Dynamic Decision Problems: Hybrid Use of Decision Trees & System Dynamics Models

Nathaniel Osgood
(with some slides courtesy of Karen Yee)

CMPT 858

March 24, 2011

Local Context

- Saskatchewan suffered the highest incidence of WNV in Canada in 2003 and 2007
- Saskatoon Health Region (SHR) reported 6.5% and 25% of the provincial cases in 2003 and 2007, respectively

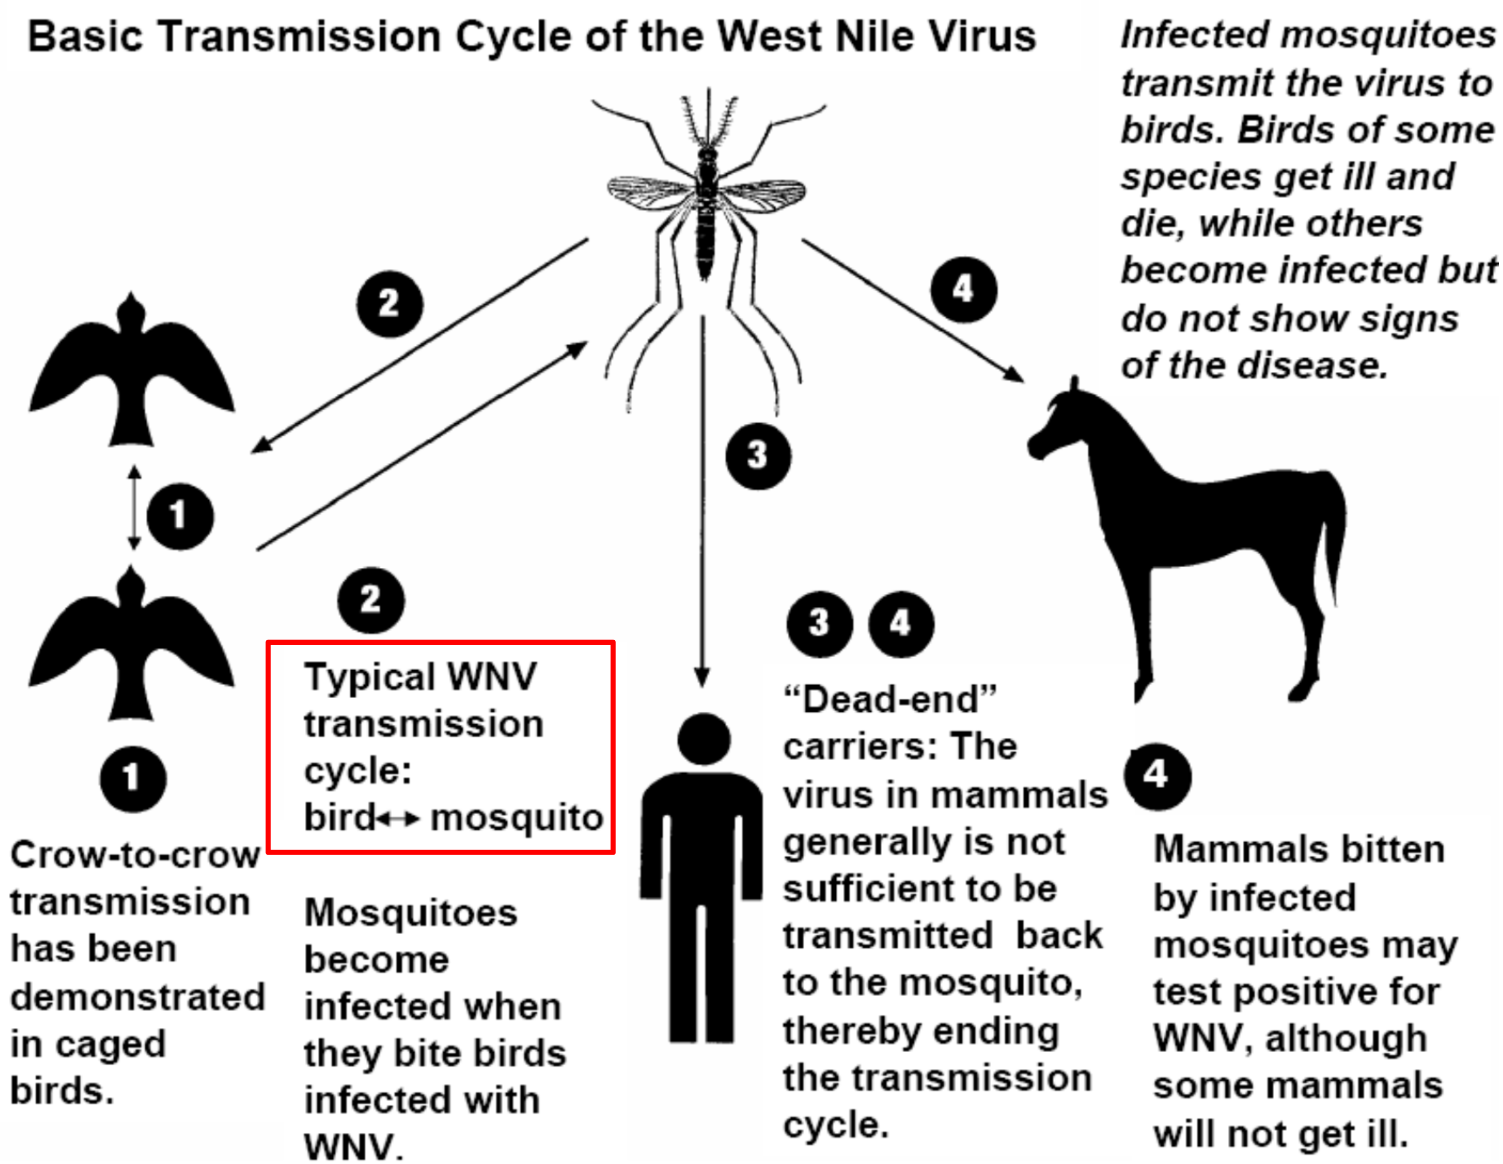

Mosquito species *Culex tarsalis* primarily responsible for spreading WNV in Saskatchewan

Source: Penn State University

Adult

Source: www.azstarnet.com/metro/295104

Terrestrial

Aquatic

Mosquito Lifecycle

Egg Raft

Source: http://www.flickr.com/photos/lordv/207198441/

Larvae

Source: unknown

Pupa

Source: www.comosquitocontrol.com/Mosquito_Biology.html

Mosquito Environmental Dependencies

For increasing their numbers:

- Temperature (Average number of night temperatures above 15ºC; heat accumulated days)
- Habitat availability
- Rainfall

Human Dependencies

For using protective measures:

- Perceived risk
- Knowledge of WNV
- Temperature

Health Managers' Dilemma

Need to make decisions now taking into account uncertainties regarding:

1. Mosquito population
   - abundance
   - WNV prevalence
2. Environmental conditions
   - current
   - forecasted
3. Human behaviour

Different Levels of Challenge in Dynamic Decision-Making

- **Type A:** Making complex dynamic choices given some expected/typical course of important factors outside our control
  – Here, the focus is centred on building models that help us understand the complex impact of our choices given this 'expected course'
  – **Tough**
- **Type B:** Making complex dynamic choices when we can't anticipate the course of the important factors outside our control
  – Focus on both dynamic model and adaptive planning given uncertainty
  – **Tougher**

Implications

- **Type A:** When important exogenous conditions are known, we often seek to identify & stick to an 'optimal' pre-set plan
  – Don't have to worry much about unfolding external conditions – they are known or unimportant
- **Type B:** Rather than putting "all our eggs in one basket", it is typically best to avoid a pre-set plan, and instead to *adaptively* make choices over time
  – What we will do over time will depend on what is observed

Decision making under dynamic uncertainty: Adaptive Planning

- **Type A:** When important exogenous conditions are known, we often seek to identify & stick to an 'optimal' pre-set plan
  - Don't have to worry much about unfolding external conditions – they are known or unimportant

The presentation focuses on this type of "dynamic decision" problems

- **Type B:** Rather than putting "all our eggs in one basket", we typically seek to avoid a pre-set plan, and instead to *adaptively* make choices over time
  - What we will do over time will depend on what is observed

Adaptive Decision Problems: Relevant Questions

How do we make decisions now, *when the choice of the best decision depends so much on what plays out (unfolds) in things beyond our control?*

- Temperature trends
- Precipitation
- Prevalence of infection in migratory bird populations

Should we make our decisions now despite these uncertainties? Or should we wait to see how things are trending before making decisions?

