta(x) = b$ *for all* $x \in g_n(T')$*, than* b controls the nth step *of the decision process.*

Definition Sequential control. *If for any* $n = 0, \ldots, ht(T')$ *there is a decision maker* $b \in B$ *such that* b *controls the* n*th step then the decision process is* sequentially controlled.

Example Generalised sequential game. *Let* $N_j \subset \mathbb{N}_0, j \in \{1, \ldots, J\}$*, be a partition of* $\mathbb{N}_0$*,* $B = \{b_1, \ldots, b_J\}$ *and for* $x \in T'$ *let* $\beta(x) = b_j$ *if* $x \in g_{N_j}$*.*

Complementary Ownership

Ownership does not have to be tied to the step of the decision process.

Definition Complementary ownership. *Let B be a set of decision makers and $b_1, b_2 \in B$. Two decision makers b_1, b_2 own complementary parts of a decision process if*

$$\beta(x) = b_1 \quad \iff \quad \beta(x) \neq b_2 \quad \text{for all } x \in T'.$$

Example: Most stages in UG admissions in the UK

Example: Conditional offer

Example *Let $T \in \mathcal{T}$ with $ht(T) = 3$. Let b_1 and b_2 be human decision makers and b_3 be an external force. In the first step of this decision process, β_2 decides whether or not to make a conditional offer to β_1. In the second step of the decision process, β_1 decides whether or not to accept it. In the third step, an external force decides whether or not the condition of the offer is fulfilled. The control tree is given by*

$$\beta(x) = \begin{cases} b_1 & \textit{for all } x \in \rho, \\ b_2 & \textit{for all } x \in g_1(T'), \\ b_3 & \textit{for all } x \in g_2(T'). \end{cases}$$

All DMs are complementary to each other.

