From Prediction to Policy

Joseph Rollinson
Advisor: Emma Brunskill

2 models

similar accuracy

different instruction

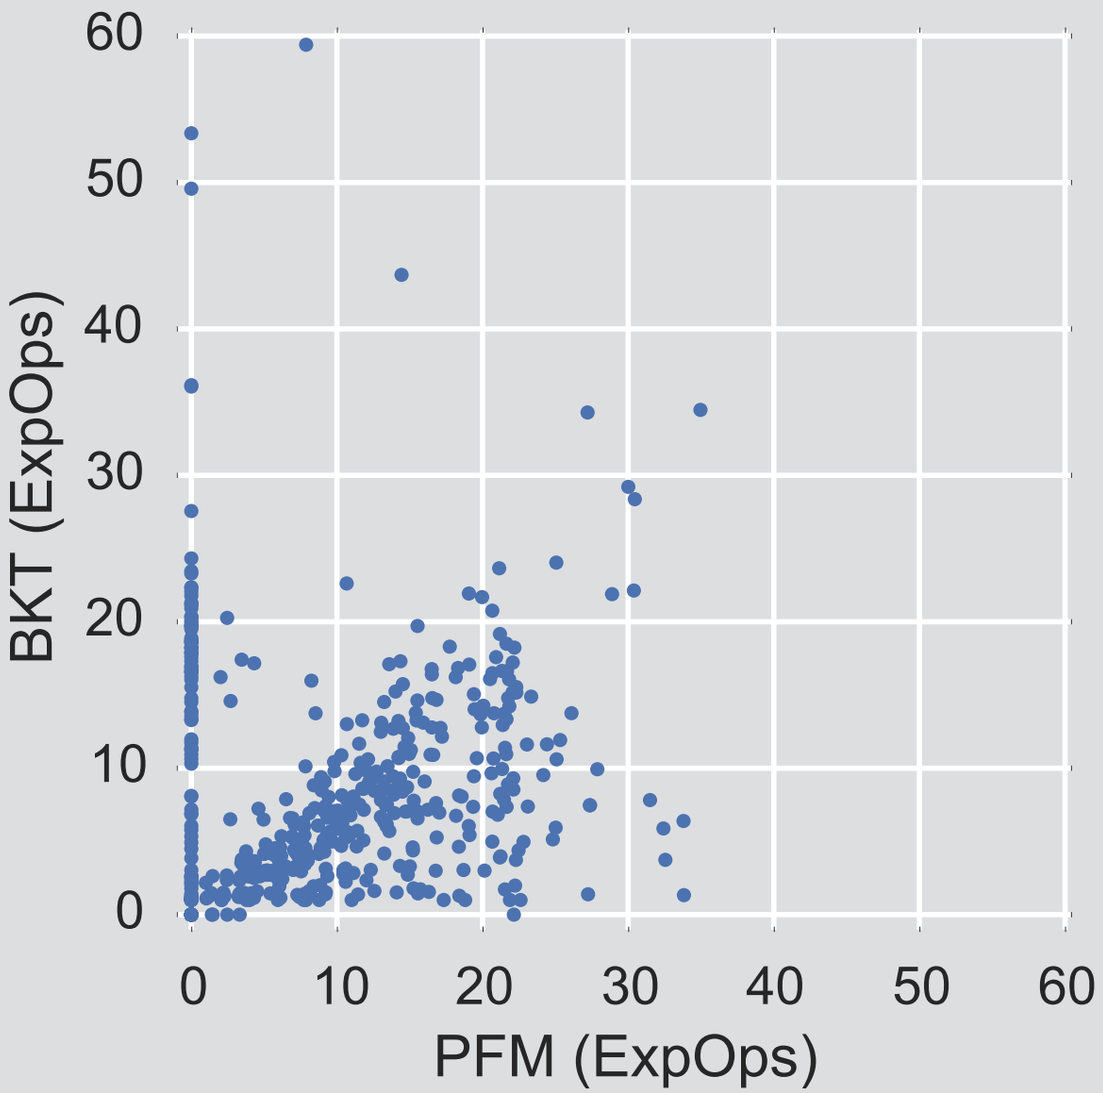

choosing multiple skills

looking ahead helps student learn more

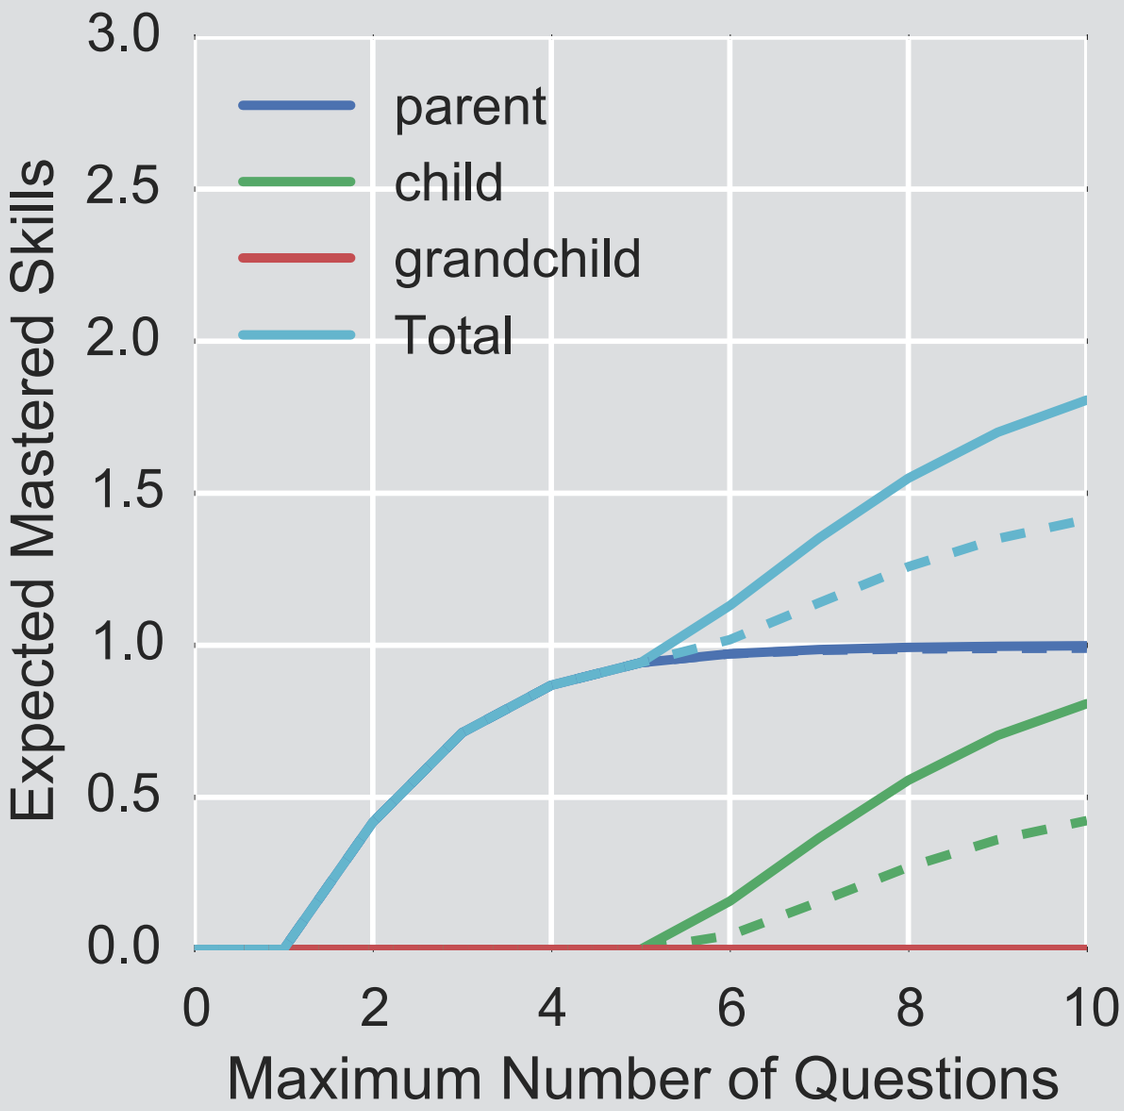

Why does this work matter?

Education Matters

Issues in Education:

1. Access
2. Quality

Solution: Intelligent Tutoring Systems (ITS)