Characteristics of "Adaptive Decision" Problems

- Can't count on one particular future trajectory unfolding for things outside our control
- Choosing decisions *now* requires considering the different possibilities of what might unfold in the future
- We must make decisions over time, as we observe things unfold
  - The later we wait, the more information we'll have
- It may be advantageous to decide to "wait and see" as to how things play out until a later decision point

In these Conditions…

- What decision we make at a particular point in time will depend on
  - Our current situation
    - What we've observed as happening to this point (what we've "learned" – e.g. recent levels of growth)
    - State (as given by stocks & derived quantities)
  - Possible future eventualities, in light of what we've already seen (e.g. future levels of growth, given recent growth)
  - Our possible decision points in the future
- Here, we are balancing two desires:
  - To "seize the moment" and act early
  - To "wait and see" what happens, and decide on the basis of this

A Hybrid System Architecture to Address these "Tougher" Problems

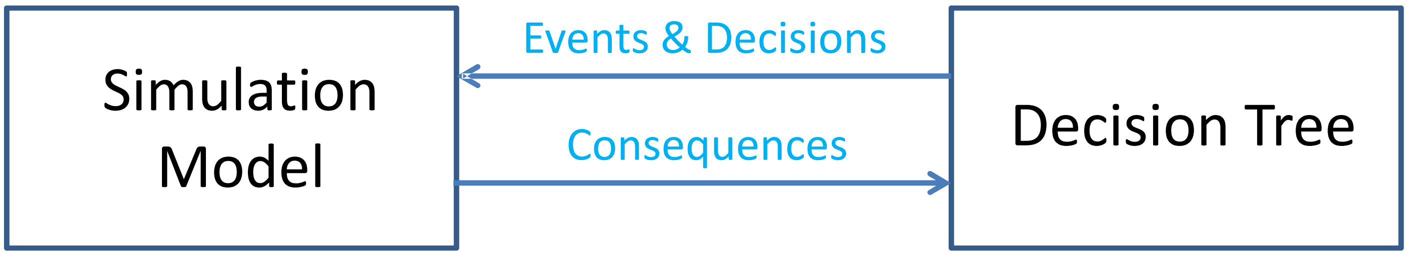

Introduction to Decision Trees

- We will use decision trees both for
  - Diagrammatically illustrating decision making w/uncertainty
  - Quantitative reasoning
- Represent
  - Flow of time
  - Decisions
  - Uncertainties (via events)
  - Consequences (deterministic or stochastic)

Decision Tree Nodes

- Decision (choice) Node

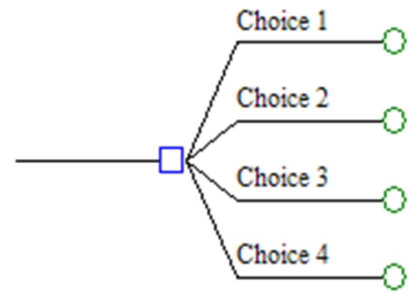

- Chance (event) Node

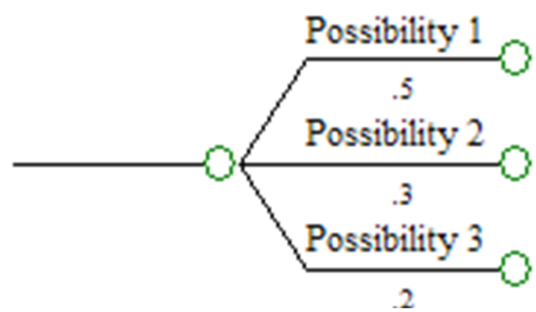

- Terminal (consequence) node

Identifying the Optimal Decision Rule

- To select decision rules, we perform a "rollback" of the tree (dynamic programming)
  - For *terminal* nodes, pass up *value*
  - For *event* nodes, pass up *expected value* of children
  - For *decision* nodes, *select whichever child offers highest value* and pass up that value for this node

Example Tree
Feature Decision Making

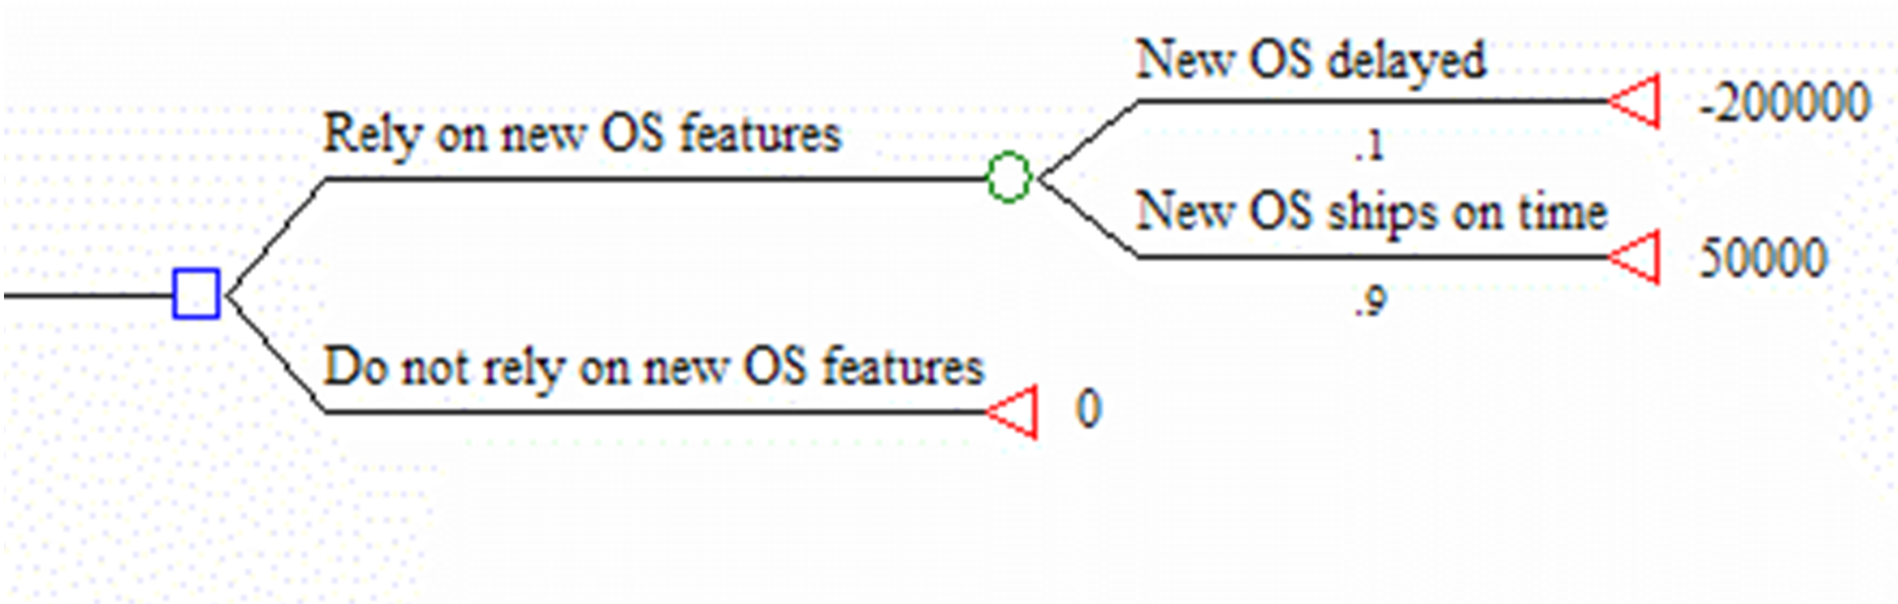

Best Option

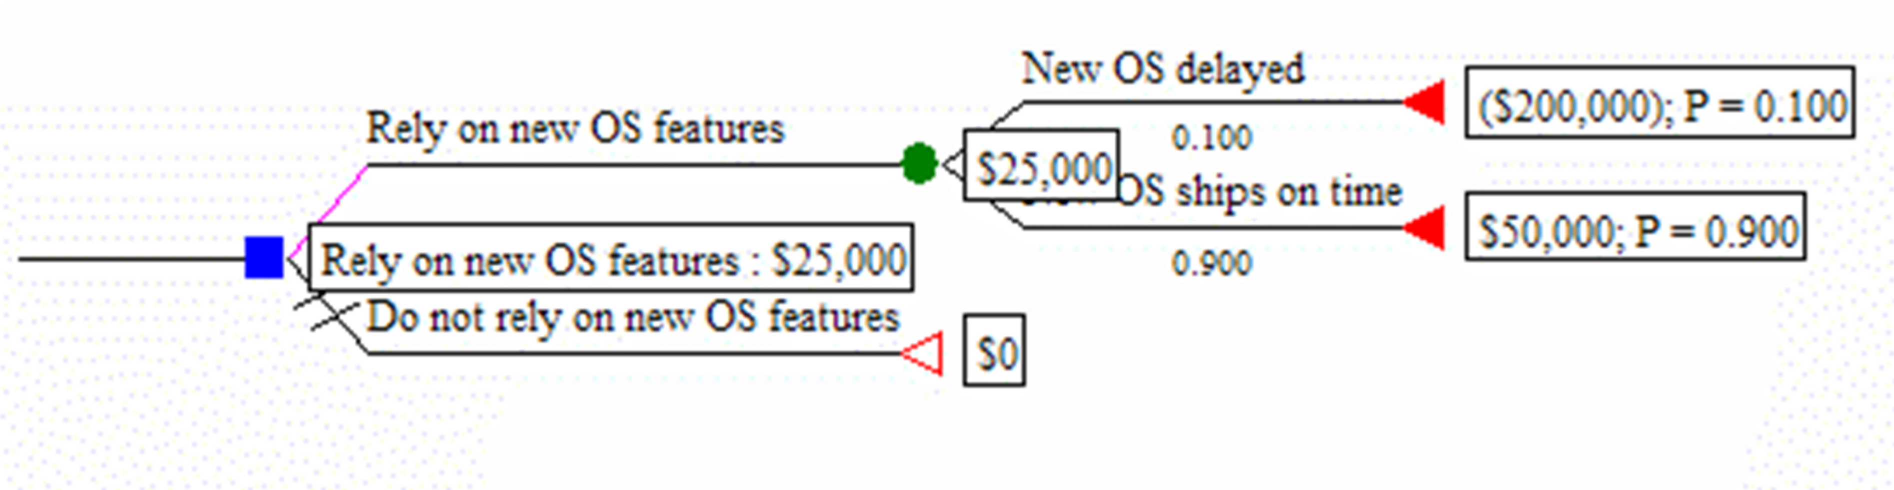
New OS delayed
($200,000); P = 0.100
0.100
Rely on new OS features
$25,000
OS ships on time
$50,000; P = 0.900
Rely on new OS features : $25,000
0.900
Do not rely on new OS features
$0