Shafer’s Decision Trees

Intermediate situations: **partial control** of a decision by DM

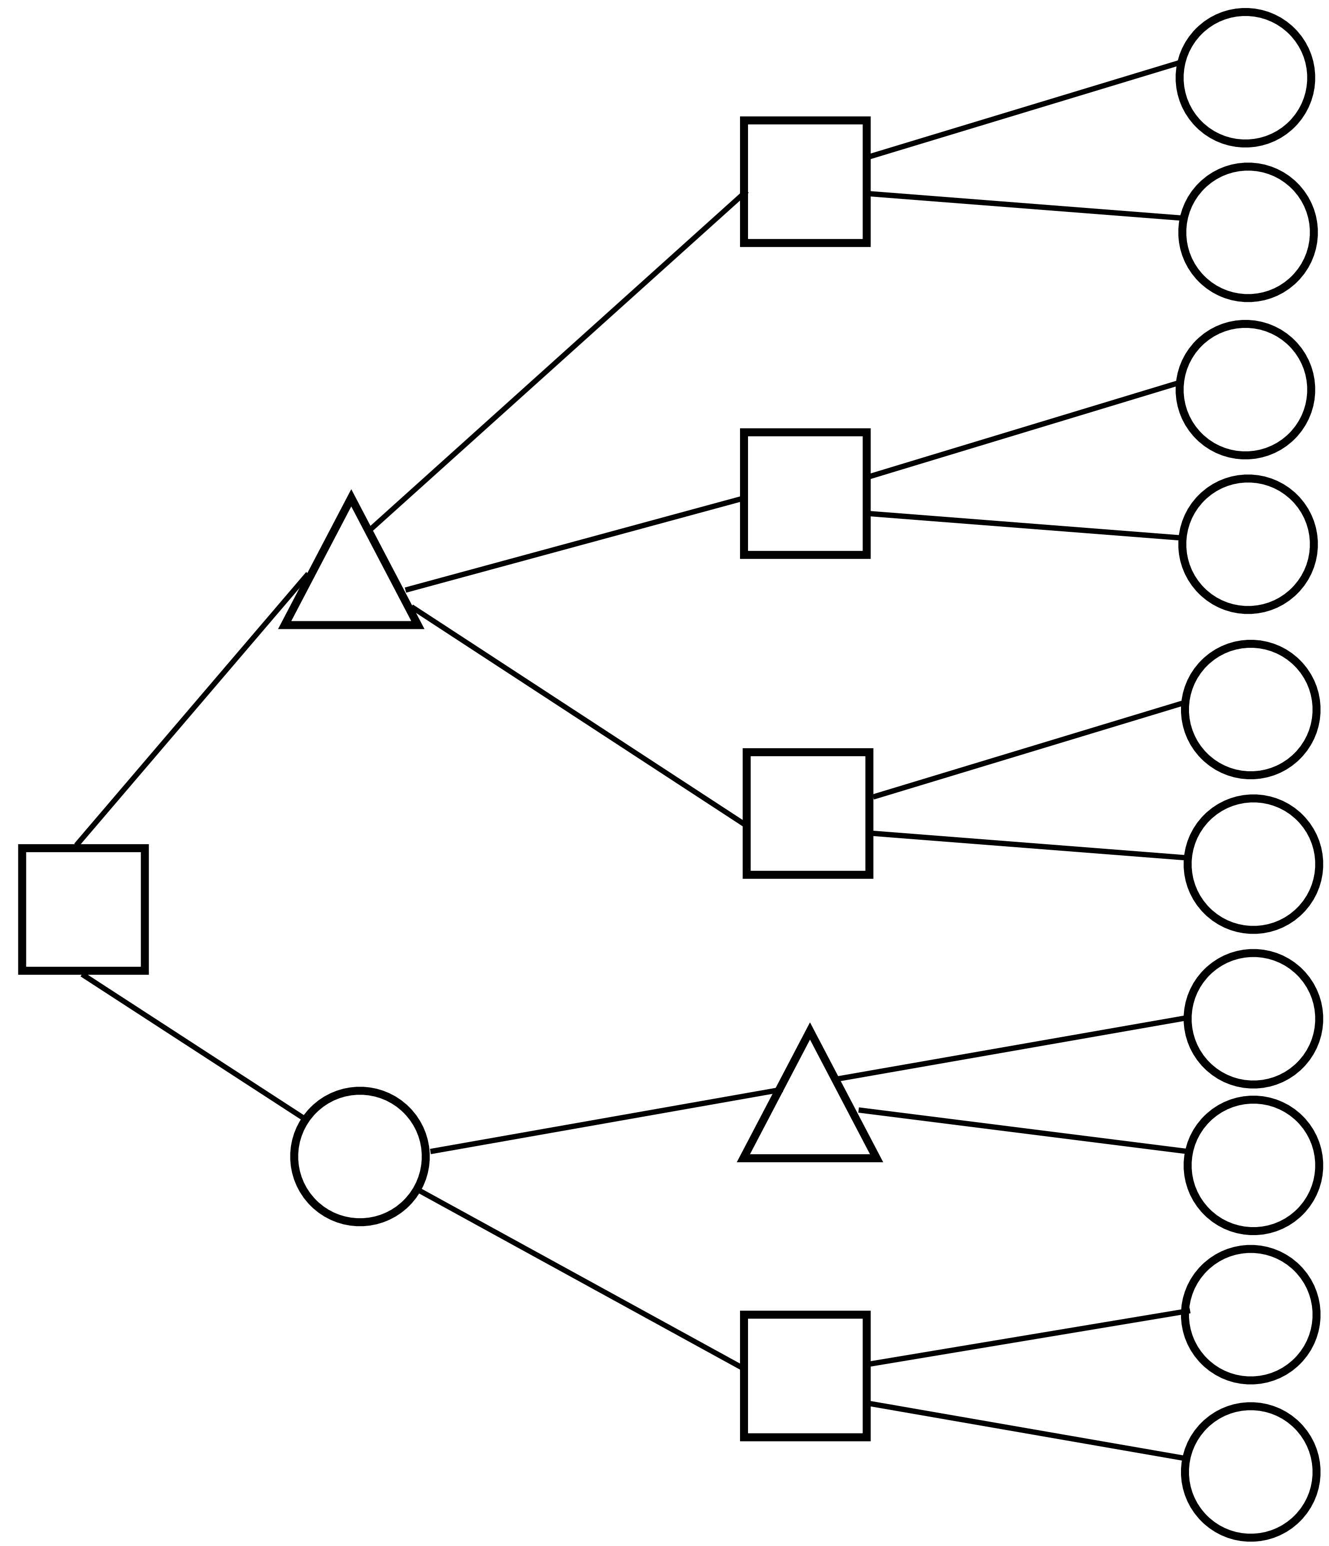

Decision situations

Chance situations

Intermediate situations

Influence

Intermediate situations decision task:

Combine choices from more than one decision maker at one decision node.

- Average of preferred choices (assumes algebraic structure)
- Voting models (algorithm to select group preference from individual decision makers' preferences)
- Probability distributions to *share ownership* in each knot

Definition 11. Influence distribution and influence tree. *Let $T \in \mathcal{T}$ be a decision tree and $P = (p_x)_{x \in T'}$ be a family of probability distributions on a set of decision makers B. For each $x \in T'$ let β_{p_x} be a random variable with distribution p_x. Then P is called* influence distribution *and β_P defined by $\beta_P(x) = \beta_{p_x}$ $(x \in T')$ is called* influence tree *generated by P.*

Interpretation

$(\beta^{(i)})_{i\in\mathbb{N}}$ be a sequence of independent realisations of the same influence tree β_P

Then, by Borel's law of large numbers, with probability 1,

$$\frac{1}{n}\left|\left\{i \in \mathbb{N} \,\middle|\, 1 \leq i \leq n, \beta^{(i)}(x) = b\right\}\right| \longrightarrow p_x(b) \quad \text{for } n \to \infty$$

for all $b \in B$ and for each $x \in T'$.

Asymptotically:

probability that the decision in x is taken by decision maker b is $p_x(b)$

Interpretation:

b has an influence of $p_x(b)$

Utility and knowledge trees

Definition **Utility tree.** *Let $T \in \mathcal{T}$ be a decision tree equipped with an influence tree β_P generated by P. For each $b \in B$ let $u_b : \mathcal{R} \longrightarrow \mathbb{R}$ be a the utility of decision maker b. Then the* utility tree U *is defined as*

$$U(x) = u_{\beta_{px}}(r(x)) \qquad (x \in T).$$

Definition **Knowledge tree.** *Let $T \in \mathcal{T}$ be a decision tree equipped with an influence tree β_P generated by P. For each $b \in B$ let $\kappa_b : T \longrightarrow \mathcal{S}$ be a function that assigns each $x \in T$ the knowledge available to decision maker b in that decision point. Then the* knowledge tree K *is defined as*

$$K(x) = \kappa_{\beta_{px}}(x) \qquad (x \in T).$$

Utility Tree: Extreme Examples

Example **Final outcomes utility.** *A decision maker b who only takes into account the utility of the final outcome is modelled by a utility function of the form*

$$u_b : T \times \mathbb{R} \longmapsto \mathbb{R} \text{ with } u(x,r) = 0 \text{ for all } x \in T \setminus L(T).$$

The following class of examples captures the opposite situation.

Example **Elephant utility.** *Assume the utility tree U has the property*

$$U(x) = \sum_{\rho \preceq y \prec x} U(y) \text{ for all } y \in T.$$

This describes a utility that is build up by summing up all utility accumulated along the way. Concrete examples for this can easily be constructed iteratively.

Memory and Foresight

Not all available knowledge and utility may be taken into account.

Definition **Memory.** *A function* $\psi^- : T \longrightarrow \mathcal{P}(T)$ *on a decision tree* $T \in \mathcal{T}$ *is called* memory function *if for any* $x \in T$, $\psi^-(x)$ *is connected and* $x \in \psi^-(x)$.

Definition **Foresight** *A function* $\psi^+ : T \longrightarrow \mathcal{T}$ *on a decision tree* $T \in \mathcal{T}$ *is called* foresight function *if for any* $x \in T$, x *is the root of* $\psi^+(x)$ *and* $\psi^+(x) \subseteq T$.

