Week 10: Markov Models I

Marcelo Coca Perraillon

University of Colorado
Anschutz Medical Campus

Cost-Effectiveness Analysis
HSMP 6609
20120

Outline

- Elements of Markov models
- Example: HIV transitions
- State transition diagrams
- Adding costs and benefits
- Simulating the transitions of a cohort
- Half-cycle correction
- The memoryless property of cohort models
- Markov models as "trees"
- When should we use Markov models?

Big picture

- Last week we covered decision trees as they are used in CEA
- The new element was the introduction of **uncertainty** into the calculation of ICERs
- The ICER becomes an **expected value**
- We also saw that decision models are not explicit about time and that they get **too complicated if events are recurrent**
- Markov models solve these problems
- **Confusion alert**: Keep in mind that Markov models can be illustrated using "trees." Also, decision trees and Markov models **are often combined**. I'll get back to this later in the class

Elements of Markov models

- The first element of Markov model are so-called **health states**, such as well, ill, dead, relief, no relief, headache, no headache
- People **transition** from one health state to the another. For this reason, Markov models are sometimes called **transition models**
- Each transition has a probability (**transition probability**)
- Transitions happen over a period of time (called a **cycle**)
- Each health state can have a **cost** and/or **benefit** associated to it (called **rewards** in Markov models; the equivalent of payoffs in decision model)
- And that's all the key elements we need

Example

- We;'ll use a now-classic example from your textbook and Briggs et al (2006) (available on Canvas)
- Two therapeutical strategies for HIV: zidovudine monotherapy and zidovudine in combination with lamivudine (for simplicity, "monotherapy" versus "combination" therapy)
- Four possible health states, some of them depending on CD4 counts:
  1. State A: CD4 from 200 to 500 (best)
  2. State B: CD4 less than 200 (not so good)
  3. State C: AIDS (bad)
  4. State D: Death (very bad)
- Cycle length is **one year**
- We can illustrate the states and all possible transitions with a **state transition diagram**

HIV state transition diagram

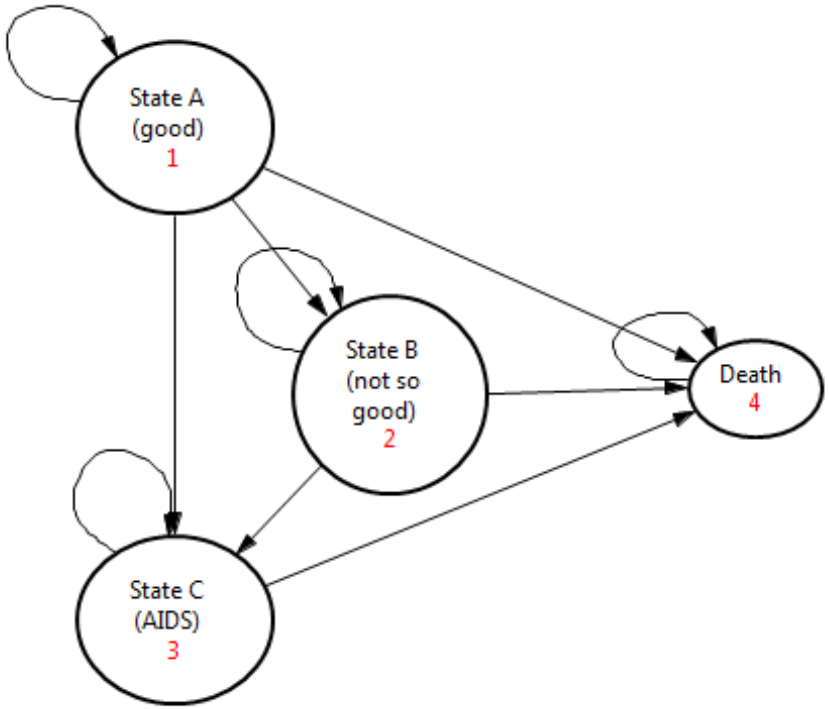

- All paths lead to death and death leads to... nowhere (aka **absorbing** state). Everybody starts in state A
- Note that in this example there is no way to get better (no longer true for HIV)
- Possible to remain in same state after a cycle