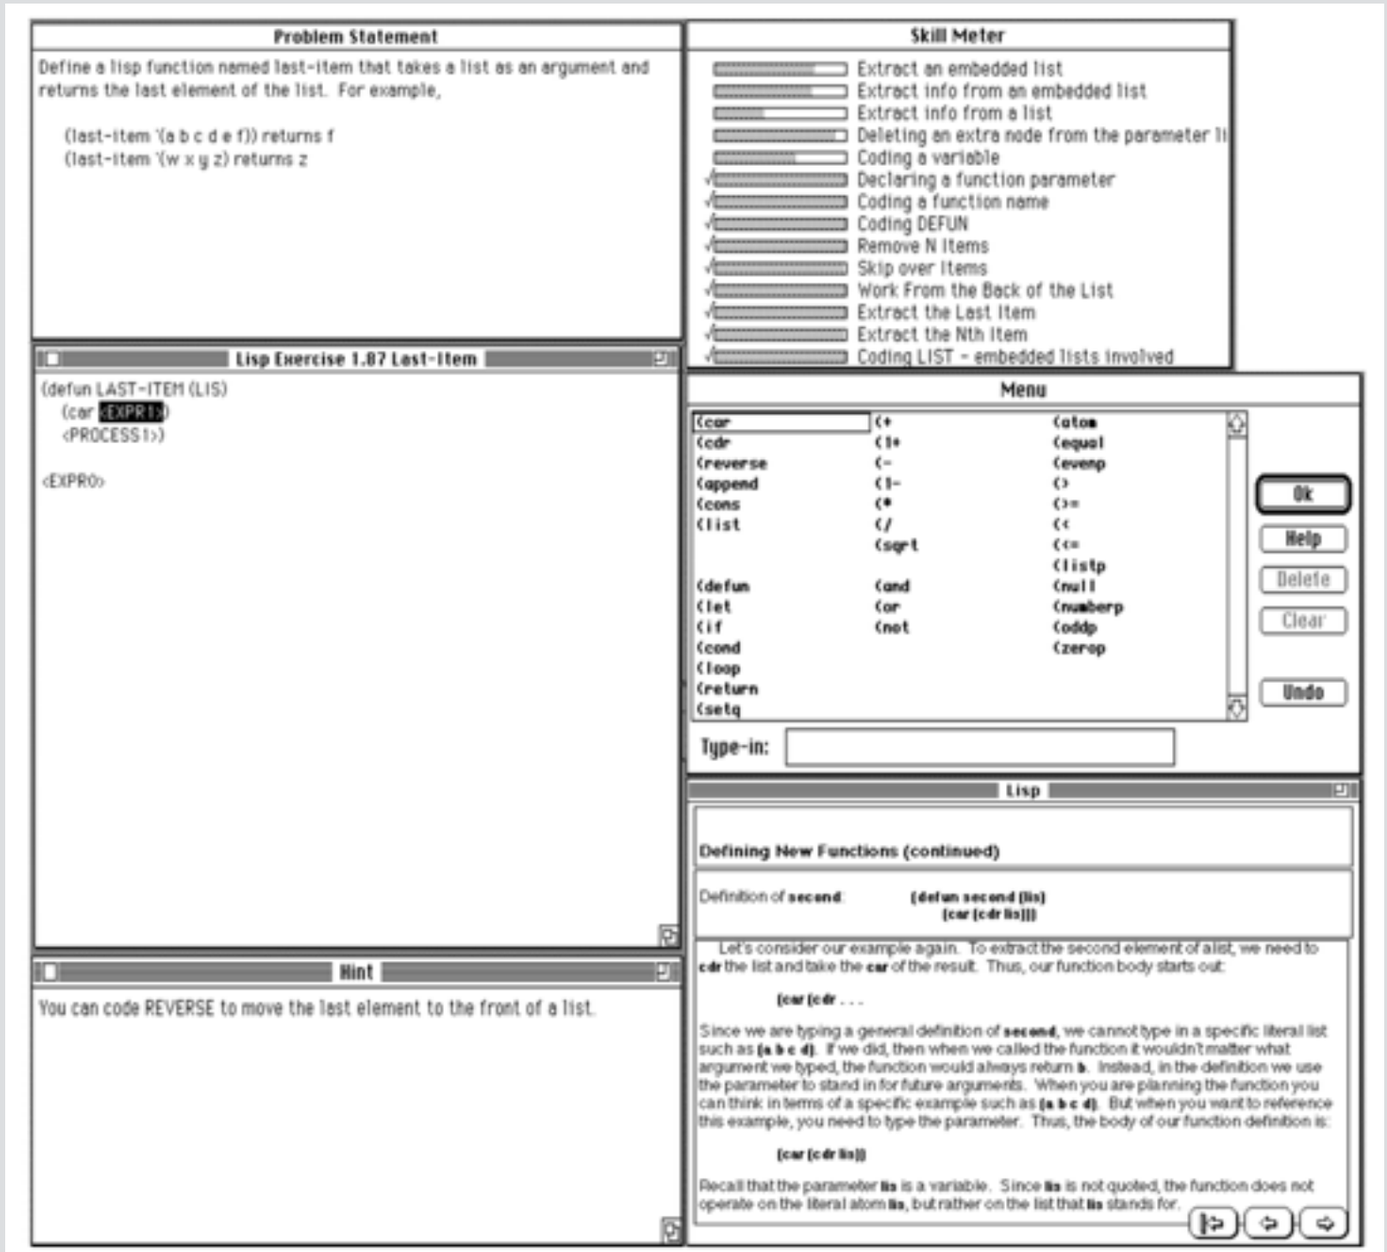