Examples:

- Forgetful (1-step past)
- Amnesia (0-step past)
- Elephant (full past)
- Myopic (1-step future)
- No future (0-step)
- Farsighted (full future)

Awareness

These concepts can be specific to DM (subjective).

Hence need to be applied accordingly using influence distribution.

Definition **Awareness.** *Let $T \in \mathcal{T}$ be a decision tree equipped with an influence tree β_P generated by P. For each $b \in B$ let ψ_b^- the memory function of band ψ_b^+ the* memory and foresight trees *are defined as*

$$\Psi^-(x) = \psi^-_{\beta_{px}}(x) \text{ and } \Psi^+(x) = \psi^-_{\beta_{px}}(x) \text{ for } (x \in T).$$

The combination $\Psi = (\psi^-, \psi^+)$ is called awareness range.

Agriculture

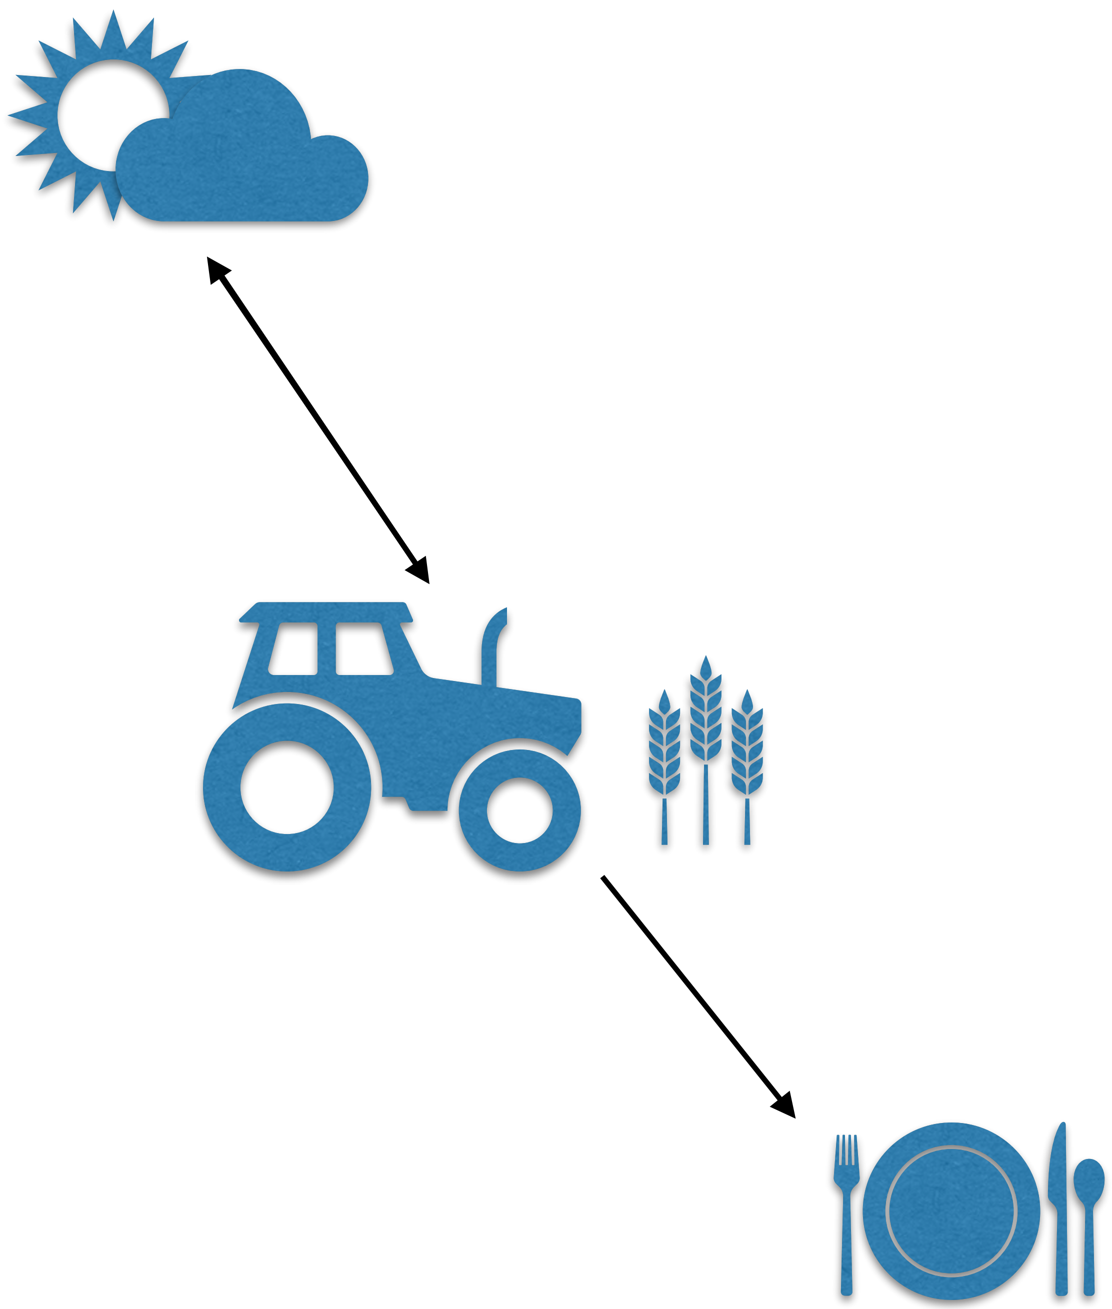

Econometrics perspective

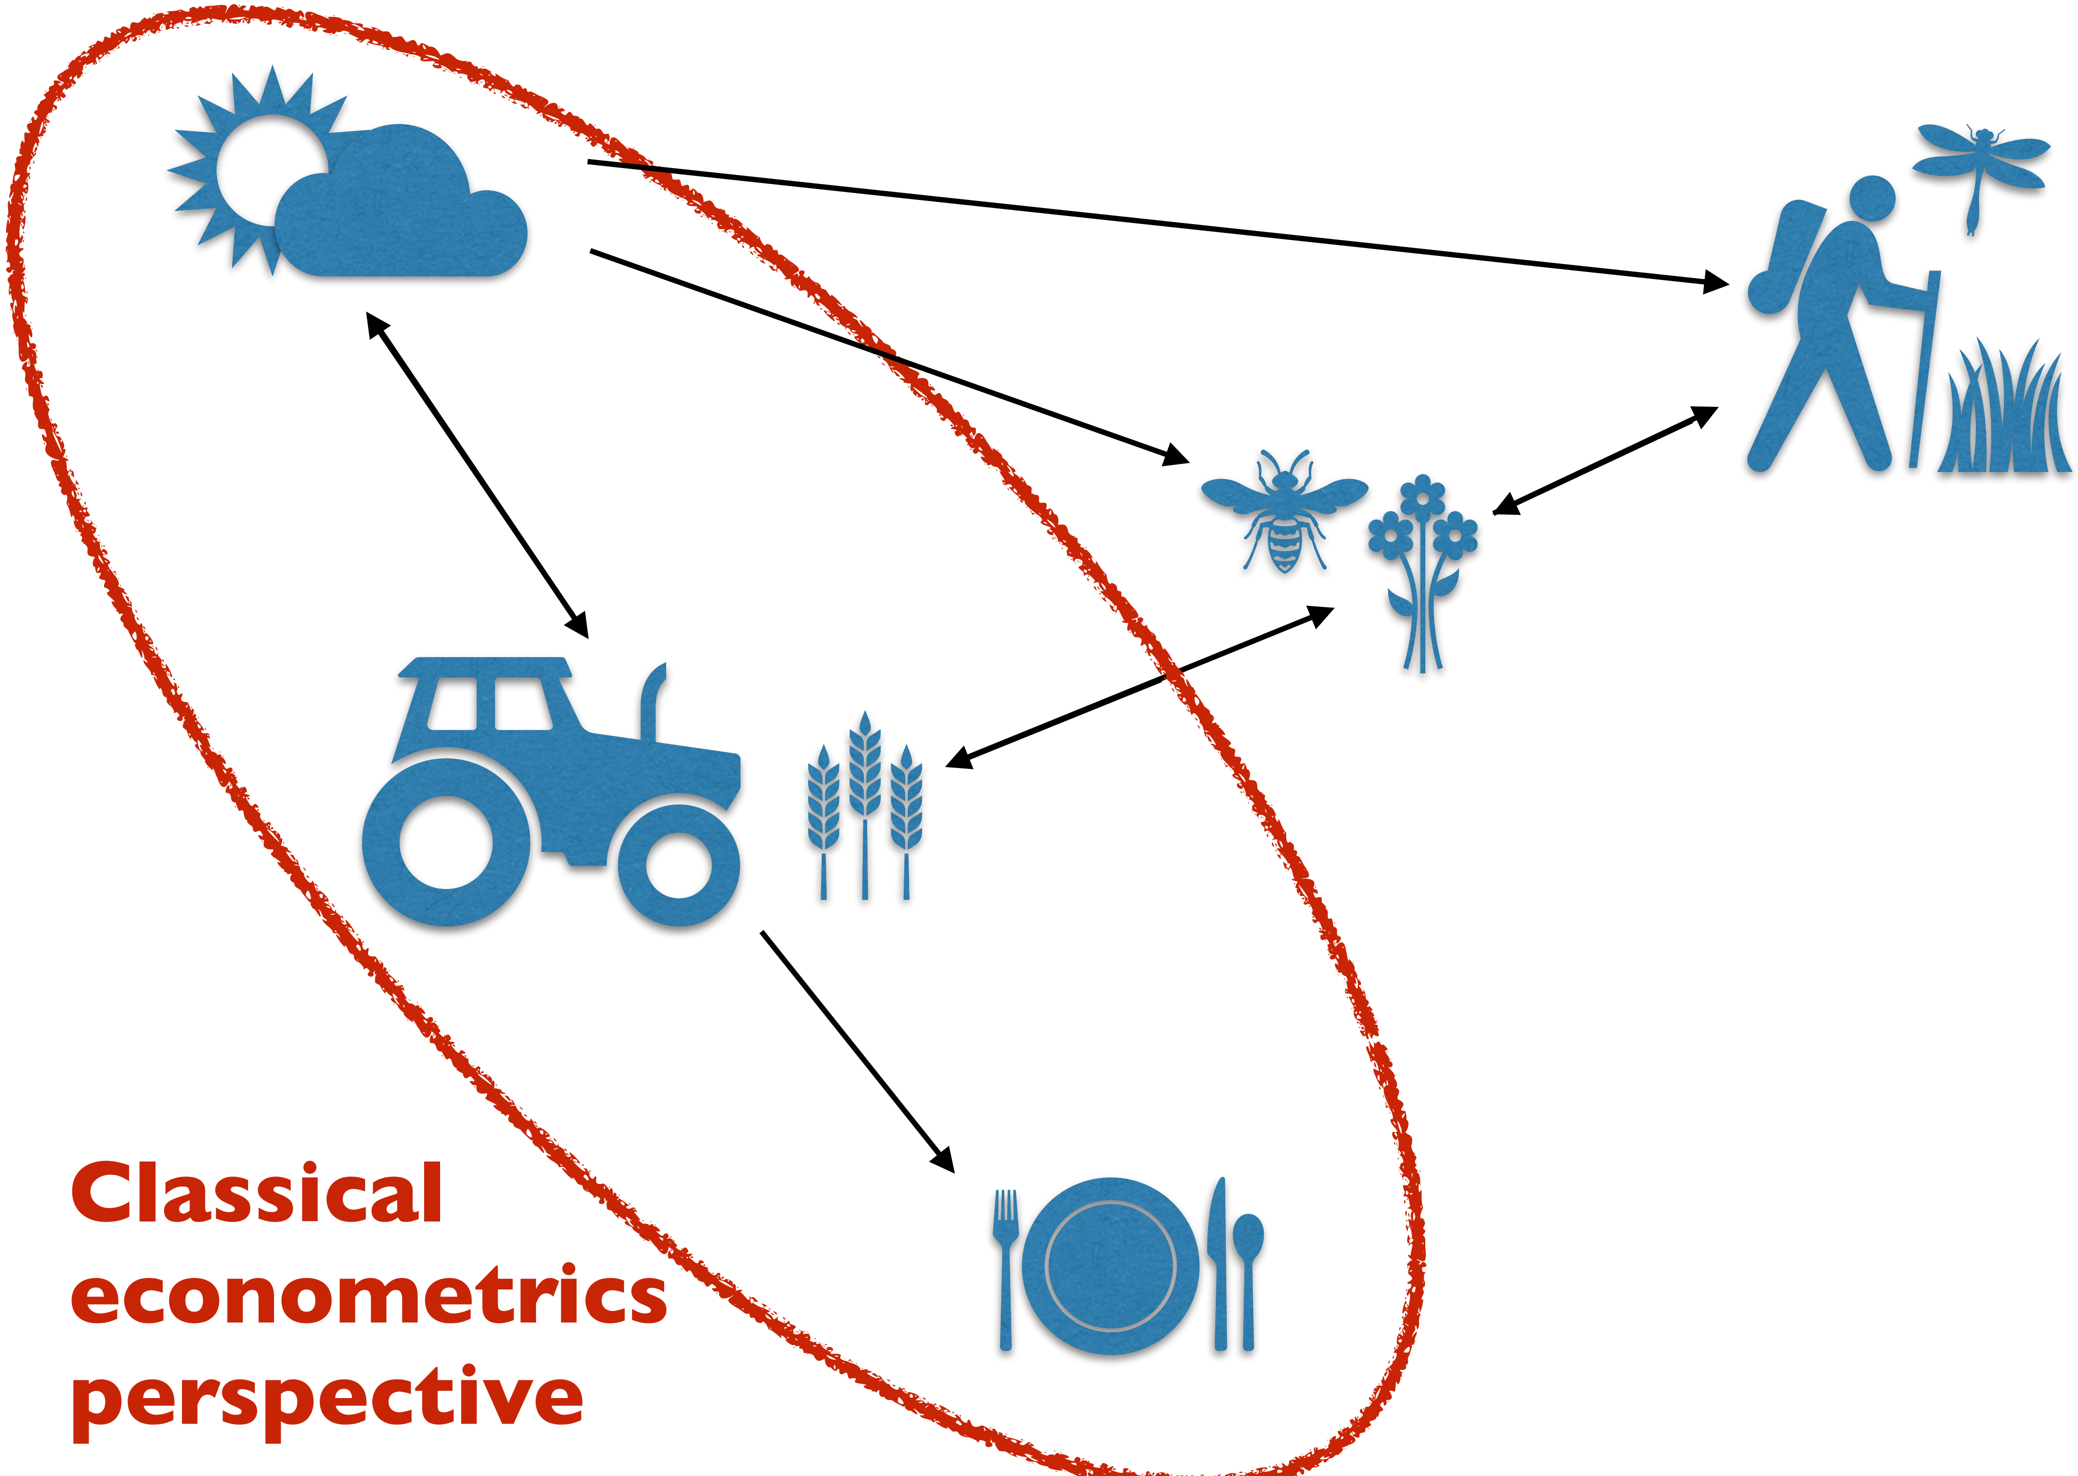

Ecometrics

Metrika, December 1969, Volume 14, Issue 1, pp 293–301

Ecometrics: An Ideal for Economics and Ecology

By J. F. Bennett, Vienna [1])

The concept is based on three main sources of inspiration: physical science, particularly as it illuminates the inanimate world with the comprehensive idea of energy; modern biology; and Professor Sagoroff's energy-balance economics [2]). The term "ecometrics" [3]) is introduced provisionally. In full generality, the ideal of ecometrics can be so expressed: to know all the world's deoxyribonucleic acid (DNA), as to its whereabouts and the energy-transformations which it is catalyzing, at all times. This ideal is unattainable, I believe in principle as well as in practice, but not unapproachable. The main concern of the first section following is with the approach.

---

[1]) Prof. John F. Bennett, University of Pennsylvania, Dept. of History and Philosophy of Science, 103 General Laboratories, Philadelphia, Pa. 19104, USA. This paper contributed while serving as Fullright Lecturer, Institute for Statistics, University of Vienna, 1967—68.