Extended Example

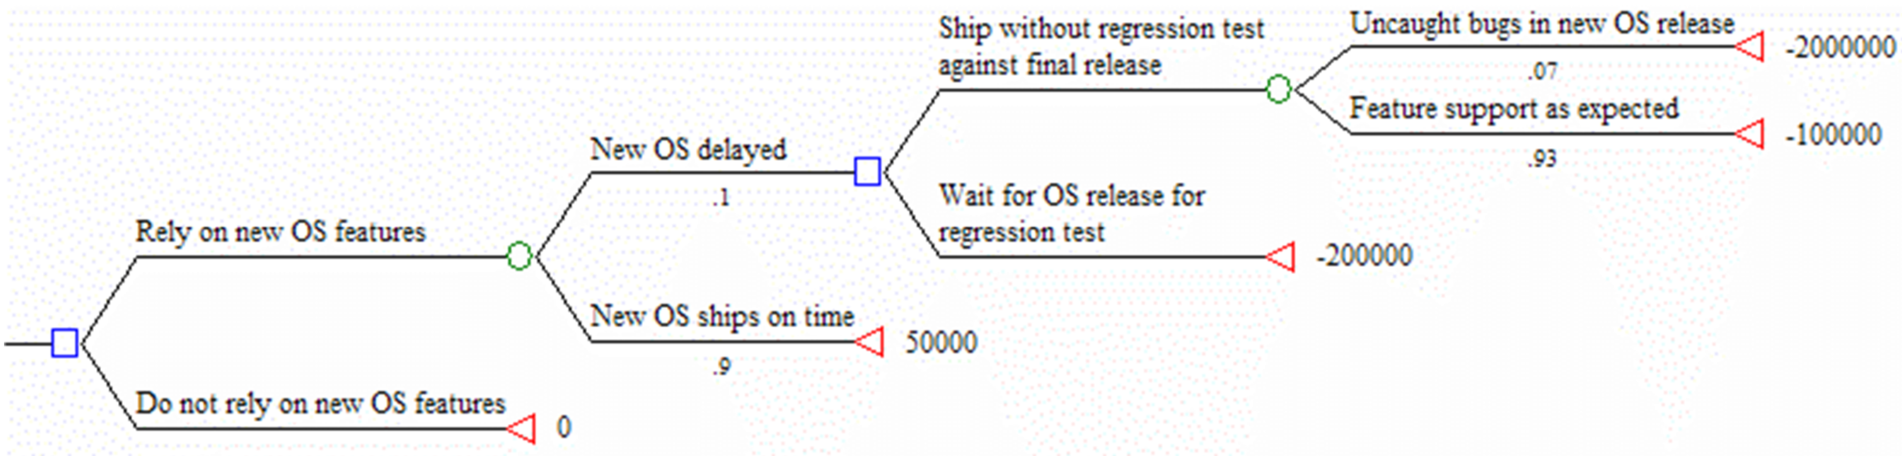

Extended Rule

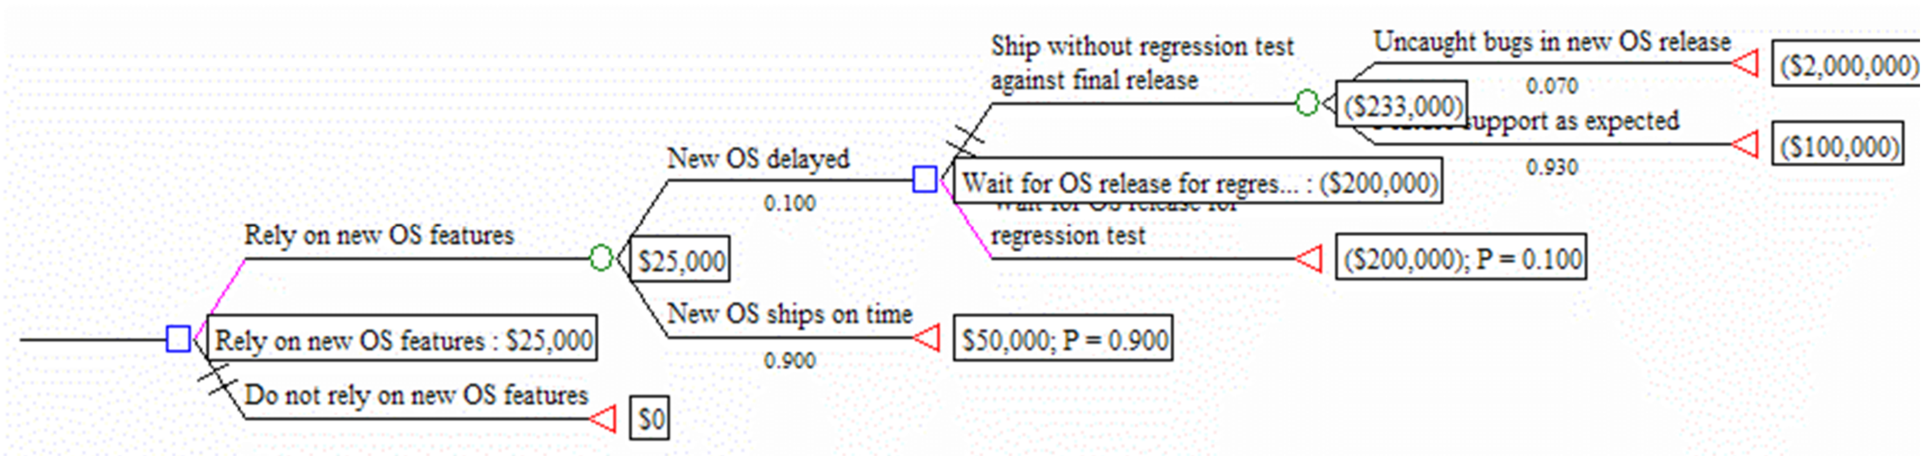

Decision Rules

- Decision trees can be used to identify "optimal" decision rules
  - Remember: Optimality is in light of (simplified) assumptions!
- A decision rule specify what we should do given any possible eventuality