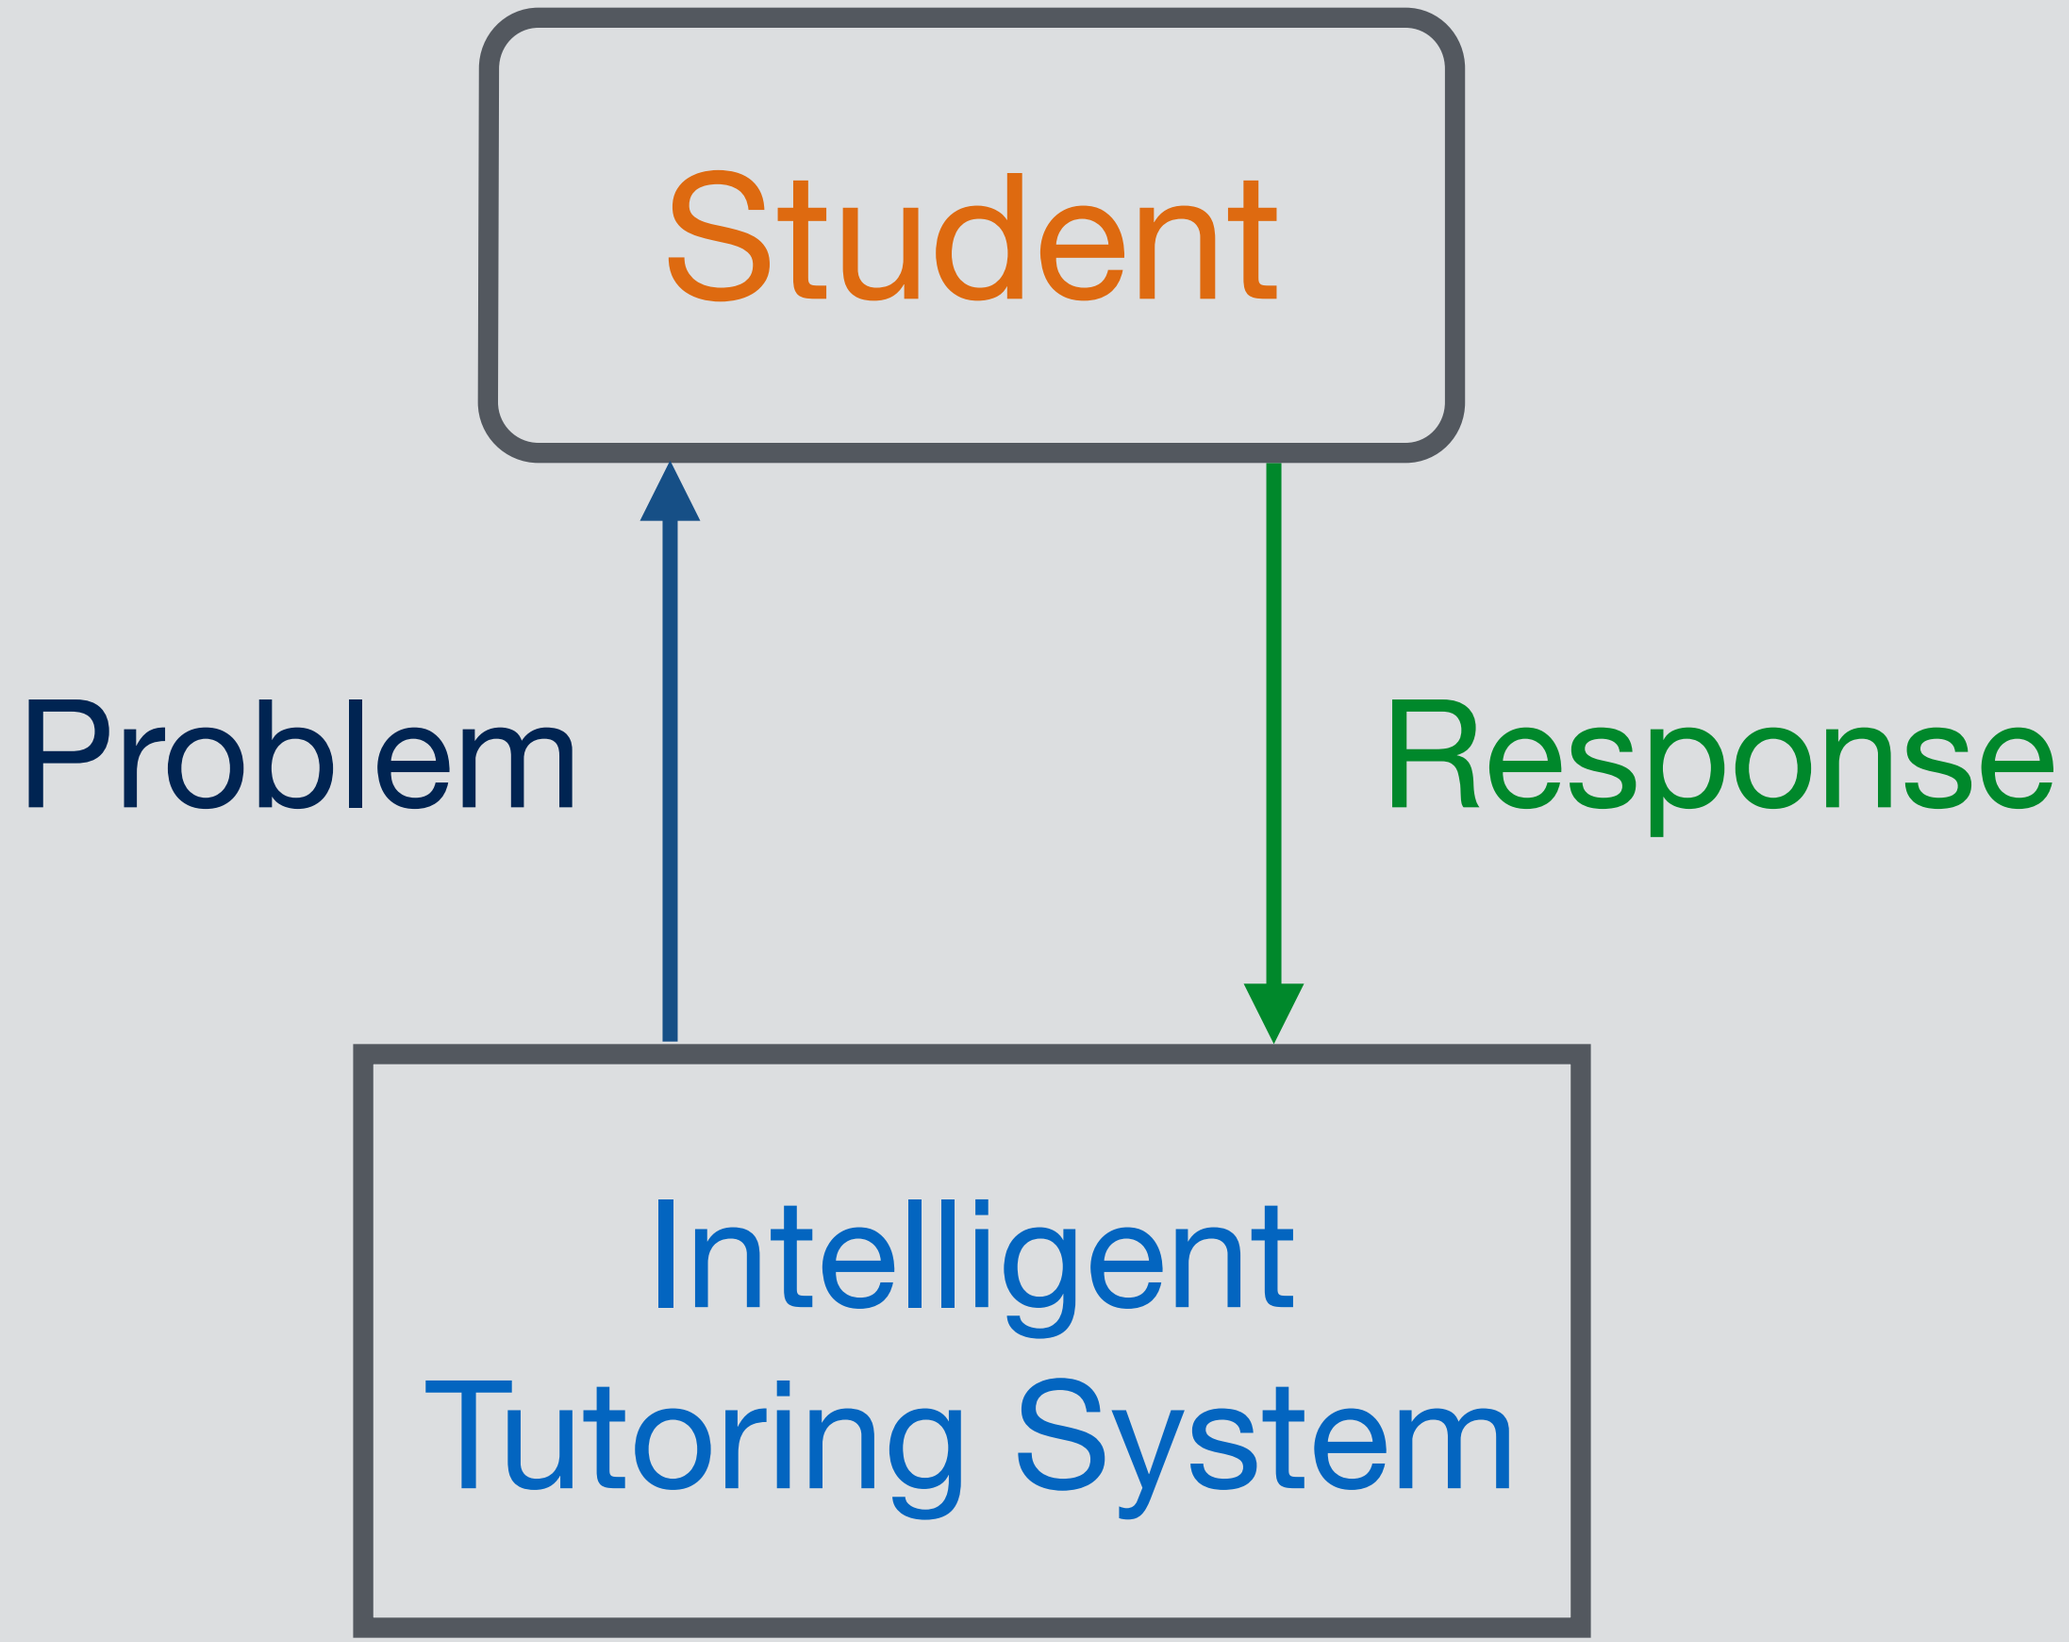
Student
Problem
Response
Intelligent
Tutoring System