Agri-environmental-societal perspective

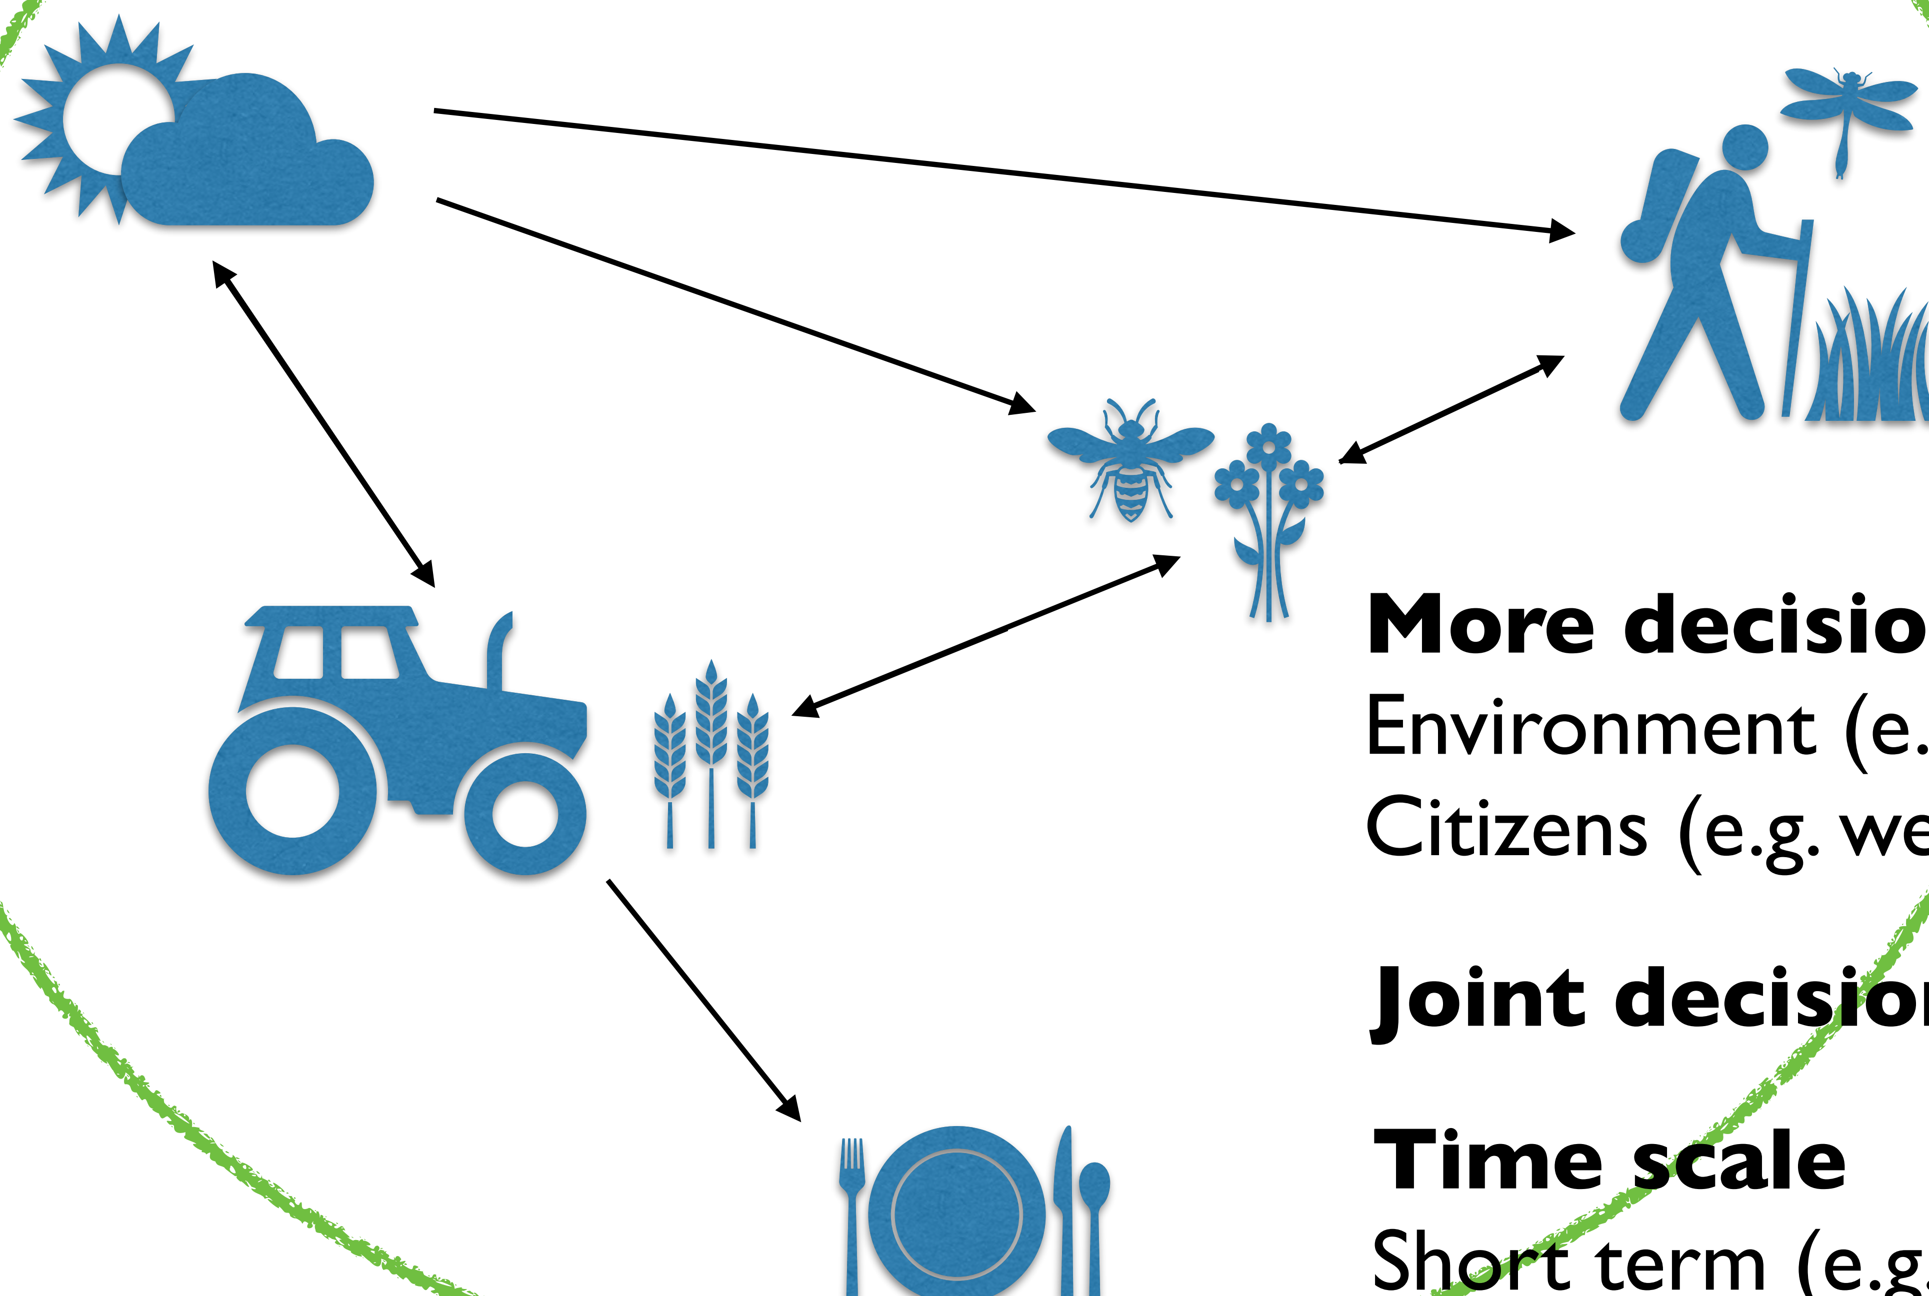

More decision makers

Environment (e.g. pollinators)