Decision Tree To Structure Policy Space

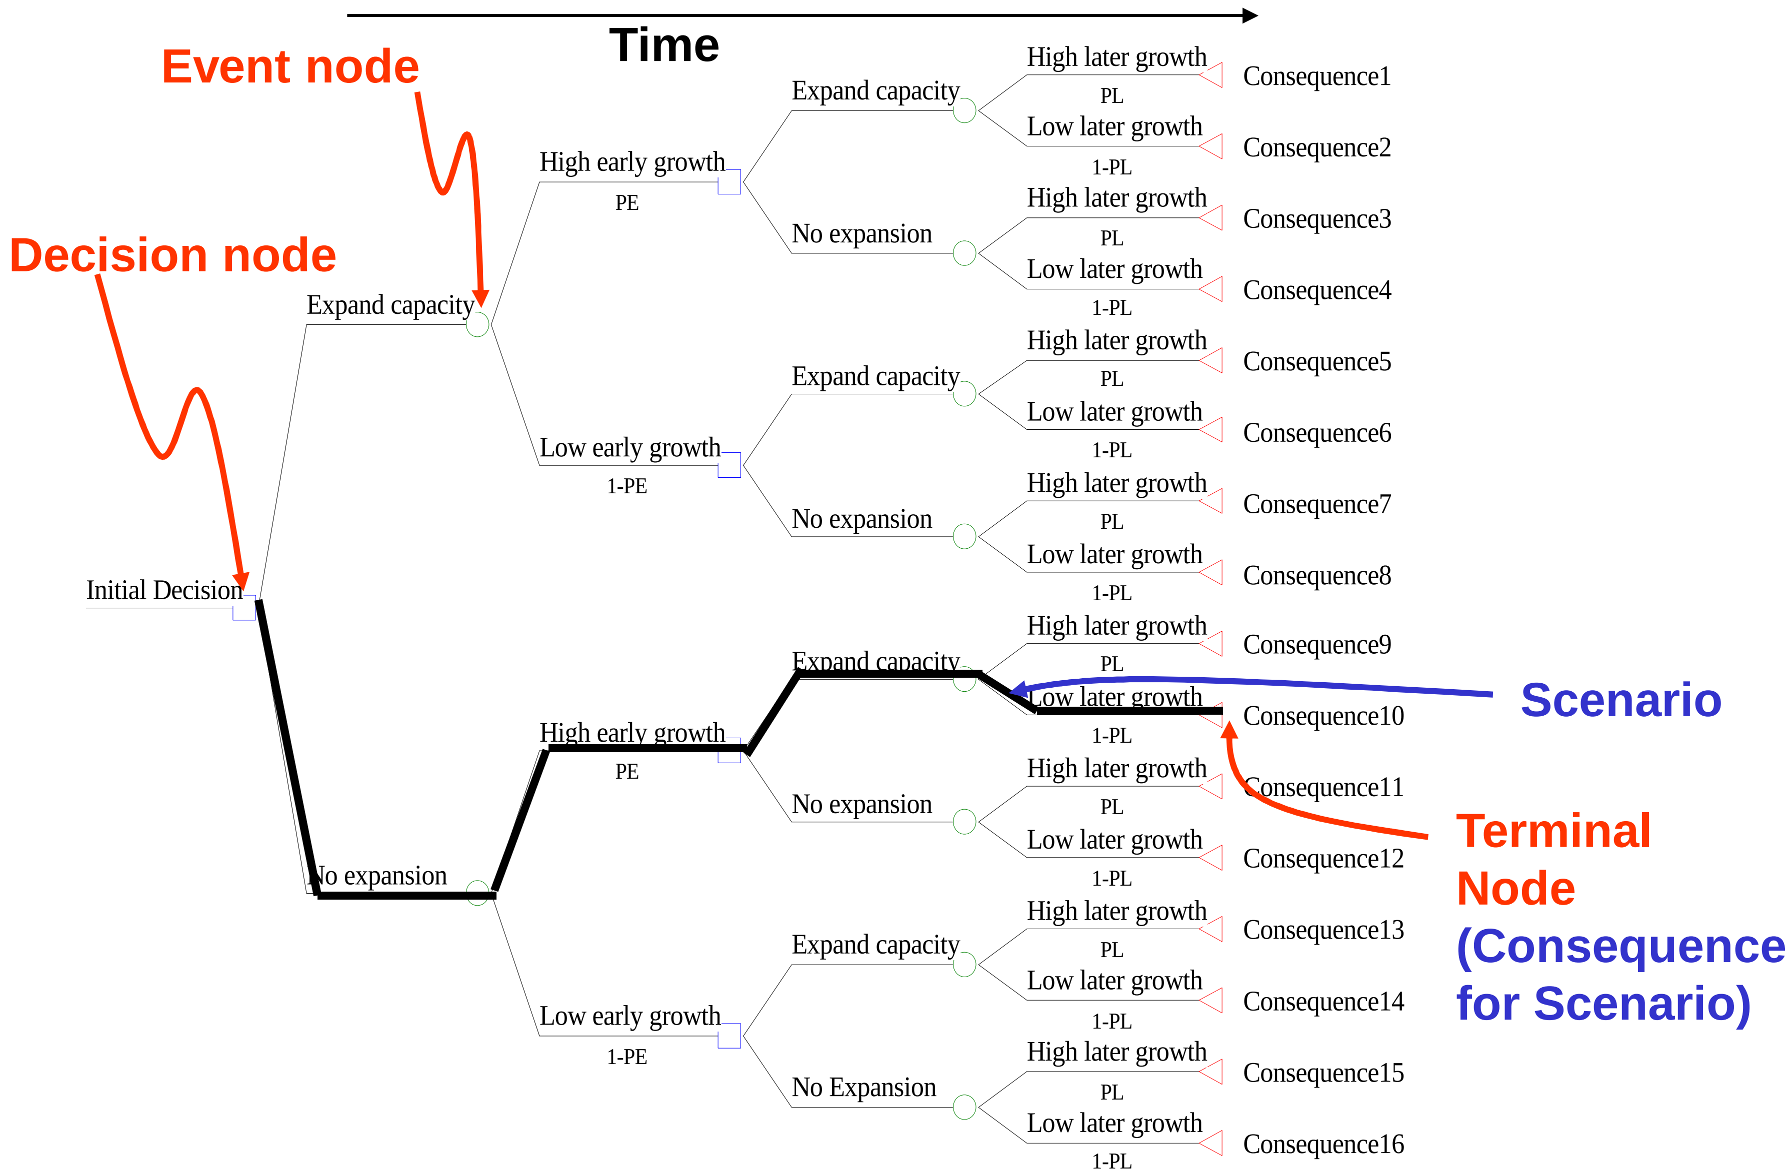

Terminology

- A *static* decision rule pursues the same predetermined decisions (actions) regardless of eventualities
- An *adaptive* decision rule varies its decisions (actions) based on which events have occurred
- *Observation: Static decision rules are rarely optimal*

A Static Decision Rule

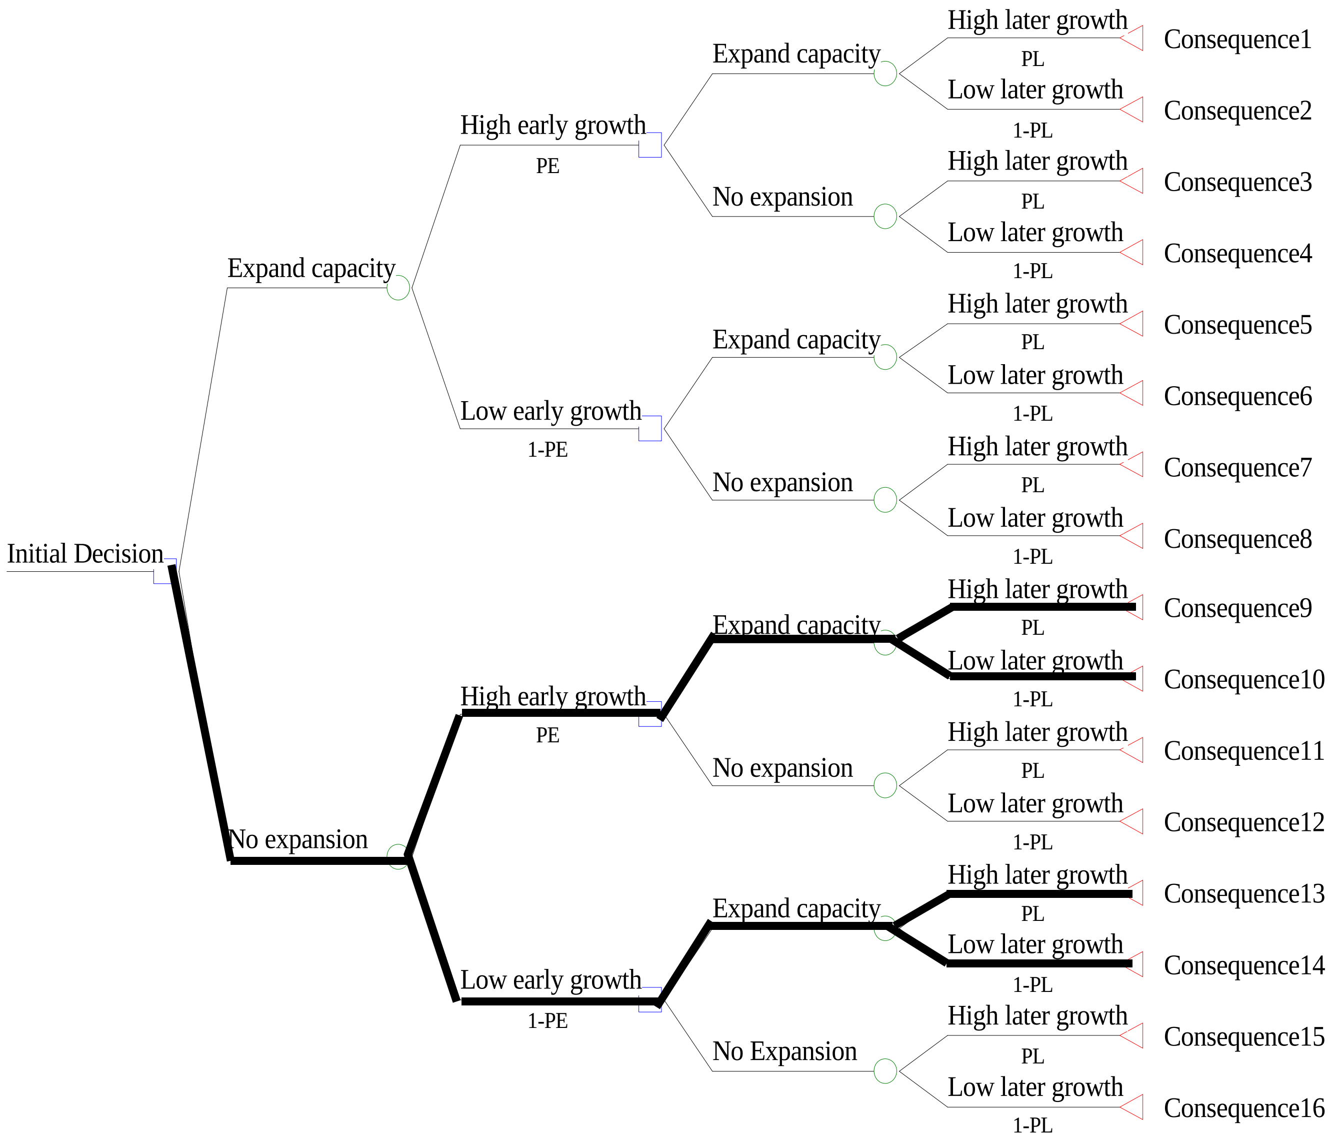

Observation:
The consequences observed at a particular terminal node are a function of the associated scenario (***particular sequence of decisions and events on the path leading to that terminal node)*** – and are the same regardless as to which decision rule that gives rise to this sequence