Transition probabilities (matrix)

Monotherapy

	State A	State B	State C	Death
State A	0.721	0.202	0.067	0.01
State B	0	0.581	0.407	0.012
State C	0	0	0.75	0.25
Death	0	0	0	1

Combination

	State A	State B	State C	Death
State A	0.858	0.103	0.034	0.005
State B	0	0.787	0.207	0.006
State C	0	0	0.873	0.127
Death	0	0	0	1

- Read them from left to right: probability of transitioning from C to Death is 0.25
- Note that **combo therapy has better outcomes**. To be precise, combo therapy's risk is reduced by about half (0.509)
- See Excel file for actual probabilities

Transition probabilities (matrix)

Monotherapy

	State A	State B	State C	Death
State A	0.721	0.202	0.067	0.01
State B	0	0.581	0.407	0.012
State C	0	0	0.75	0.25
Death	0	0	0	1

Combination

	State A	State B	State C	Death
State A	0.858	0.103	0.034	0.005
State B	0	0.787	0.207	0.006
State C	0	0	0.873	0.127
Death	0	0	0	1

- Note that probabilities must up to one (horizontally): $0.721 + 0.202 + 0.067 + 0.01 = 1$ (always check!!)
- Note that this implies that we always must "account" for people transitioning. They can either stay in the same state or they must go somewhere in each cycle
- Similar to decision trees, probabilities are exhaustive and mutually exclusive
- Note death. Death is final. A person who gets there stays there

Transition probabilities (matrix)

Monotherapy

	State A	State B	State C	Death
State A	0.721	0.202	0.067	0.01
State B	0	0.581	0.407	0.012
State C	0	0	0.75	0.25
Death	0	0	0	1

Combination

	State A	State B	State C	Death
State A	0.858	0.103	0.034	0.005
State B	0	0.787	0.207	0.006
State C	0	0	0.873	0.127
Death	0	0	0	1

- Note death. Death is final. A person who gets there stays there
- Note the zeroes. This implies that a transition is not allowed. Nobody can get better in this example. Nobody can go from state C to B, or B to A
- The transition matrix agrees with the transition diagram

Where do probabilities come from?

- As with decision models, probabilities come from clinical trials, observational data, meta-analyses, expert panels, surveys...
- Note that we could have used a decision tree instead
- More precisely, a recursive decision tree with one tree per year but it would be too complicated (the "bushy" tree problem)
- Now we need to add **rewards** (i.e. payoffs)

The language of trees vs Markov models

- Now probabilities are called **transition** probabilities. Now events are **states**. There are no "branches;" but going from one state to another has a probability
- Now there is time: everything happens in a **cycle**, which can be 1 week, 1 year, etc
- Payoffs are now **rewards** which are part of each state (more in a second)
- Long history in statistics. Introduced by Andrew Markov in 1906
- **Careful when googling**. We are covering Markov or transition models, which are examples of a Markov process. But many things come under the name "Markov process." Same with decision trees. **Decision trees in machine learning have nothing to do with decision trees in decision theory**. They are actually regression trees, not decision trees. The similarity is that in both cases you can draw something that sort of looks like a tree. That's where the similarity ends

Costs

- We will first do a cost analysis (we will add life years later)
- The HIV study collected health care, community, and medication costs
- Drug costs: zidovudine (£2,278); lamivudine (£2,086); combination (£4364)
- Cost per state (health care and community): A: £2756; B: £3052 C: £9007; D: £0
- Costs are **per cycle** (i.e. year)
- Keep in mind a couple of things: combination therapy is more expensive but also more effective (first quadrant in the cost-effectiveness plane)
- The worse the state the higher the costs, but once in state "death" there is no cost. If end-of-life costs were significant, we could have added another state

"Solving" Markov models

- We now have all ingredients: **health states**, **transition probabilities**, and **rewards** for each cycle
- "Solving" the Markov model means that we will **simulate** what would happen to a group (cohort) of people over time (that's why they are called **cohort models** sometimes)
- We will then calculate **expected costs**
- In this example, we will simulate the transitions of 1,000 patients in each type of therapy over **20 years**
- Why 20 years? Because most of our population won't live longer. So it's really a **lifetime time horizon**
- **Pay attention!** It's actually **very easy** but it's also very easy to get confused