Problem

Response

Instructional Policy

Student Model

Predictive Student Models

Student models are also frequently used for predicting student performance

Problem

Response

Instructional Policy

?

Student Model

Solution: Model Agnostic Instructional Policies

Instructional policies that can use any underlying predictive student model.

Case 1:
When-To-Stop Problem

When should the system stop providing problems for the given skill to the student?

Background:
Bayesian Knowledge Tracing (BKT)

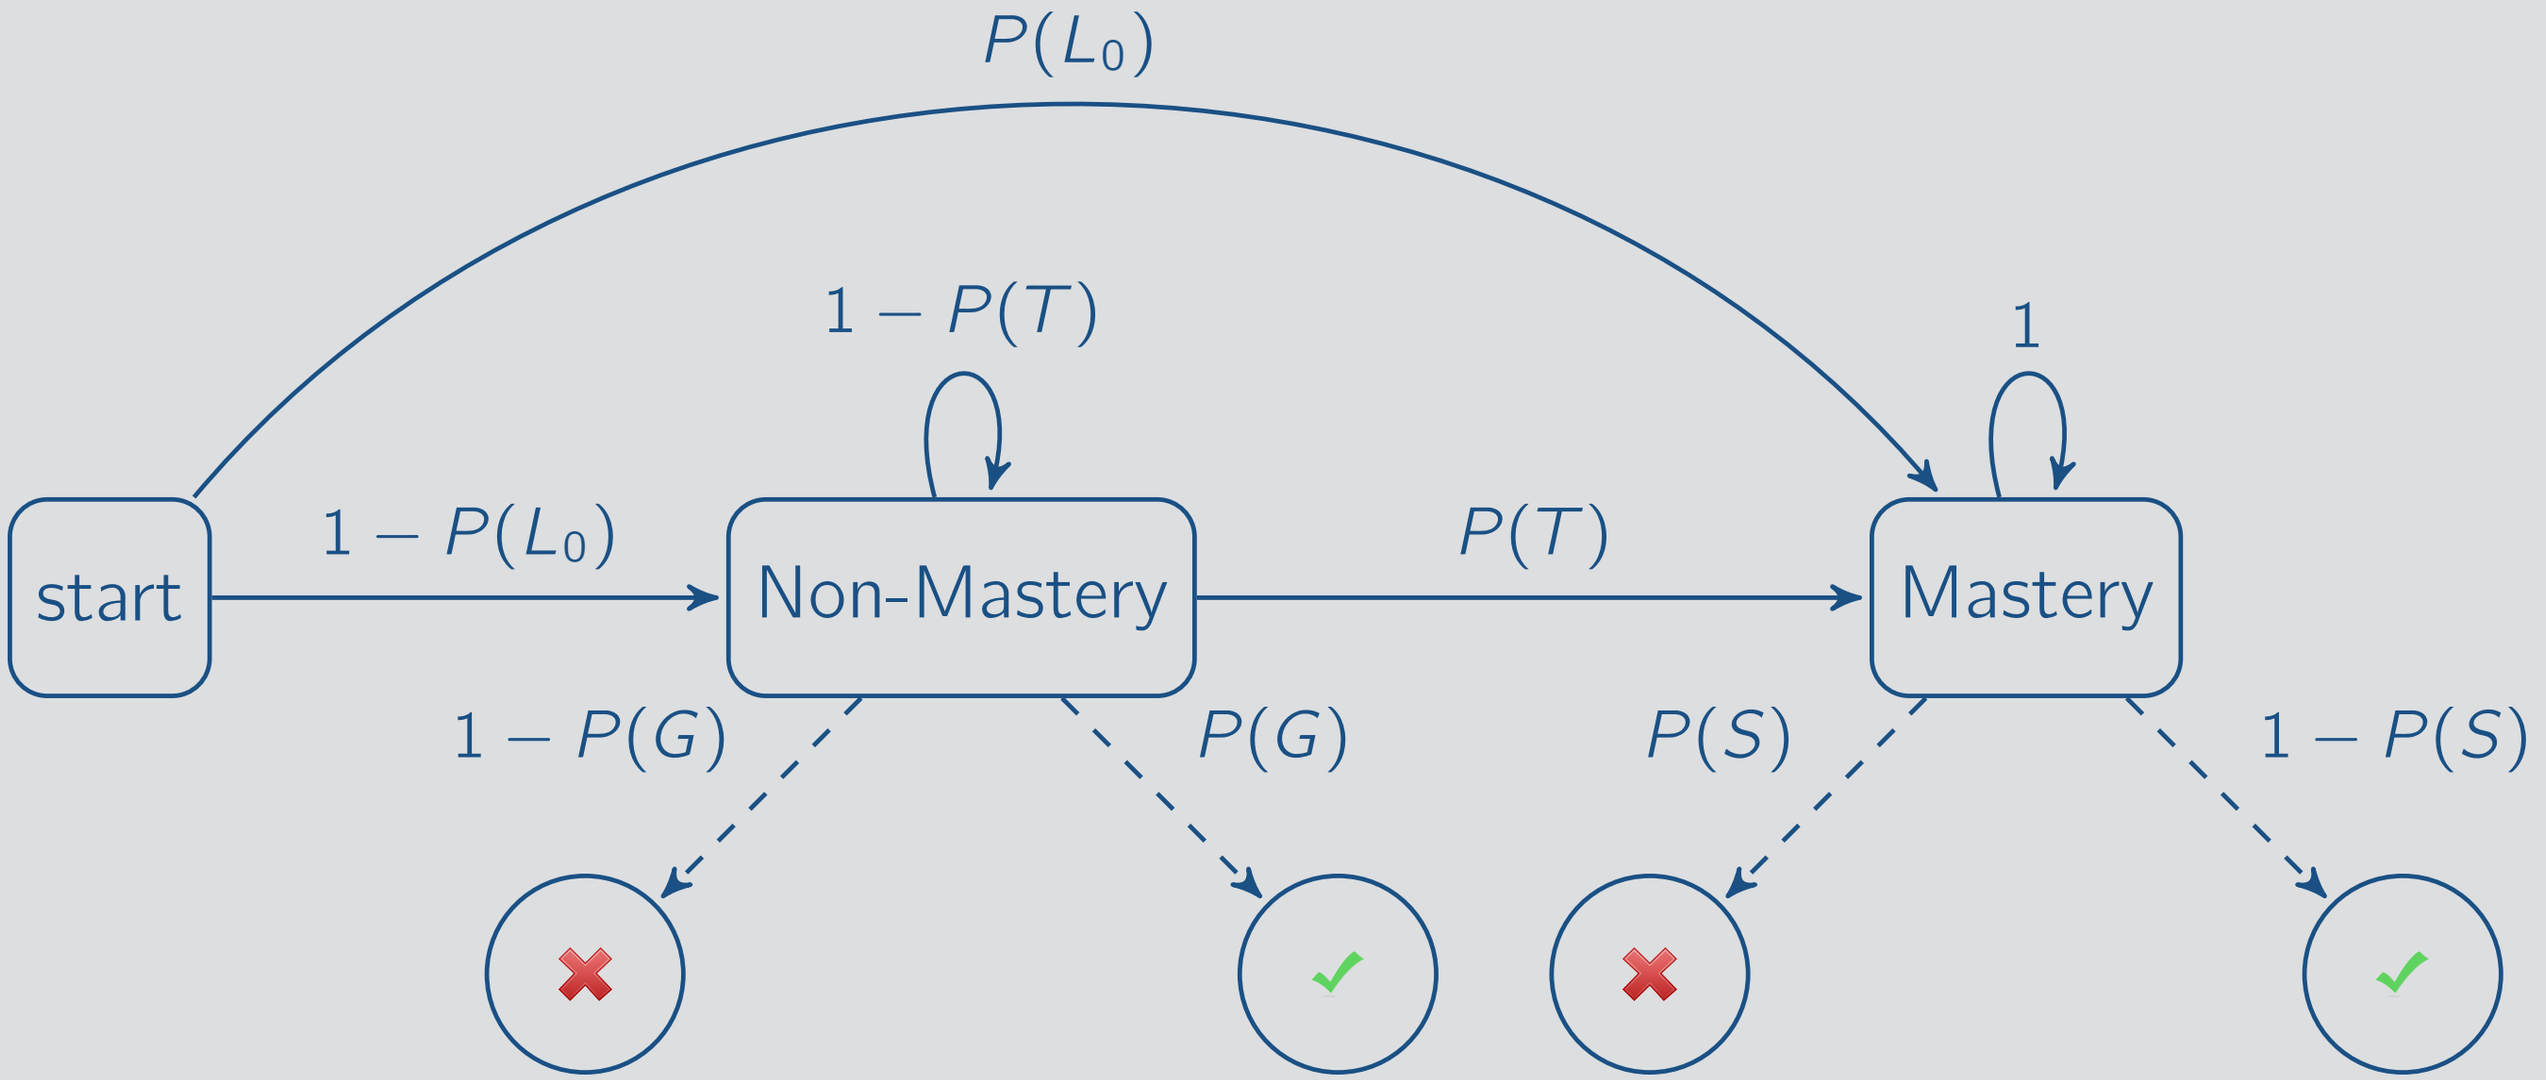

Corbett, A. T., & Anderson, J. R. (1995). Knowledge tracing: modeling the acquisition of procedural knowledge. User Modeling and User-Adapted Interaction, 4, 253–278

Background: Performance Factors Model (PFM)

Logistic Regression Model for predicting student performance.