An Adaptive Decision Rule

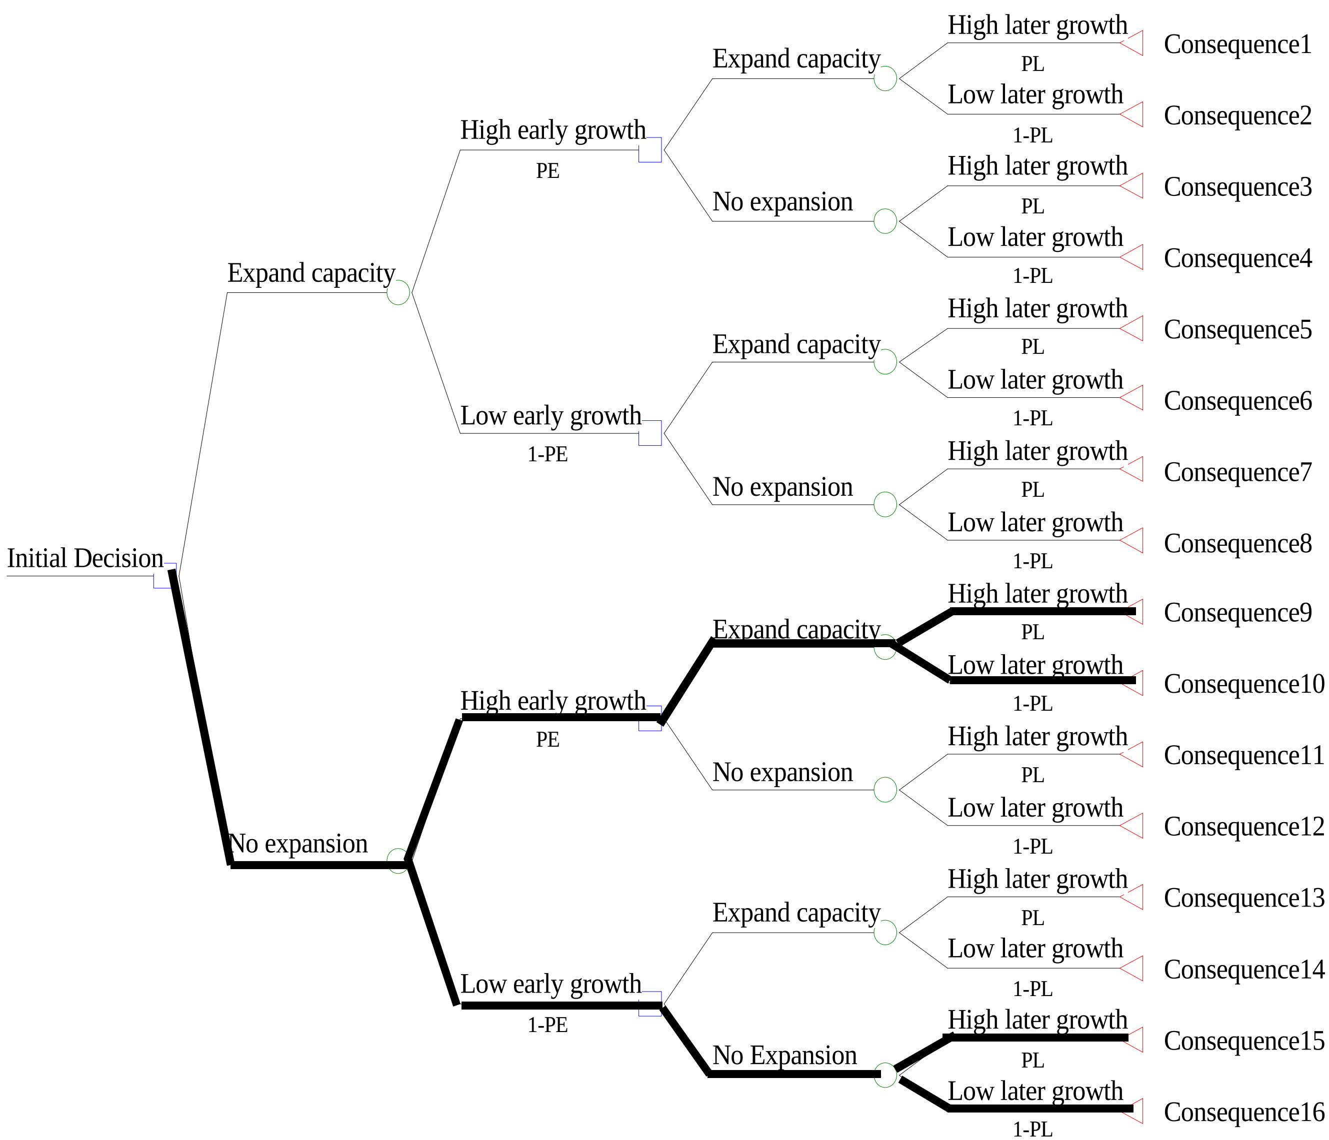

Observation:
The consequences observed at a particular terminal node are a function of the associated scenario (***particular sequence of decisions and events on the path leading to that terminal node)*** – and are the same regardless as to which decision rule that gives rise to this sequence

Analysis Using Decision Trees

- Decision trees are a powerful analysis tool
- Addition of symbolic components to decision trees greatly expand power
- Example analytic techniques
  - Strategy selection
  - One-way and multi-way sensitivity analyses
  - Value of information

Decision Tree w/Variables

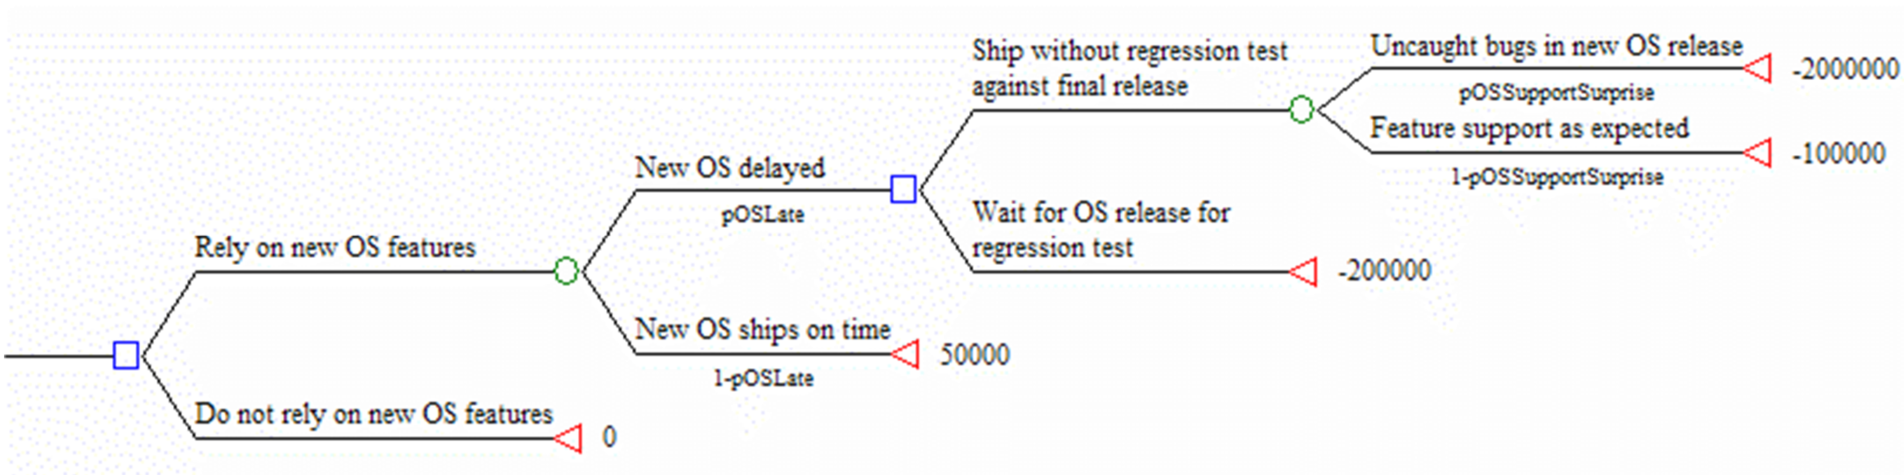

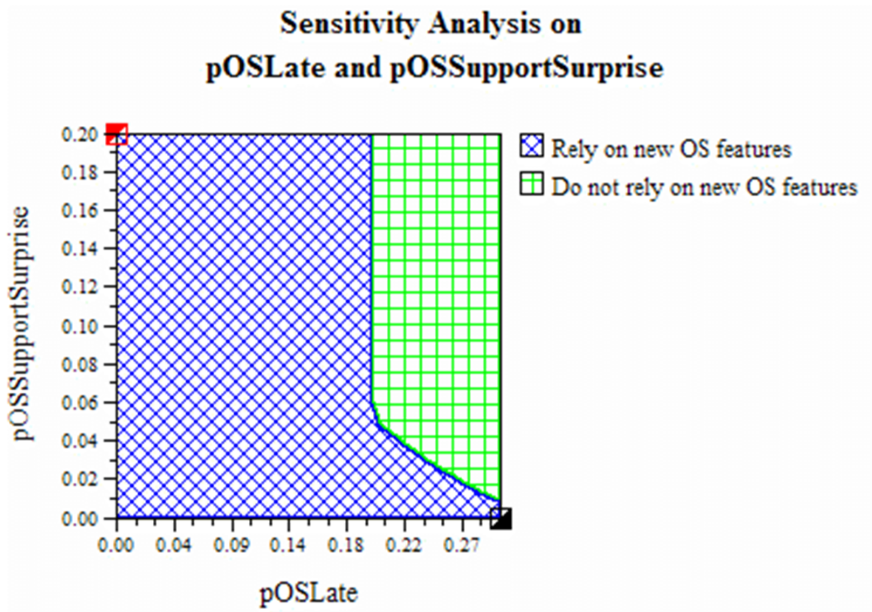
Sensitivity Analysis on
pOSLate and pOSSupportSurprise
Rely on new OS features
Do not rely on new OS features
pOSSupportSurprise
0.20
0.18
0.16
0.14
0.12
0.10
0.08
0.06
0.04
0.02
0.00
0.00 0.04 0.09 0.14 0.18 0.22 0.27
pOSLate

Risk Preference

- People are not indifferent to uncertainty
  - Lack of indifference from uncertainty arises from uneven preferences for different outcomes
  - E.g. someone may
    - dislike losing $x far more than gaining $x
    - value gaining $x far more than they disvalue losing $x.
- Individuals differ in comfort with uncertainty based on circumstances and preferences
- Risk averse individuals will pay “risk premiums” to avoid uncertainty

Risk Preference
(Decision Tree Preview)