Solving Markov HIV model

- Monotherapy: we start with 1000 patients in the healthy state A
- After 1 cycle, what will happen?
- We know that the probability of staying in state A is 0.721, so afer one year, 721 will remain in state A (1000×0.721)
- Same logic for all other states 202 in B, 67 in C, and 10 in D
- See, **super easy**. We are just "allocating" the 1,000 people into the four possible states after one cycle base on the probabilities of transition from one state to the other

Solving Markov HIV model

Monotherapy

	Proportion of Cohort in each year			
Year	**A**	**B**	**C**	**D**
0	1000			
1	721	202	67	10

- Now we need to repeat the process for the next cycle
- Of the 721 in state A, how many will stay in A? **520** (0.721×721)
- How many in B? We need to take into account that some will move from A to B but also that some in B will stay in B: $721 \times 0.202 + 202 \times 0.581 = 263$ (**here is where you are likely to make mistakes**)
- Easier to see it graphically

Transition probabilities

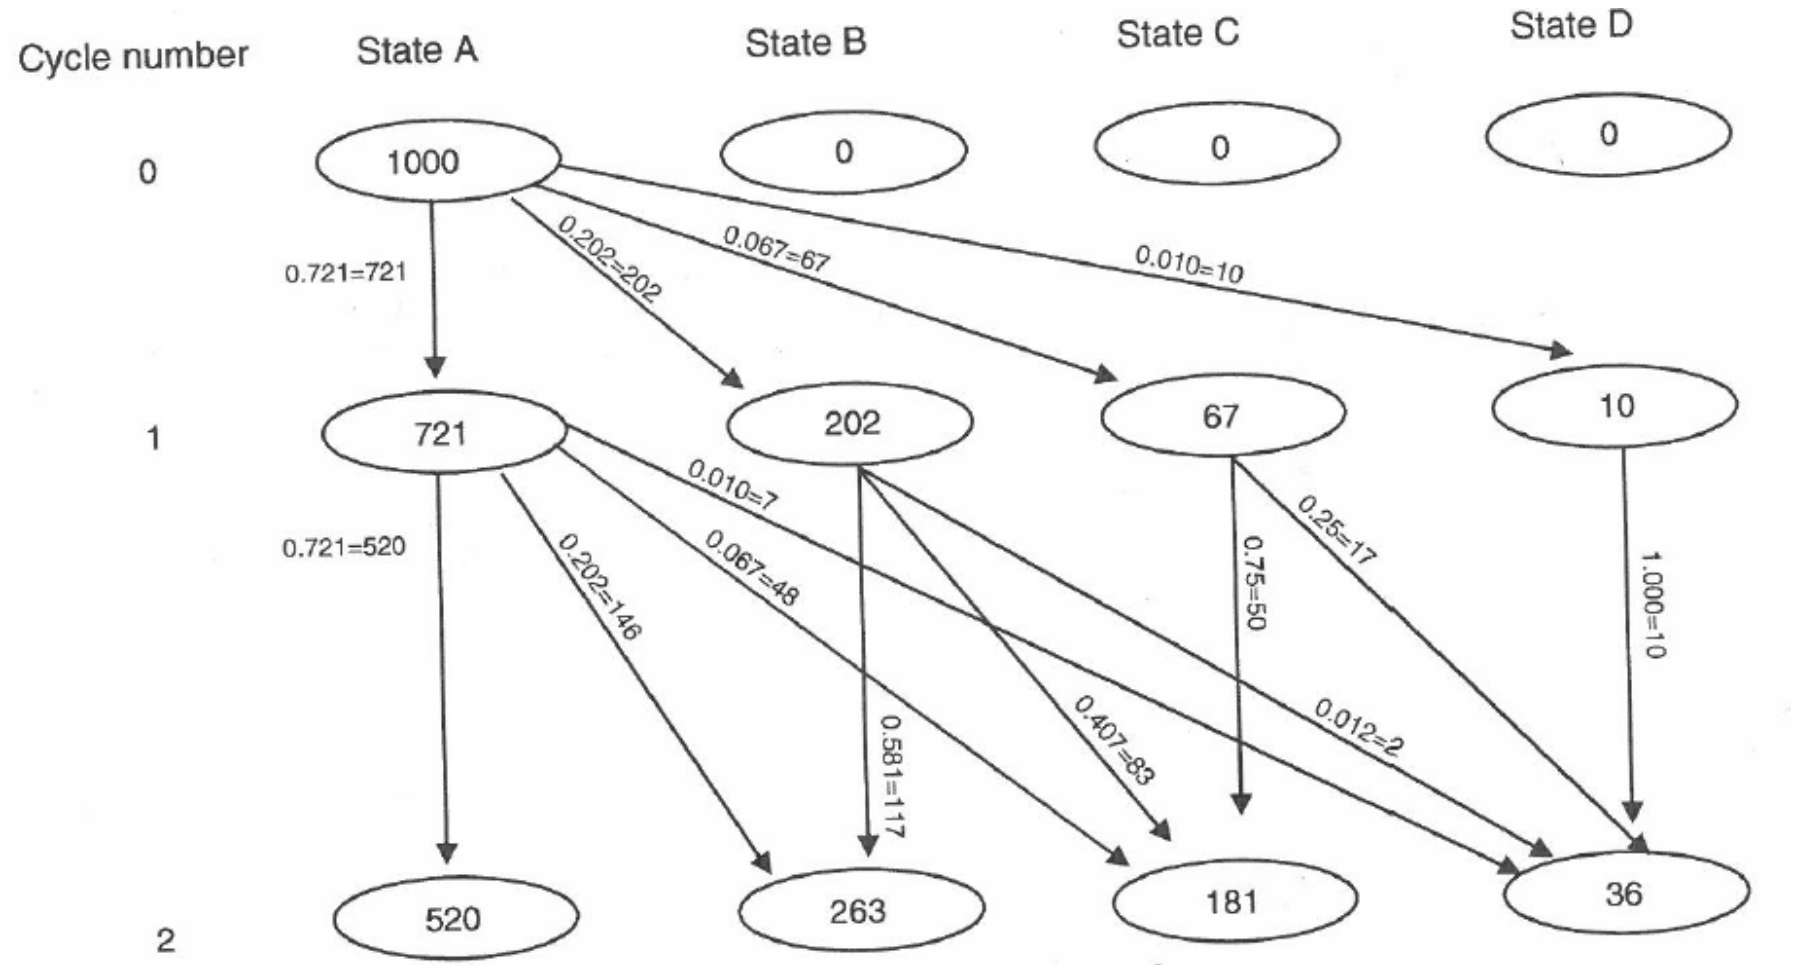

Solving Markov HIV model

Monotherapy

	Proportion of Cohort in each year			
Year	A	B	C	D
0	1000			
1	721	202	67	10
2	520	263	181	36

- After the second year, copy-pasting formula in Excel will do it
- Just be careful and make sure some of the cells' references are "fixed" (press F4 or double $ sign)
- See Excel file for this lecture

Big picture

- We start with a group of patients (the number doesn't matter, we could use 1 person, but you need to use the same number of people in both or do it by person)
- Cycle by cycle, we transition them to different health states
- **That's it**. Really, that's all. We have simulated **disease progression**
- Note that we are transitioning a group of patients. We are **NOT following each patient** (more on this later as it becomes very important)
- We are moving a cohort into the simulation. Each cycle they accumulate costs and life

Costs

- Now we know what will happen to our cohort over time. What about costs?
- We just need to multiply the number of people in each health state by the corresponding cost in each cycle
- For example, in cycle 2:

  $520 * (2756 + 2278) + 263 * (3052 + 2278) + 181 * (9007 * 2278) + 36 * 0 = 6,056,510.75$
- Per person: £6060

Life years gained

- We are simulating 20 years of life for these patients
- Some of them die (for example, 10 in the first cycle, 26 in the second)
- **We can therefore calculate life years** in each cycle
- In the first cycle, 10 people died but 990 were alive so these "alive" people accumulated 990 life years in the first cycle
- In the second cycle, $520 + 263 + 181 = 964$ were alive, so they accumulated another 964 years of life
- We do the same for each cycle
- At the end of the 20 cycles, we add up all the years of life over the 20 years