Citizens (e.g. wellbeing, tourism)

Joint decisions (influence)

Time scale

Short term (e.g. harvest)

Long term (e.g. soil, air, climate)

Wild pollinators in apple orchards

Study in New York state

M. G. Park *et al.*, *Proceedings of the Royal Society B: Biological Sciences* **282**, 20150299 (2015)

16 orchards over 2 years, data before/after bloom on bloom index, pesticides, pollinators etc

Minimise impact of insecticides and herbicides (indirect) on pollinators!

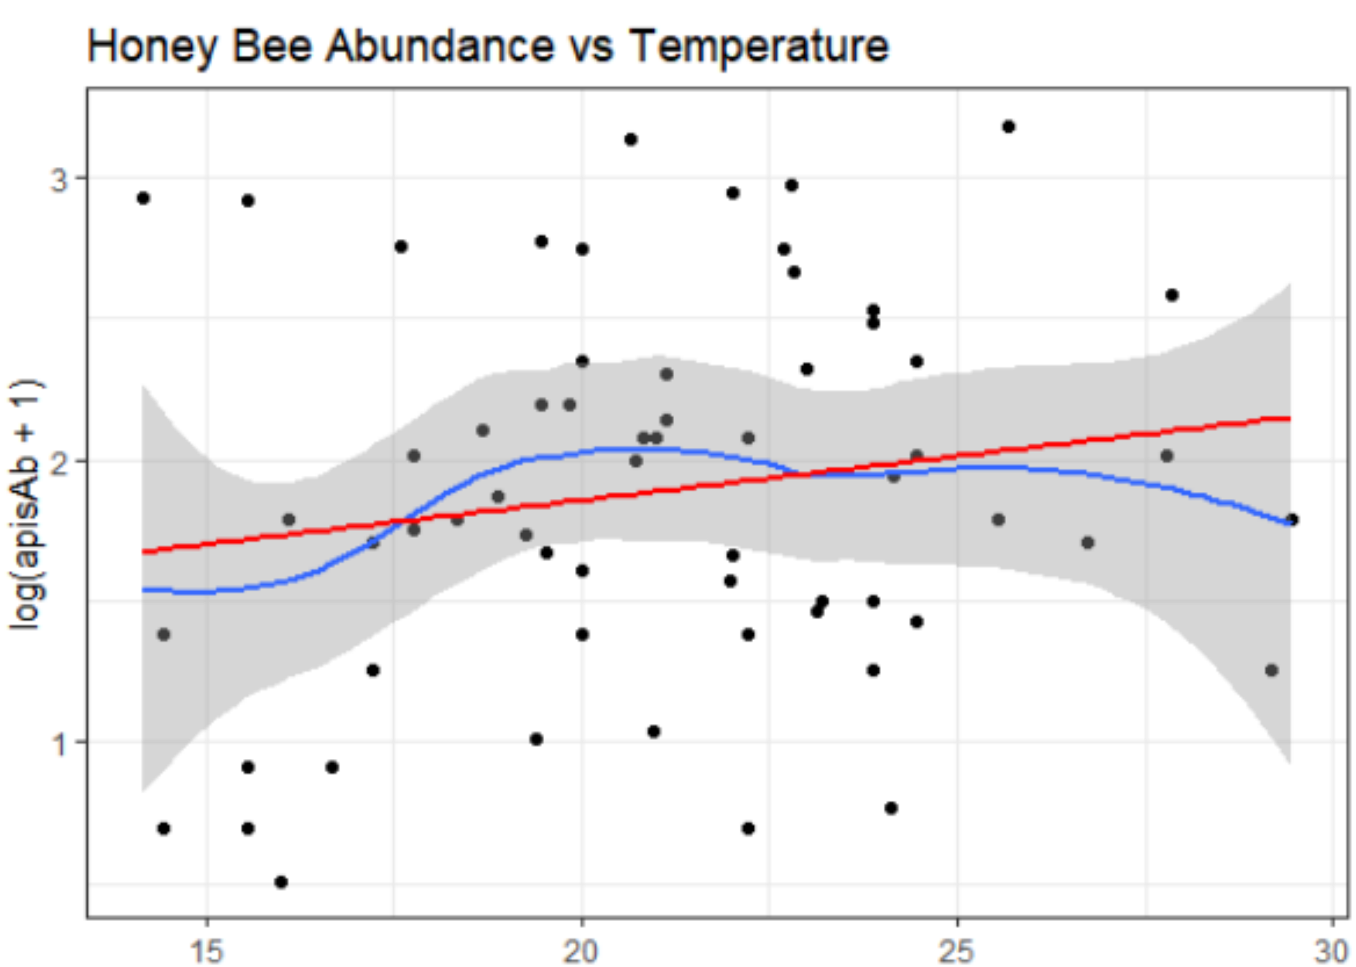

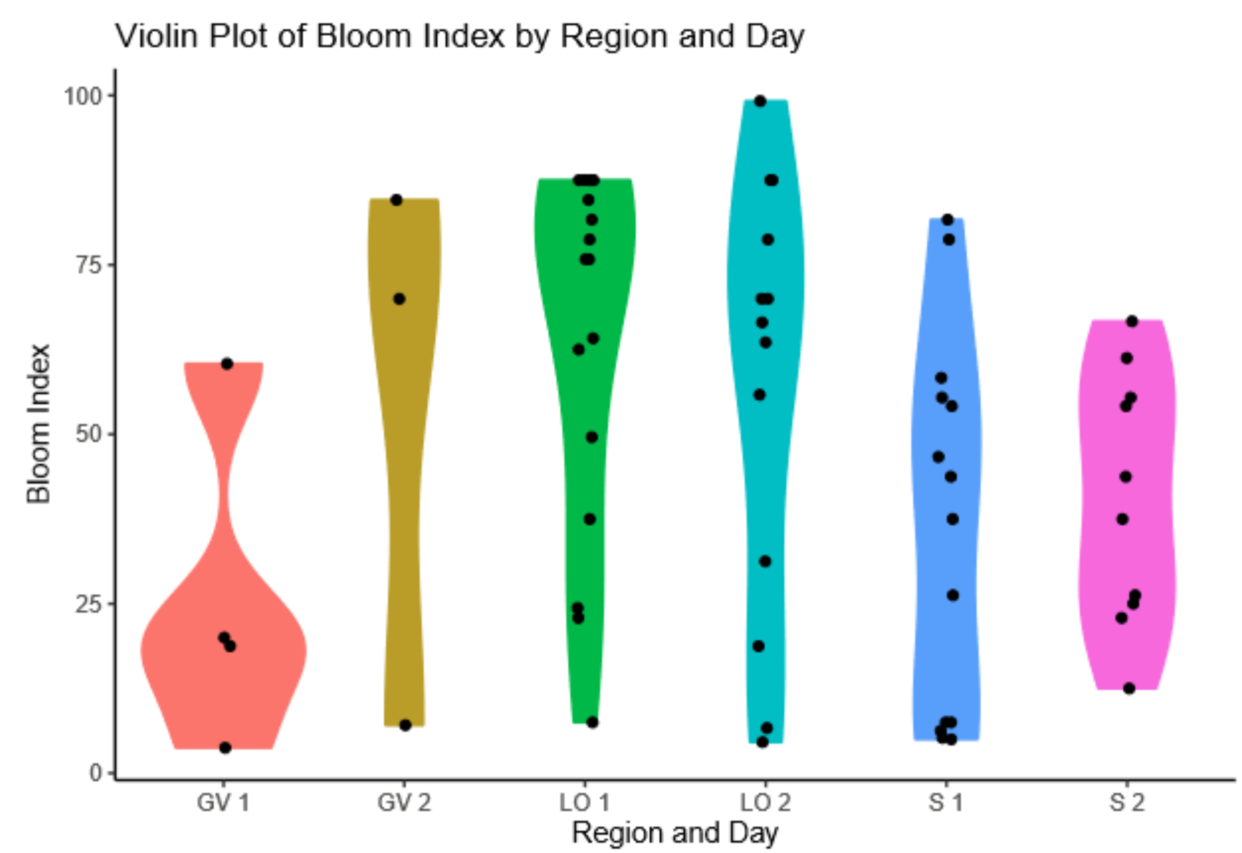

Wild pollinators in apple orchards

Modelling challenges

- Orchard management is individualised
- Covariate dependency (e.g. temperature)
- Missing data (mostly 1st year)
- No dates, only phases (“before/during/after bloom”) and bloom index varies largely within assessment days

Case study: farm scale experiment

Study in the UK

J. N. Perry *et al.*, *Journal of Applied Ecology* **40**, 17–31 (2003).

- Maize, Beet, Spring Oilseed Rape, and Winter Oilseed Rape
- Records of the impact of growing practices on biodiversity and crop yield: herbicide application timings, percentage cover of weeds, crop height, biodiversity counts, pollinator counts during the growing season, Met Office weather station data, yields
- 65 fields per crop on average
- Application of herbicides on weeds impacts pollinators

Case study: Farm scale experiment

Modelling challenges

- Individual management schemes
- Data in form of complex time courses
- Chemical quantities of pesticides given rather than environmental impact measures
- Decision rules

Current & future work

- Building trees: normative using expert judgement
- Building trees: data driven using machine learning; goes back to random forests (Leo Breiman) actually!
- Time and state dependent covariates
- Decision rules: What is optimal and for whom?