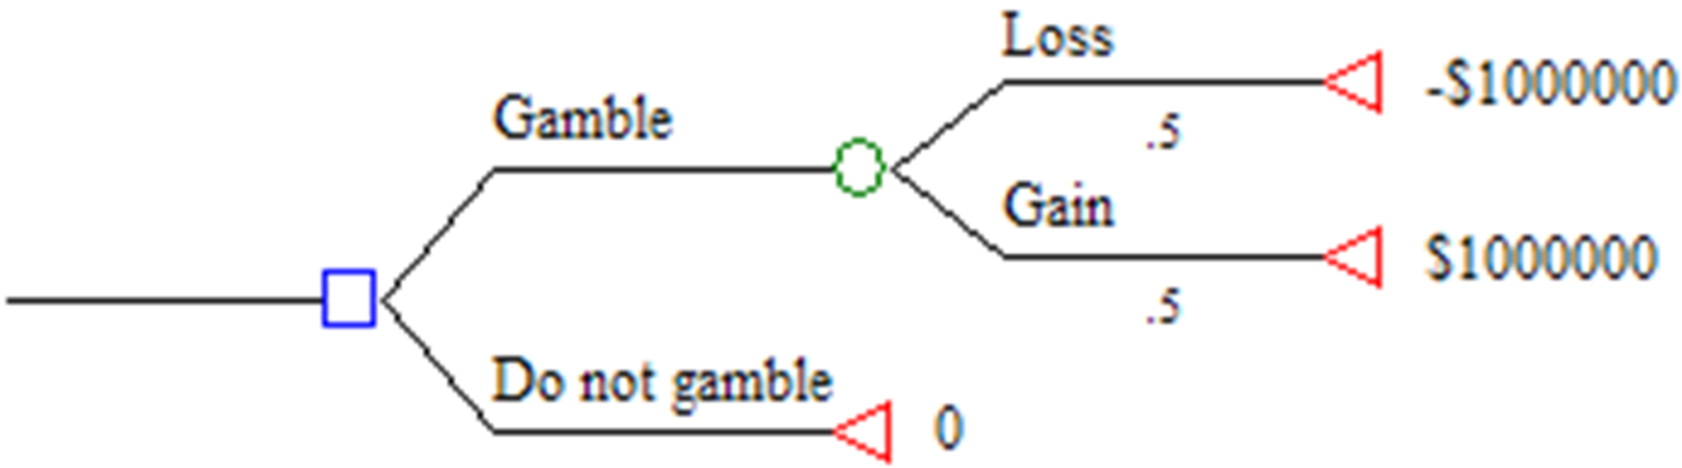

Categories of Risk Attitudes

- Risk attitude is a general way of classifying risk preferences
- Classifications
  - Risk averse fear loss and seek sureness
  - Risk neutral are indifferent to uncertainty
  - Risk lovers hope to "win big" and don't mind losing as much
- Risk attitudes change over
  - Time
  - Circumstance

Preference Function

- Formally expresses a particular party's degree of preference for (satisfaction with) different outcomes ($, time, level of conflict, quality…)
- Can be systematically derived
- Used to identify best decision when have uncertainty with respect to consequences
  - Choice with the highest mean preference is the best strategy for *that particular party*

Challenge:
Identify these Preference Functions

- (On the Board)

Risk Attitude in Preference Fns

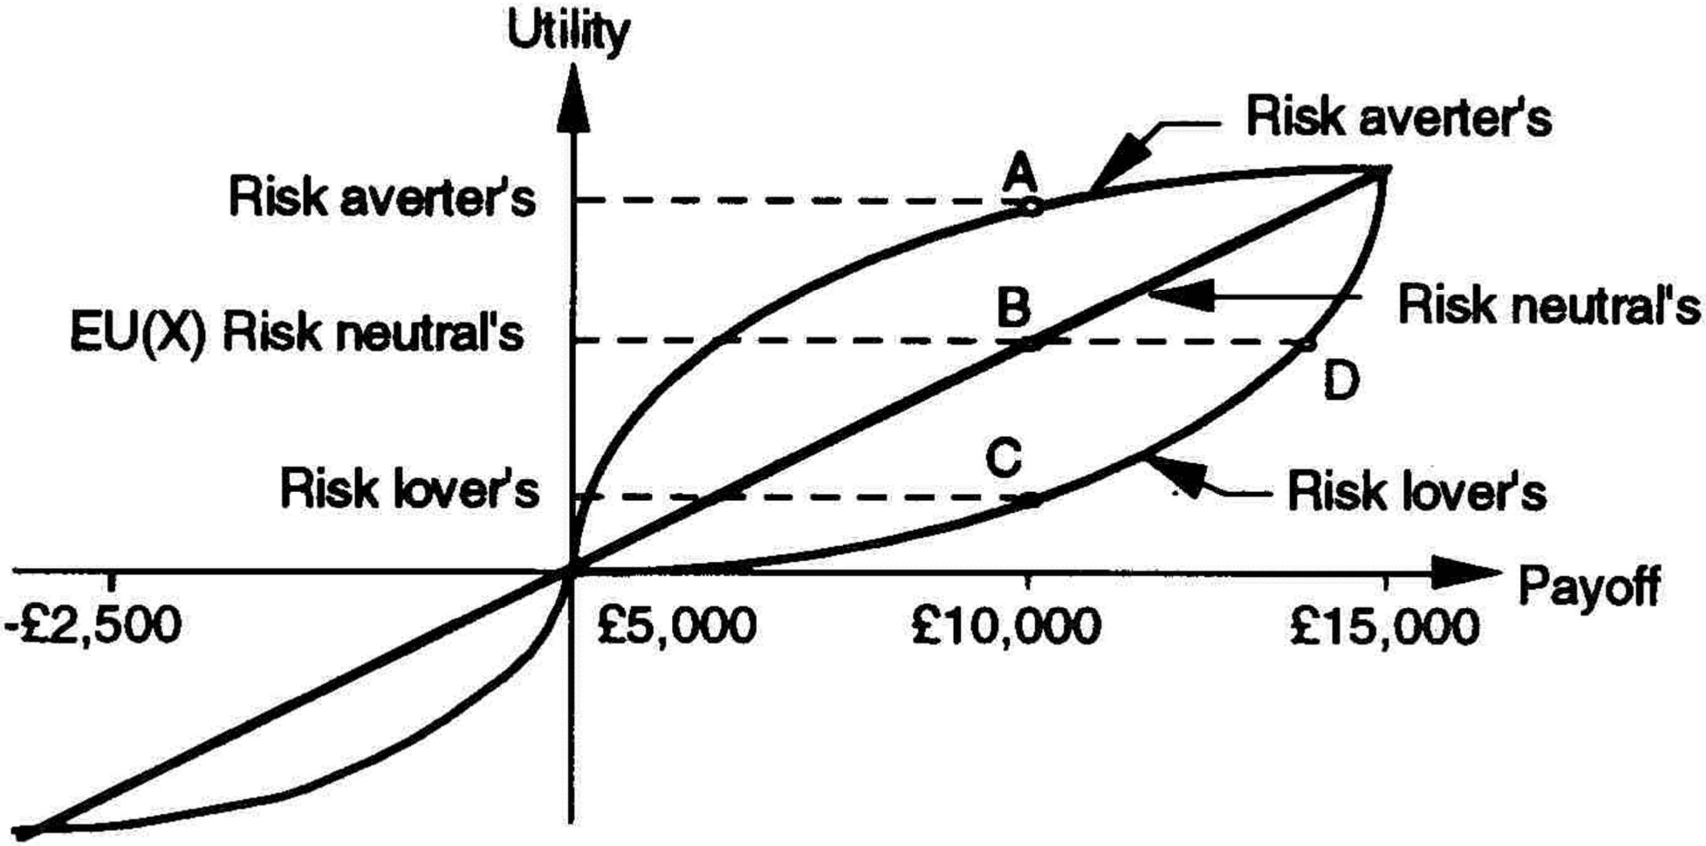

Identifying Preference Functions

- Simple procedure to identify utility value associated with multiple outcomes
- Interpolation between these data points defines the preference function

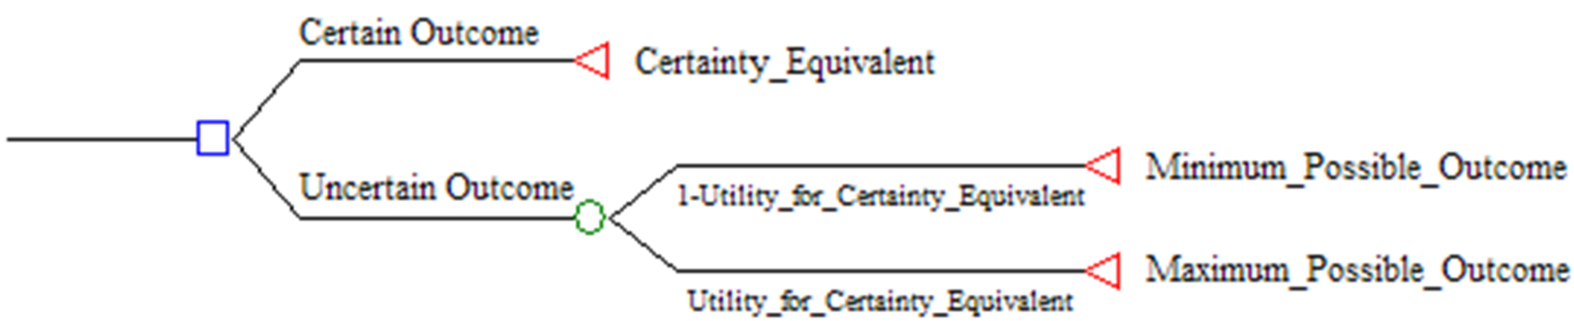

Notion of a Risk Premium

- A risk premium is the amount paid by a (risk averse) individual to avoid risk
- Risk premiums are very common – what are some examples?
  - Insurance premiums
  - Higher fees paid by owner to reputable contractors
  - Higher charges by contractor for risky work
  - Lower returns from less risky investments
  - Money paid to ensure flexibility as guard against risk

Certainty Equivalent Example

- Consider a risk averse individual with preference fn *f* faced with an investment c that provides
  - 50% chance of earning $20000
  - 50% chance of earning $0
- Average *money* from investment =
  - .5*$20,000+.5*$0=$10000
- Average *satisfaction* with the investment=
  - .5*f($20,000)+.5*f($0)=.25
- This individual would be willing to trade for a *sure* investment yielding satisfaction>.25 instead
  - Can get .25 satisfaction for a sure $f^{-1}(.25)$=$5000
    - We call this the *certainty equivalent* to the investment
  - Therefore this person should be willing to trade this investment for a sure amt of money>$5000