We can plot a survival curve

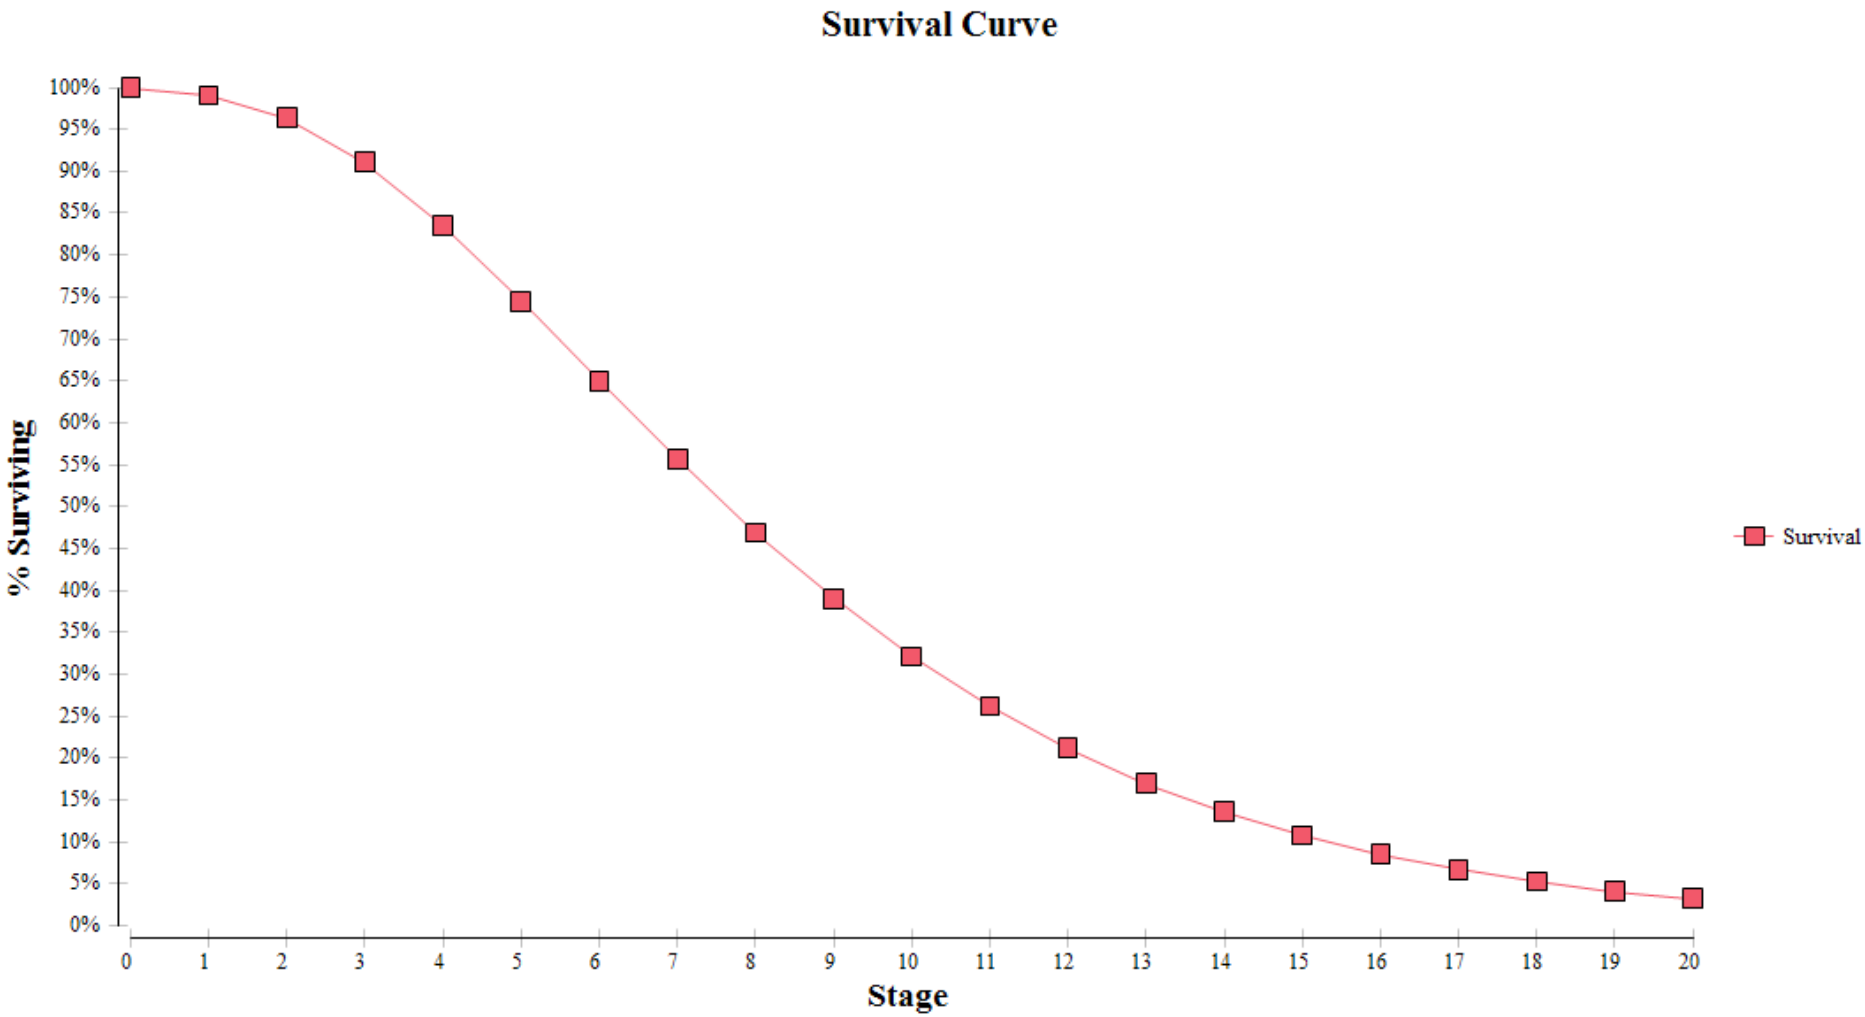

Half-cycle correction

- Why didn't we give any life years "credit" to those who died **during** each cycle?
- Because we essentially assumed that people died at the **start of the cycle**
- But we should take into account that patients die at different times; otherwise, we underestimate costs and benefits
- Not a big problem if we do the same in both treatments (the shorter the cycles the less of a problem)
- The best solution (the **unbiased** solution) is to assume that patients die in the **middle** of the cycle
- This is a result of assuming that patients **die (uniformly during the year**. That is, dying follows a uniform distribution. Or said another way, the probability of death is the same every day during the year

Half-cycle correction

- Check out the formula in the Excel example
- We add half the time of those who died during the year
- We won't worry about the half-cycle correction for the rest of the semester

CEA

- We now have costs and life years gain over the time horizon
- We need to repeat the same simulation for another cohort for the combination group (homework)
- The way I worked out this example, we need to simulate 1,000 patients in the combination group
- Not really necessary. We could calculate costs and benefits for any number or do it by person
- After calculating costs and benefits for both groups, we can obtain the ICER

What about QALYs and other features of CEAs?

- Note that **adjusting for quality** (i.e. preferences) is straightforward
- Each health state would have a preference score (the number between 0 and 1) associated with it
- We just multiply the score by the time spent in each year
- **Discounting** is easy too: we have costs and benefits per cycle so we just need to bring them into the present
- Note that with Markov models we can go from **intermediate** to **final** outcomes since we're modeling disease progression
- Note that we're ignoring something important in this basic example: **people getting older but their chance of dying is not changing**
- Easy to incorporate: increase the chances of dying of other causes by cycle (next class)

Department of Pesky Things that Cause Unnecessary Confusion

- Be careful when you read CEA papers

  1) Markov models are depicted using transition diagrams, but they can also be depicted using something that looks like a tree, but it's not a decision tree

  2) Decision trees and Marvok models can be combined
- You can have a Markov models within a decision tree