Features:

- Student (i)
- Skill (k)
- # Correct responses for skill (s)
- # Incorrect responses for skill (f)

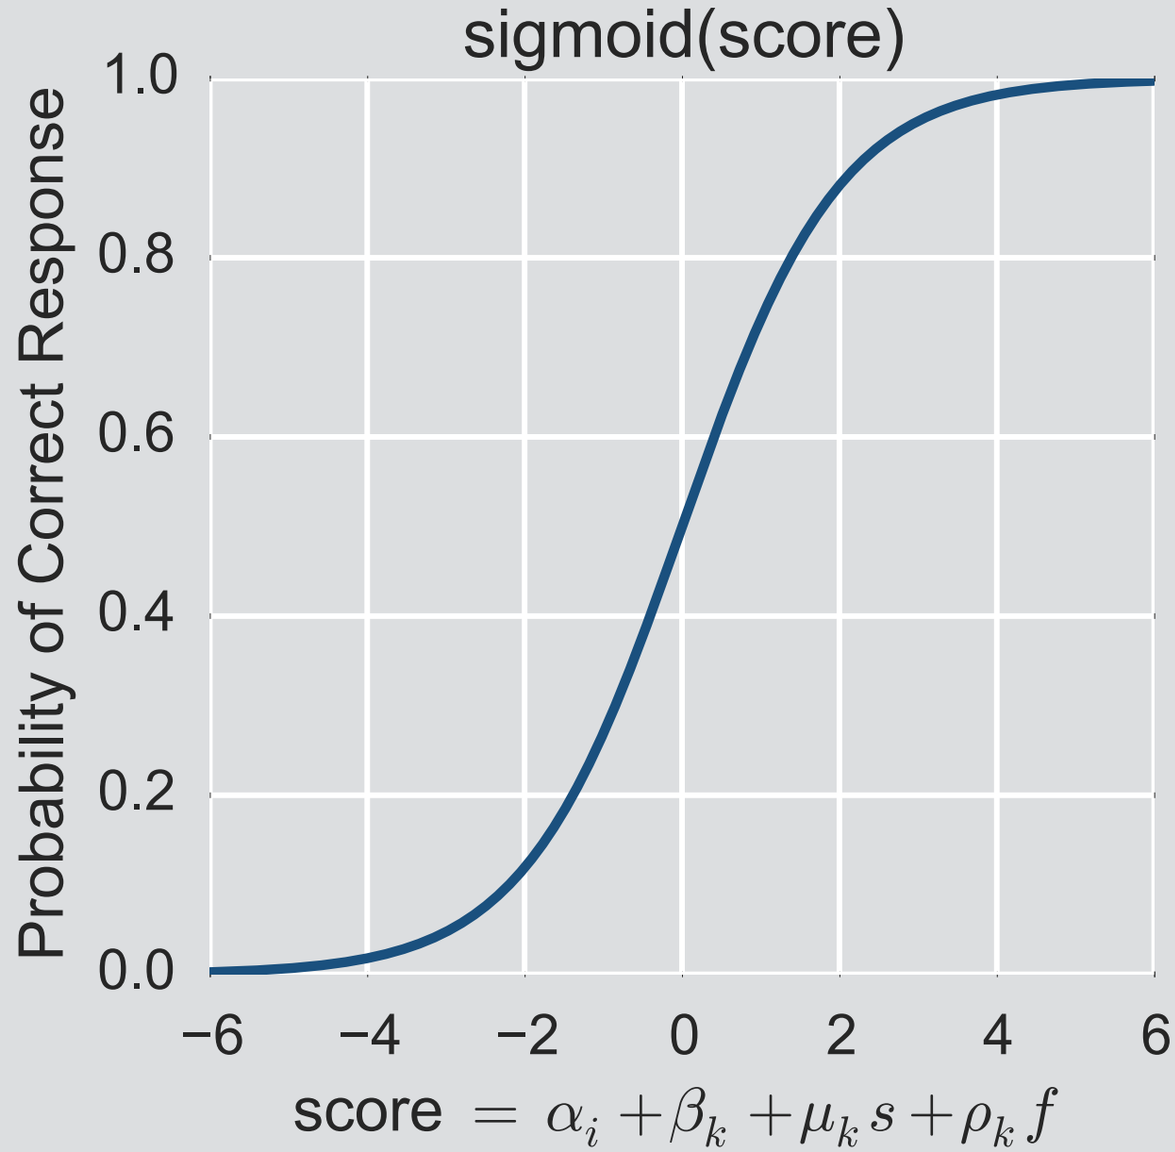

Pavlik Jr, Philip I., Hao Cen, and Kenneth R. Koedinger. "Performance Factors Analysis--A New Alternative to Knowledge Tracing." Online Submission (2009).

Prior Work:
Mastery Threshold Policy

Stop if we are confident that the student has mastered the skill.

$$P(M) > \Delta$$

Issues with the Mastery Threshold Policy

1. Requires student model with concept of mastery.
2. Will not stop if student cannot progress with given instruction (wheel-spinning).

Beck, Joseph E., and Yue Gong. "Wheel-spinning: Students who fail to master a skill." Artificial Intelligence in Education. Springer Berlin Heidelberg, 2013.

New Policy: Predictive Similarity Policy

Stop if we are confident that our prediction of the student's performance will not change much.

$$P\left(\left|P(C_{t+1}) - P(C_t)\right| < \epsilon\right) > \delta$$

Data: Algebra I

> 3000 students

505 skills

scenario

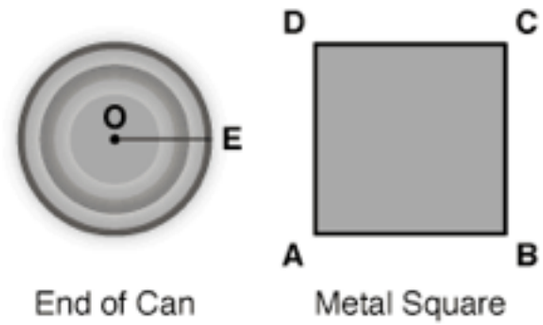

To make metal cans, the ends for the cans are stamped out of square pieces of metal. The part of the square that is left over is then recycled as scrap. The manufacturer needs to know the area of the scrap for each end. Then the total weight of the scrap can be figured out.

1. The can end has a radius of 4 inches. If an end is punched out of a square piece of metal measuring 8 inches on a side, find the square inches of the scrap.

2. The can end has a radius of 8 inches. If an end is punched out of a square piece of metal measuring 16 inches on a side, find the square inches of the scrap.

3. The can end has a radius of 12 inches. If an end is punched out of a square piece of metal measuring 24 inches per side, find the square inches of the scrap.

NOTE: To find the area of the scrap metal remaining, you might have to first find the area of the can end, and the area of the metal square

For this problem use an approximate value for pi. $\pi \approx 3.14$

Problem Making Cans

How much does the underlying student model affect the Predictive Similarity Policy?