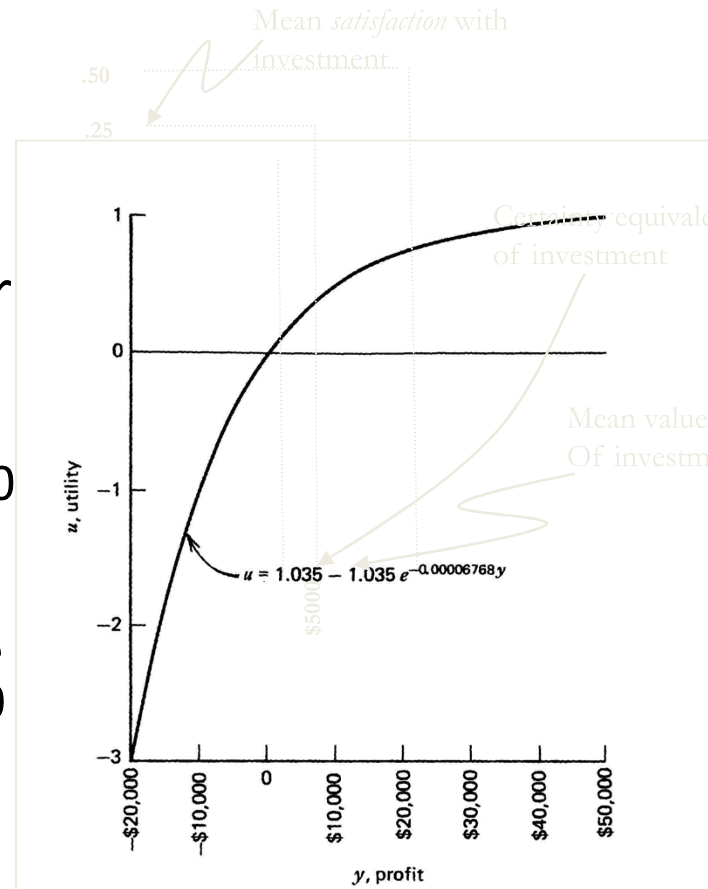

Figure 7.14. Utility function for profit.

Example Cont'd (Risk Premium)

- The risk averse individual would be willing to trade the uncertain investment c for any certain return which is > \$5000
- Equivalently, the risk averse individual would be willing to pay another party an amount *r* up to \$5000 =\$10000-\$5000 for other less risk averse party to guarantee \$10,000
  - Assuming the other party is not risk averse, that party wins because gain *r* on average
  - The risk averse individual wins b/c more satisfied

Certainty Equivalent

- More generally, consider situation in which have
  - Uncertainty with respect to consequence *c*
  - Non-linear preference function *f*
- Note: E[X] is the mean (expected value) operator
- The mean *outcome* of uncertain investment c is E[c]
  - In example, this was .5*$20,000+.5*$0=$10,000
- The mean *satisfaction with* the investment is E[f(c)]
  - In example, this was .5*f($20,000)+.5*f($0)=.25
- We call $f^{-1}(E[f(c)])$ the *certainty equivalent* of c
  - Size of *sure* return that would give the same satisfaction as *c*
  - In example, was $f^{-1}(.25)=f^{-1}(.5*20,000+.5*0)=\$5,000$

Risk Attitude Redux

- The shapes of the preference functions means can classify risk attitude by comparing the certainty equivalent and expected value
  - For risk *loving* individuals, $f^{-1}(E[f(c)])>E[c]$
  - For risk *neutral* individuals, $f^{-1}(E[f(c)])=E[c]$
  - For risk *averse* individuals, $f^{-1}(E[f(c)])<E[c]$

Motivations for a Risk Premium

- Consider
  - Risk averse individual A for whom $f^{-1}(E[f(c)])<E[c]$
  - Less risk averse party B
- A can lessen the effects of risk by paying a risk premium r of up to $E[c]-f^{-1}(E[f(c)])$ to B in return for a *guarantee* of $E[c]$ income
  - The risk premium shifts the risk to B
  - The net investment gain for A is $E[c]-r$, but A is more satisfied because $E[c] - r > f^{-1}(E[f(c)])$
  - B gets average monetary gain of r

Multiple Attribute Decisions

- Frequently we care about multiple attributes
  - Cost
  - Time
  - Quality
  - Relationship with owner
- Terminal nodes on decision trees can capture these factors – but still need to make different attributes comparable

WNV Hybrid Approach

SD Model

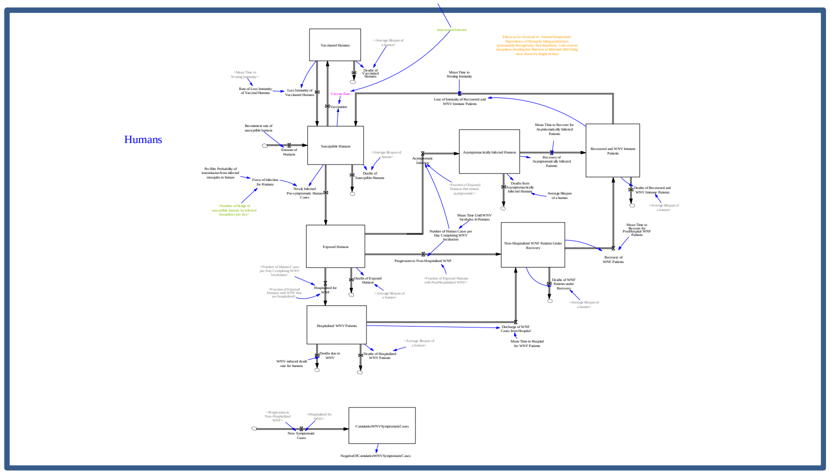

Decision Tree

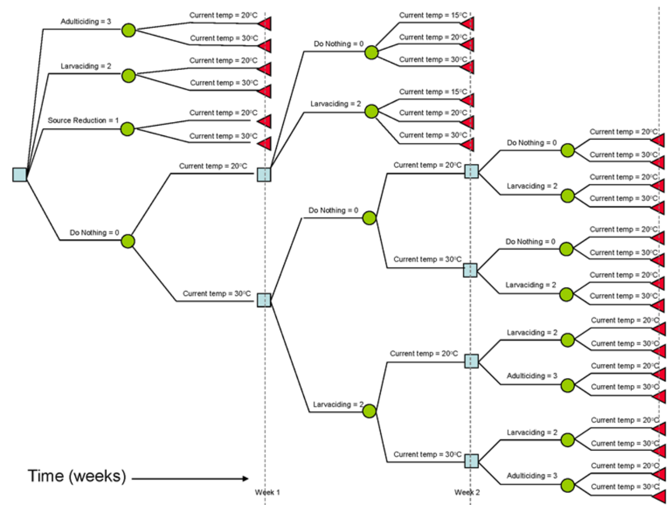

User Interface

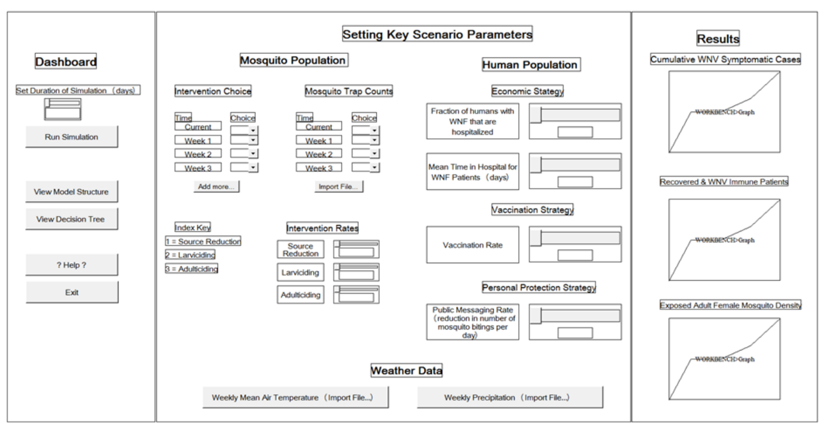

The Hybrid Approach: Critical Points

1. Is a framework geared toward an ongoing process of observation & decision making
2. Captures uncertainties as time progresses
3. Simulates a broad range of possibilities (e.g. for temperature) and not just a single scenario
4. Allows for staging of decisions over different time points – including decisions to “wait & see” (exploiting future options)
5. Could be used for diverse planning challenges (e.g. H1N1 given uncertainty regarding public reaction, vaccine availability)

Responsibilities in the Hybrid Approach

Simulation Model

- Calculates dynamic consequences of a sequence over time of
  - Events
  - Choices
- Takes care of deterministic simulation given events & decisions

Decision Tree

- Represents over time possible sequences of
  - Uncertainties (event nodes)
  - Decisions (decision nodes)
- Consequences (outcomes – e.g. Cost, quality of life, etc.)
- Takes care of encapsulating
  - Capturing all uncertainties
  - "policy space" – where policies are made over time

Example: WNV Hybrid Approach

Simulation Model

- Mosquito lifecycle (includes temperature effects)
- Bird lifecycle
- Transmission between mosquitos & bird
- Human infection & disease progression
- Future: costs & resource use (via resource intensity weights, length of stay), quality of life

Decision Tree

- Decision options over time (source reduction, larvaciding, vaccination, wait & see)
- Uncertainties (temperature)
- Consequences (all WNV cases, severe neurological cases, costs, etc.)