Markov models depicted **as** trees

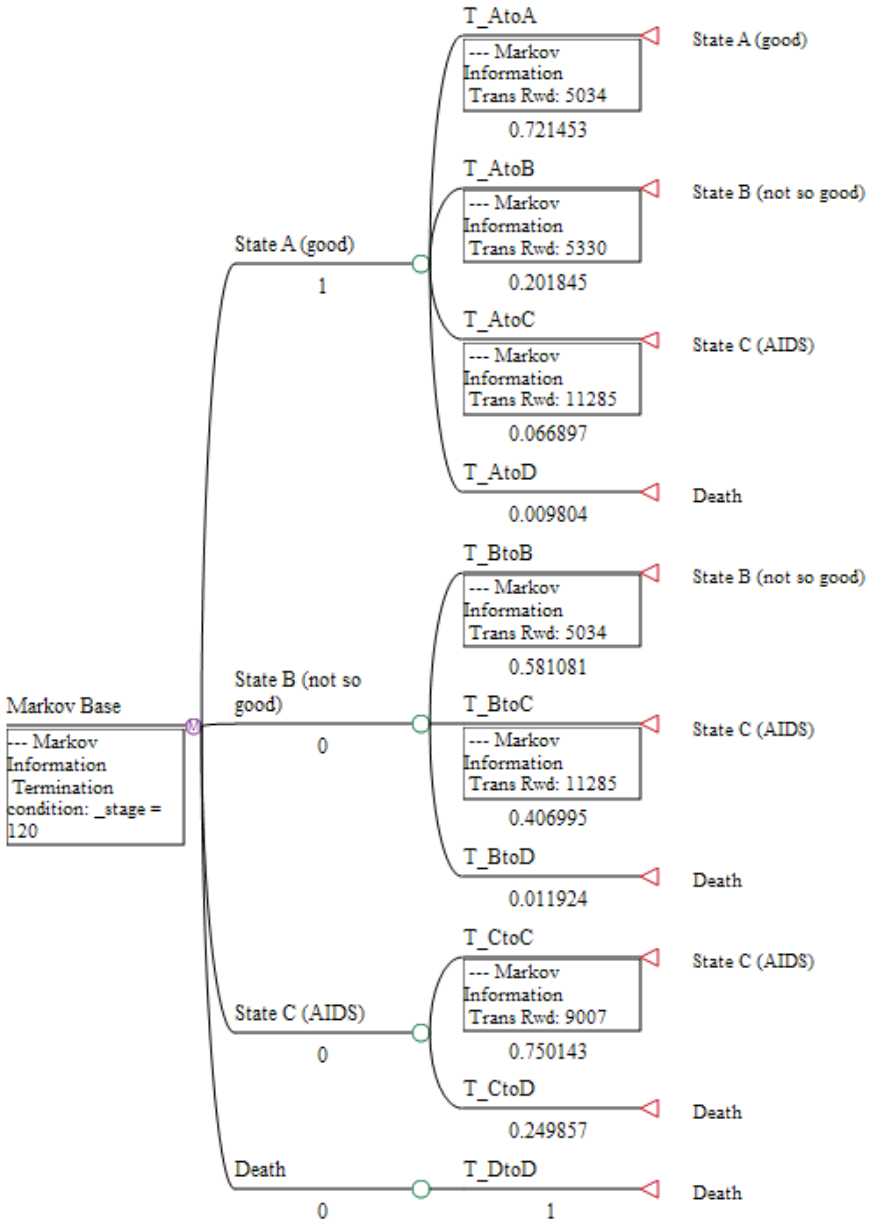

Markov models **inside** decision trees

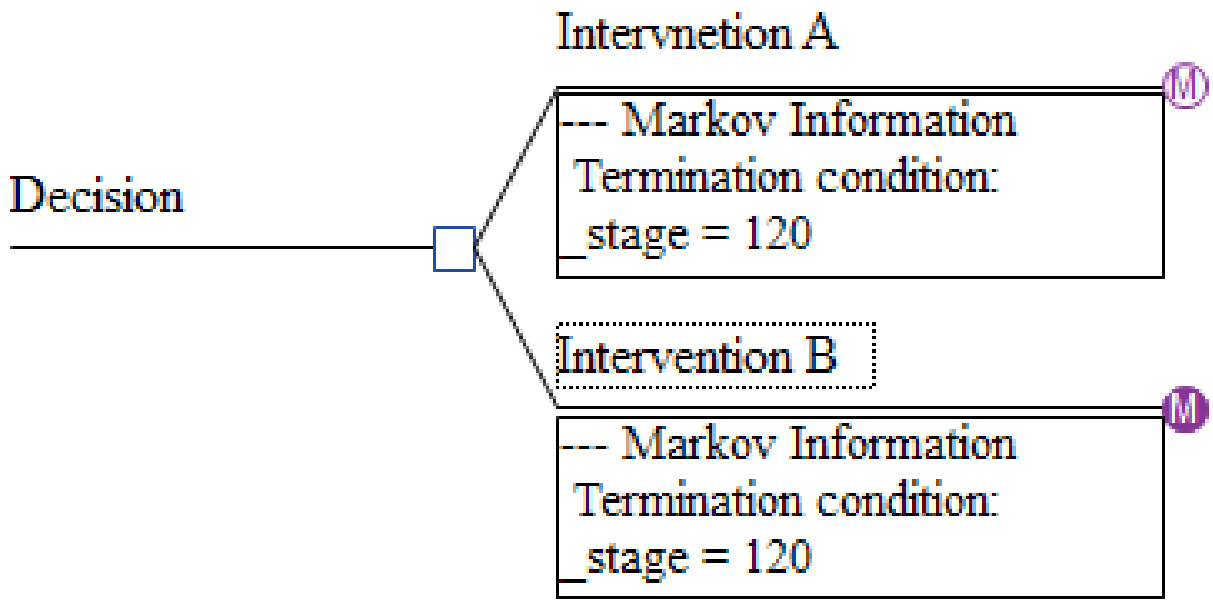

- Some parts of a decision tree could be calculated using Markov models

Big picture

- Markov models and decision models are not that different
- We have changed the language because decision trees and Markov models have different origins
- Health states could be represented by a branch in a decision tree
- In the migraine example, health states could be relief, no relief, no recurrence, recurrence, hospitalization...
- The main difference is the introduction of time in the form of **cycles** and that **recurrent events** are easily modeled
- In the migraine example, an appropriate cycle could be a day or a week

The **memoryless** property or the **Markov assumption**

- One limiting assumption of cohort Markov models is that transitions to a state **do not depend on the past** or the **time a patient has been** in a state
- In other words, once in a cycle, there is no "memory" of the past
- In many cases, how long a patient stays in a state affects the chances of an outcome
- For example, a person experiencing his third bout of depression has higher chances of worse outcomes
- There are ways to fix this limitation: adding additional transition states (second, third depression episode) and/or making transition probabilities conditional on past events
- In general, **this can be an important limitation** and we will see extensions

When should we use Markov models?

- When events are **recurrent**
- When we want to model the "**natural history**" of the disease
- Long time horizon: we want to go for intermediate outcomes to final outcomes
- When life years or QALYs are outcomes of interest
- **Markov models are the most commonly used tool in cost effectiveness**

Cycle length and some limitations

- **Important**: We haven't talked much about this but **choosing a valid cycle length matters**
- Ideally we want a cycle length in which two events usually do not happen
- Cohort Markov models are not good for modeling infectious diseases: the probability of infection depends on the number of people infected and (**herd immunity**)
- Cohort models modeling a group transitioning, but **can't follow a person**
- Next class we will see some extensions and tricks we can do to add more flexibility to Markov models but they do have limitations
- We will also cover a basic model of infectious disease using different tools but also using Excel

Summary

- Markov models allow us to model complex diseases
- Markov models better incorporate time and disease progression
- Simulating cohort models is **easy**
- As with decision analysis, the hard part is to come up with a model that isolates the key elements that need to be considered –**that's not easy**
- There are extensions to Markov models that are better for some problems (next week)

Next class

- More examples
- Incorporating time dependency
- Adding memory to the memoryless model
- Temporary states
- Other type of models in CEA