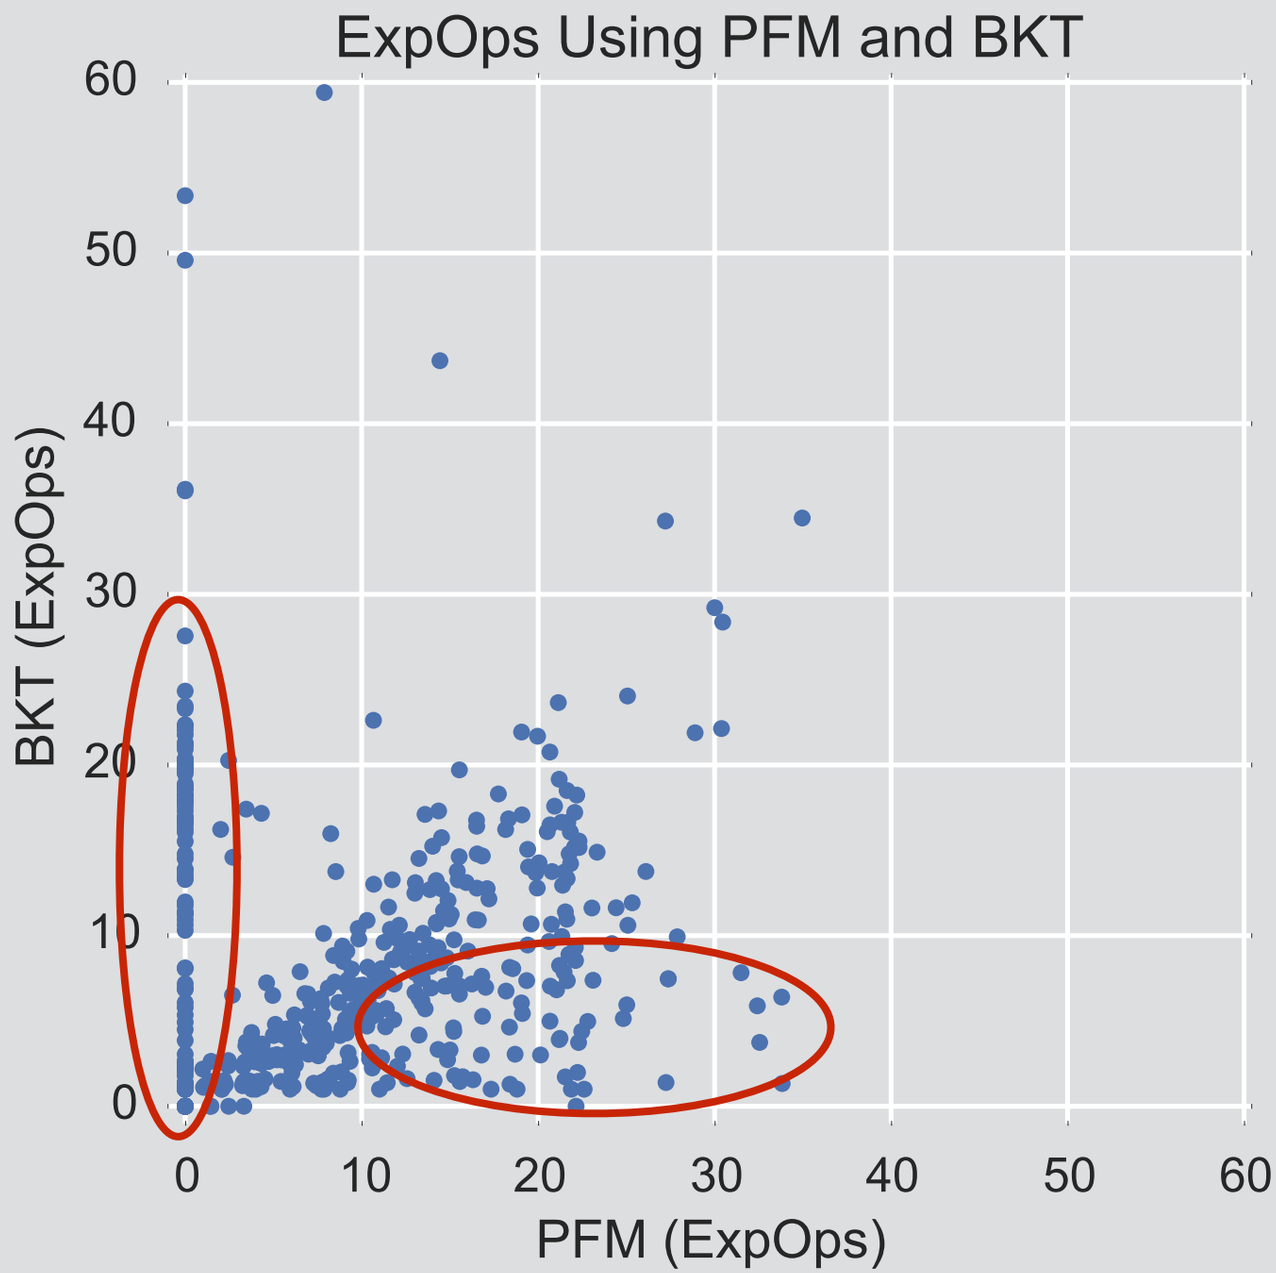

Case 2:
Skill-Choice Problem

Given T available questions, what skill should the next problem teach?

Independent Skills

French Grammar

être

je suis
tu es
il / elle est
nous sommes
vous êtes
ils / elles sont

Fraction Addition

$$\frac{2}{3} + \frac{1}{4} = \frac{11}{12}$$

Dependent Skills

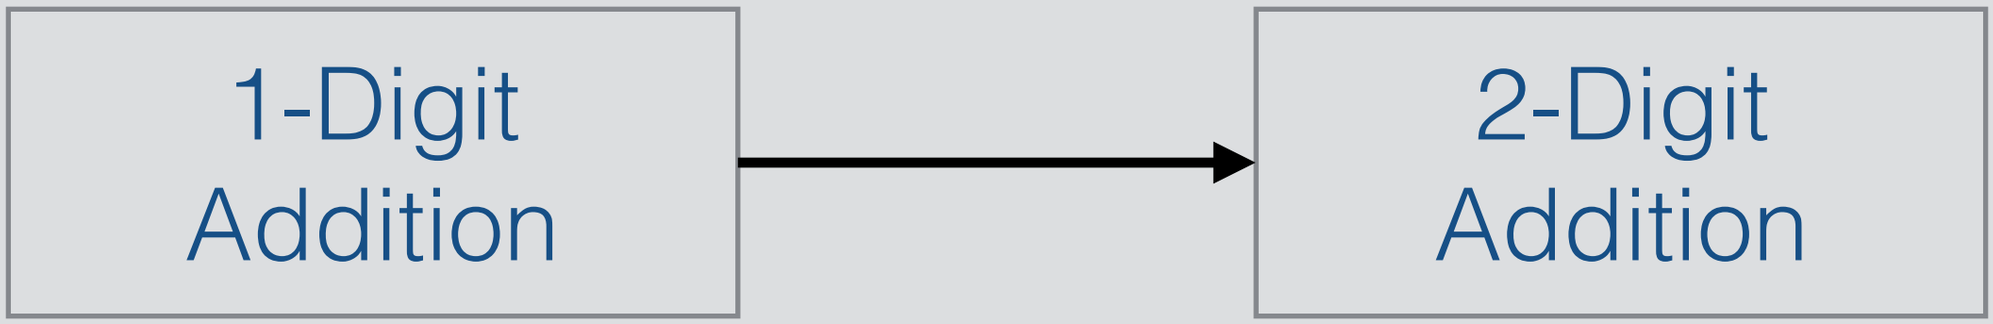

$$4 + 3 = 7$$

$$46 + 17 = 63$$

Problem

Pick the skill to maximize the final number mastered skills

What skill to teach?