- Data quality benchmarking and correction: missing data & imputation, sampling biases & adjustments
- Local vs global view
- Deviations from rationality

Resources

- ST222 https://warwick.ac.uk/fac/sci/statistics/currentstudents/modules/st2/st222 and resource page and resource page for this module
- ST301 https://warwick.ac.uk/fac/sci/statistics/currentstudents/modules/st3/st301and resource page for this module
- Parmigiani, Lourdes, "Decision theory, Principles and Approaches", Wiley & Sons, 2009.
- Koerner, "Naive Decision Making: Mathematics Applied to the Social World" (Cambridge University Press)
- Petersen, "An Introduction to Decision Theory" (Cambridge Introductions to Philosophy)
- Smith, J. Q. (2010). Bayesian Decision Analysis: Principles and Practice. Cambridge University Press.
- French, S., & Smith, J. Q. (Eds.). (1997). The Practice of Bayesian Analysis. Hodder Education.
- Keeney, R. L., & Raiffa, H. (1993). Decisions with Multiple Objectives: Preferences and Value Trade-offs. Cambridge University Press.
- DeGroot, M. H. (2005). Optimal Statistical Decisions (Vol. 82). John Wiley & Sons.

Thank you:

Bärbel Finkenstadt-Rand, Yiyuan Zhang (Warwick Statistics)

Francis Lévi (Warwick Medical School)

Pascal Guénel, Emilie Cordina Duverger (INSERM, CESP)

Terry Speed (WEHI), François Collin (Genome Health), Ben Bolstad (Affymetrics), Guilia Kennedy (Veracyte)

Thank you:

NERC Landscape Decisions Maths Small Grants

“JDec: Joint decision models for citizens, crops, and environment” (1.10.2019-30.9.2020)

Co-I: Rosemary Collier (Warwick School of Life Sciences)

Project partners: Maria Christodoulou, David Steinsaltz (Statistics & Biodemography Group, Oxford Statistics)

3rd year Data Science students: Stephen Brownsey, Elizabeth Potter, Matt Persin (Warwick Statistics)

M2D Feasibility Found award: PI Maria Christodoulou “Deciding to grow: Agriculture and forestry in a changing environment”