West Nile Virus SD Model

Mosquitoes

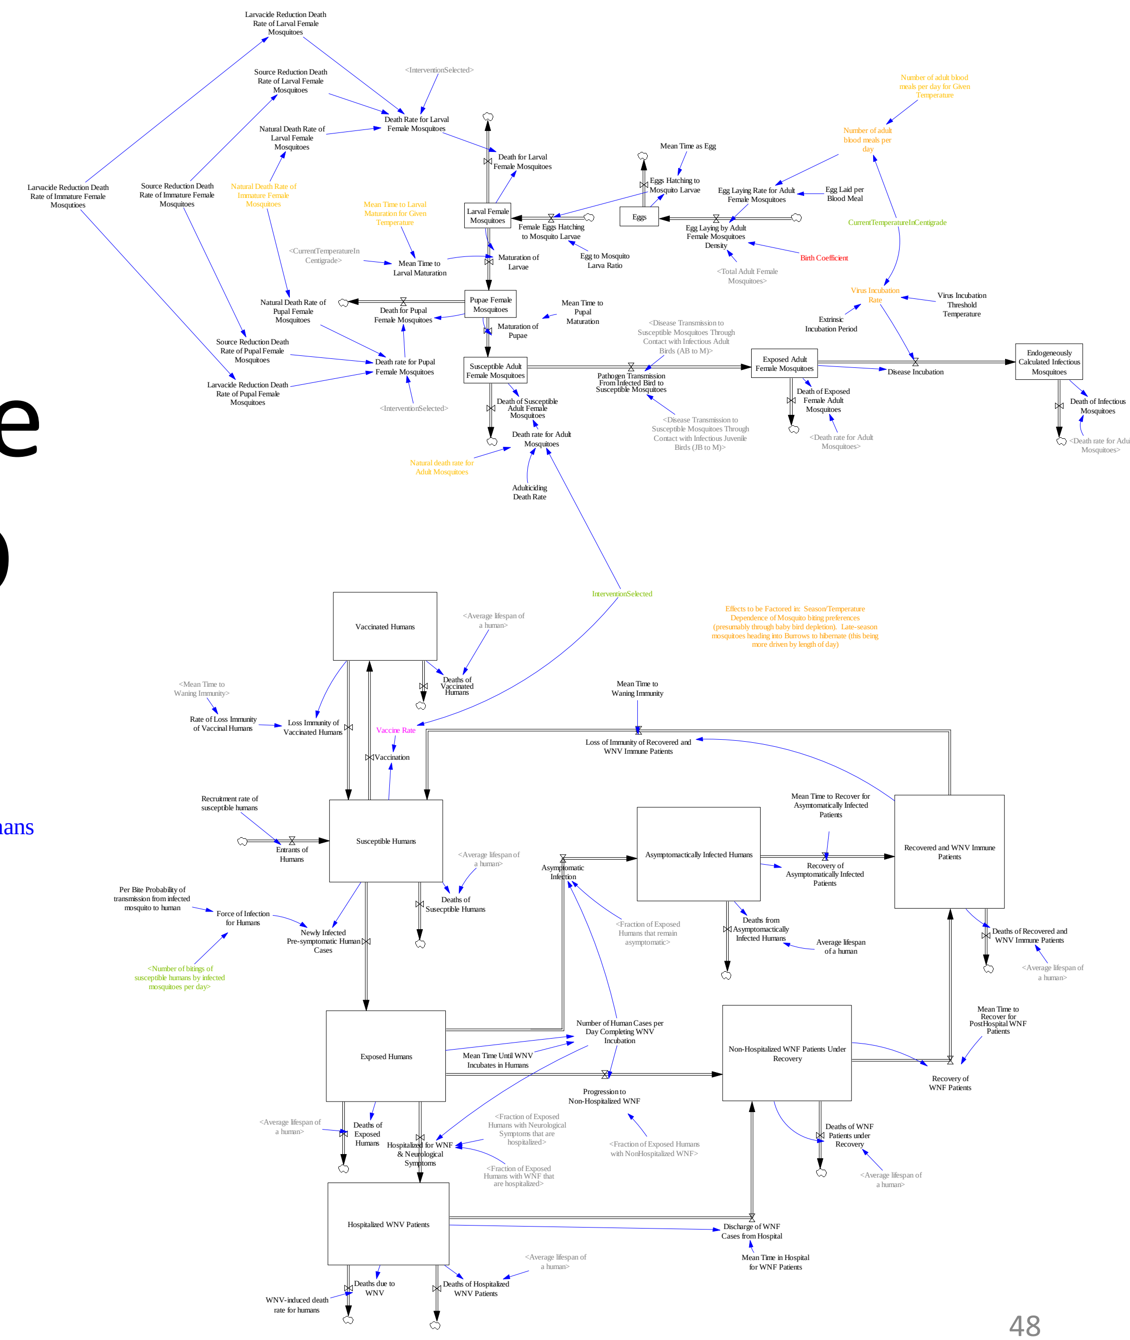

Humans

[back](#)

Decision Tree

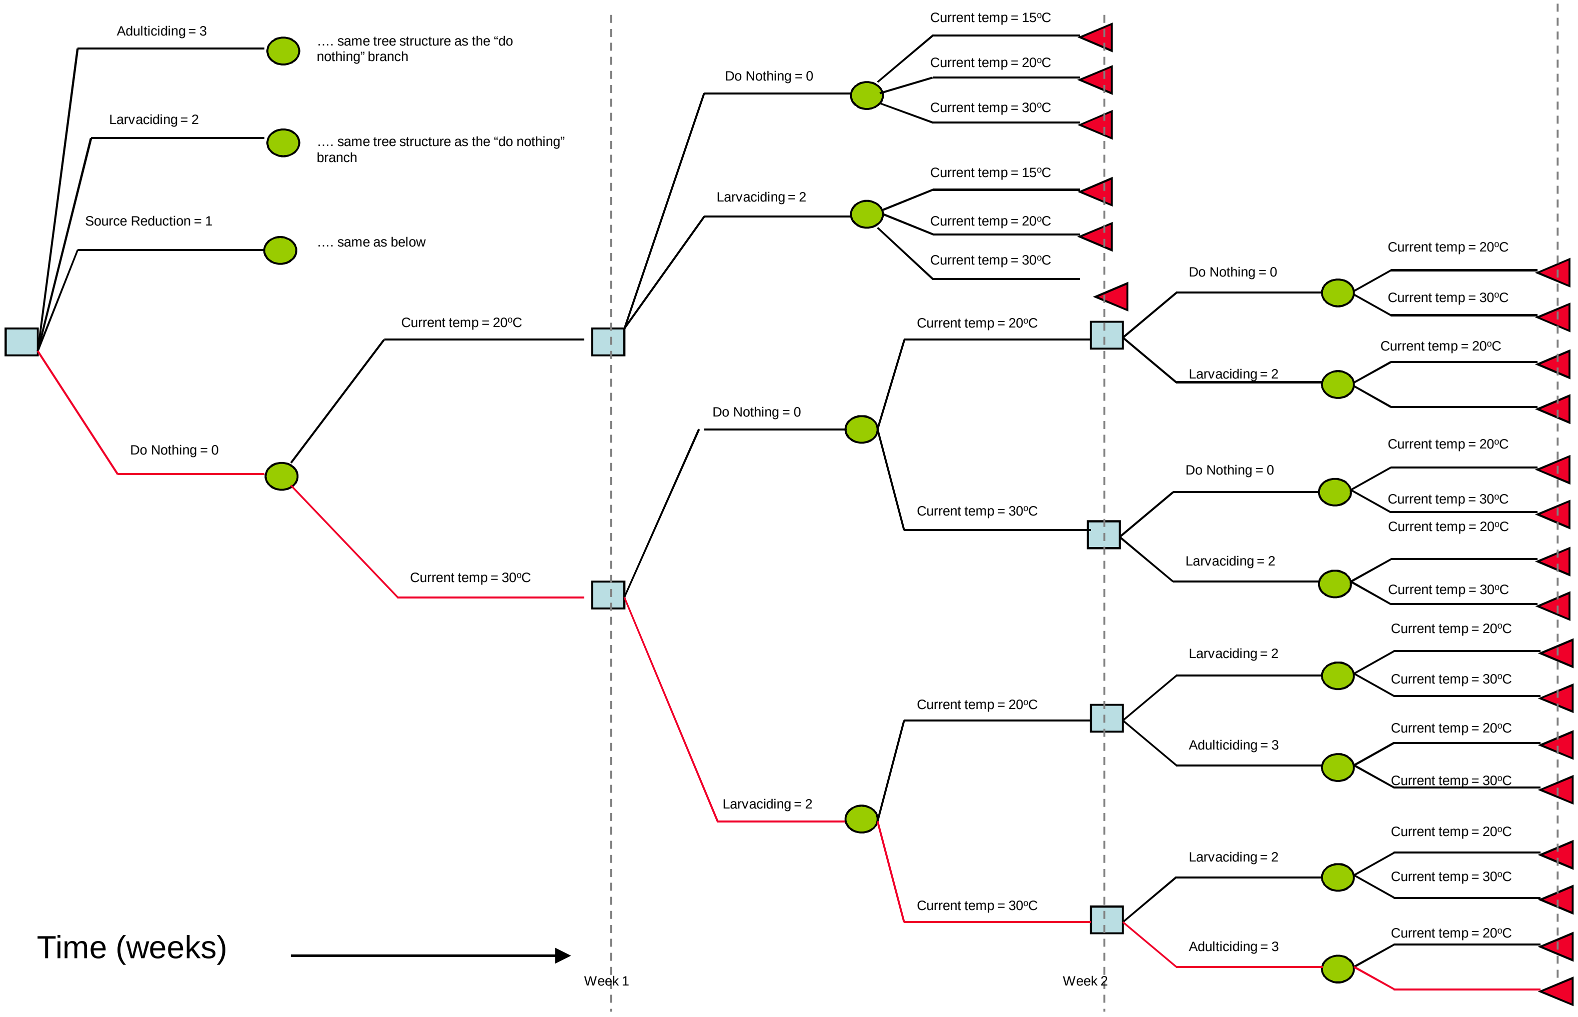

back