Maximization

Expectation

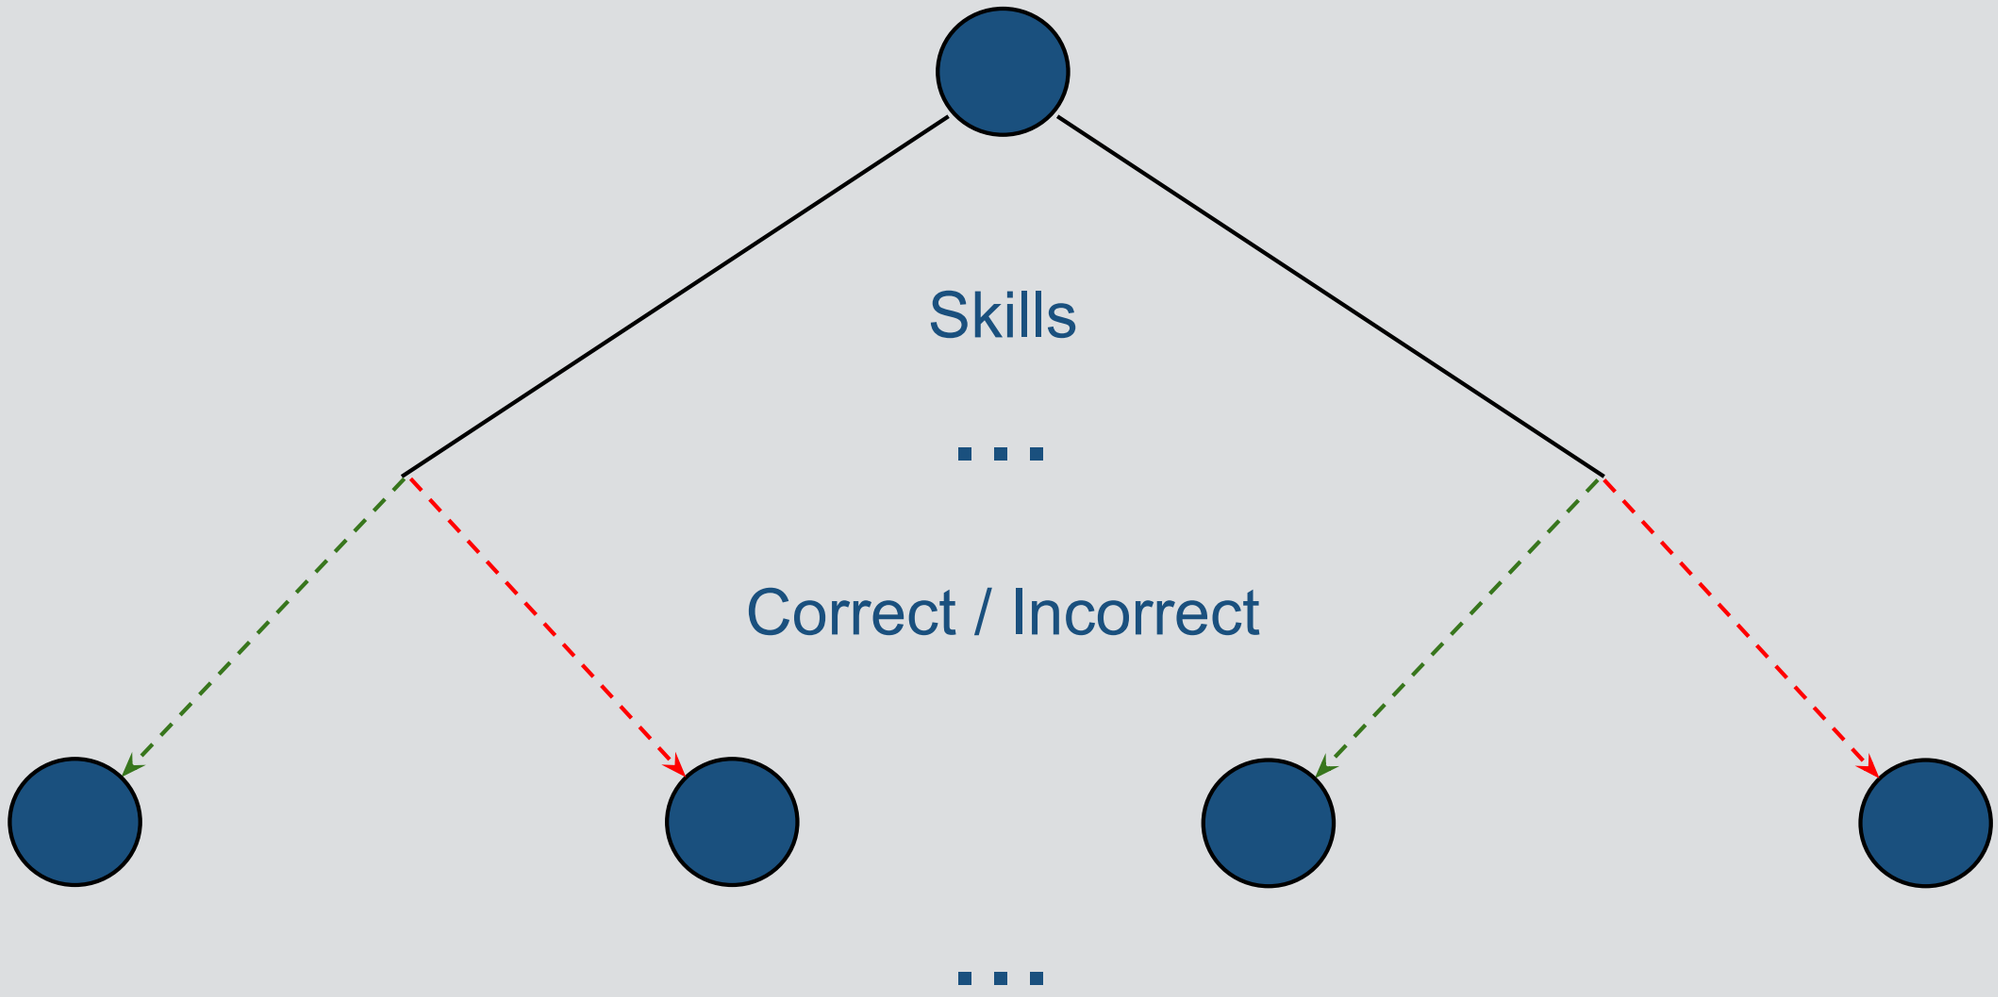

Full Horizon Policy

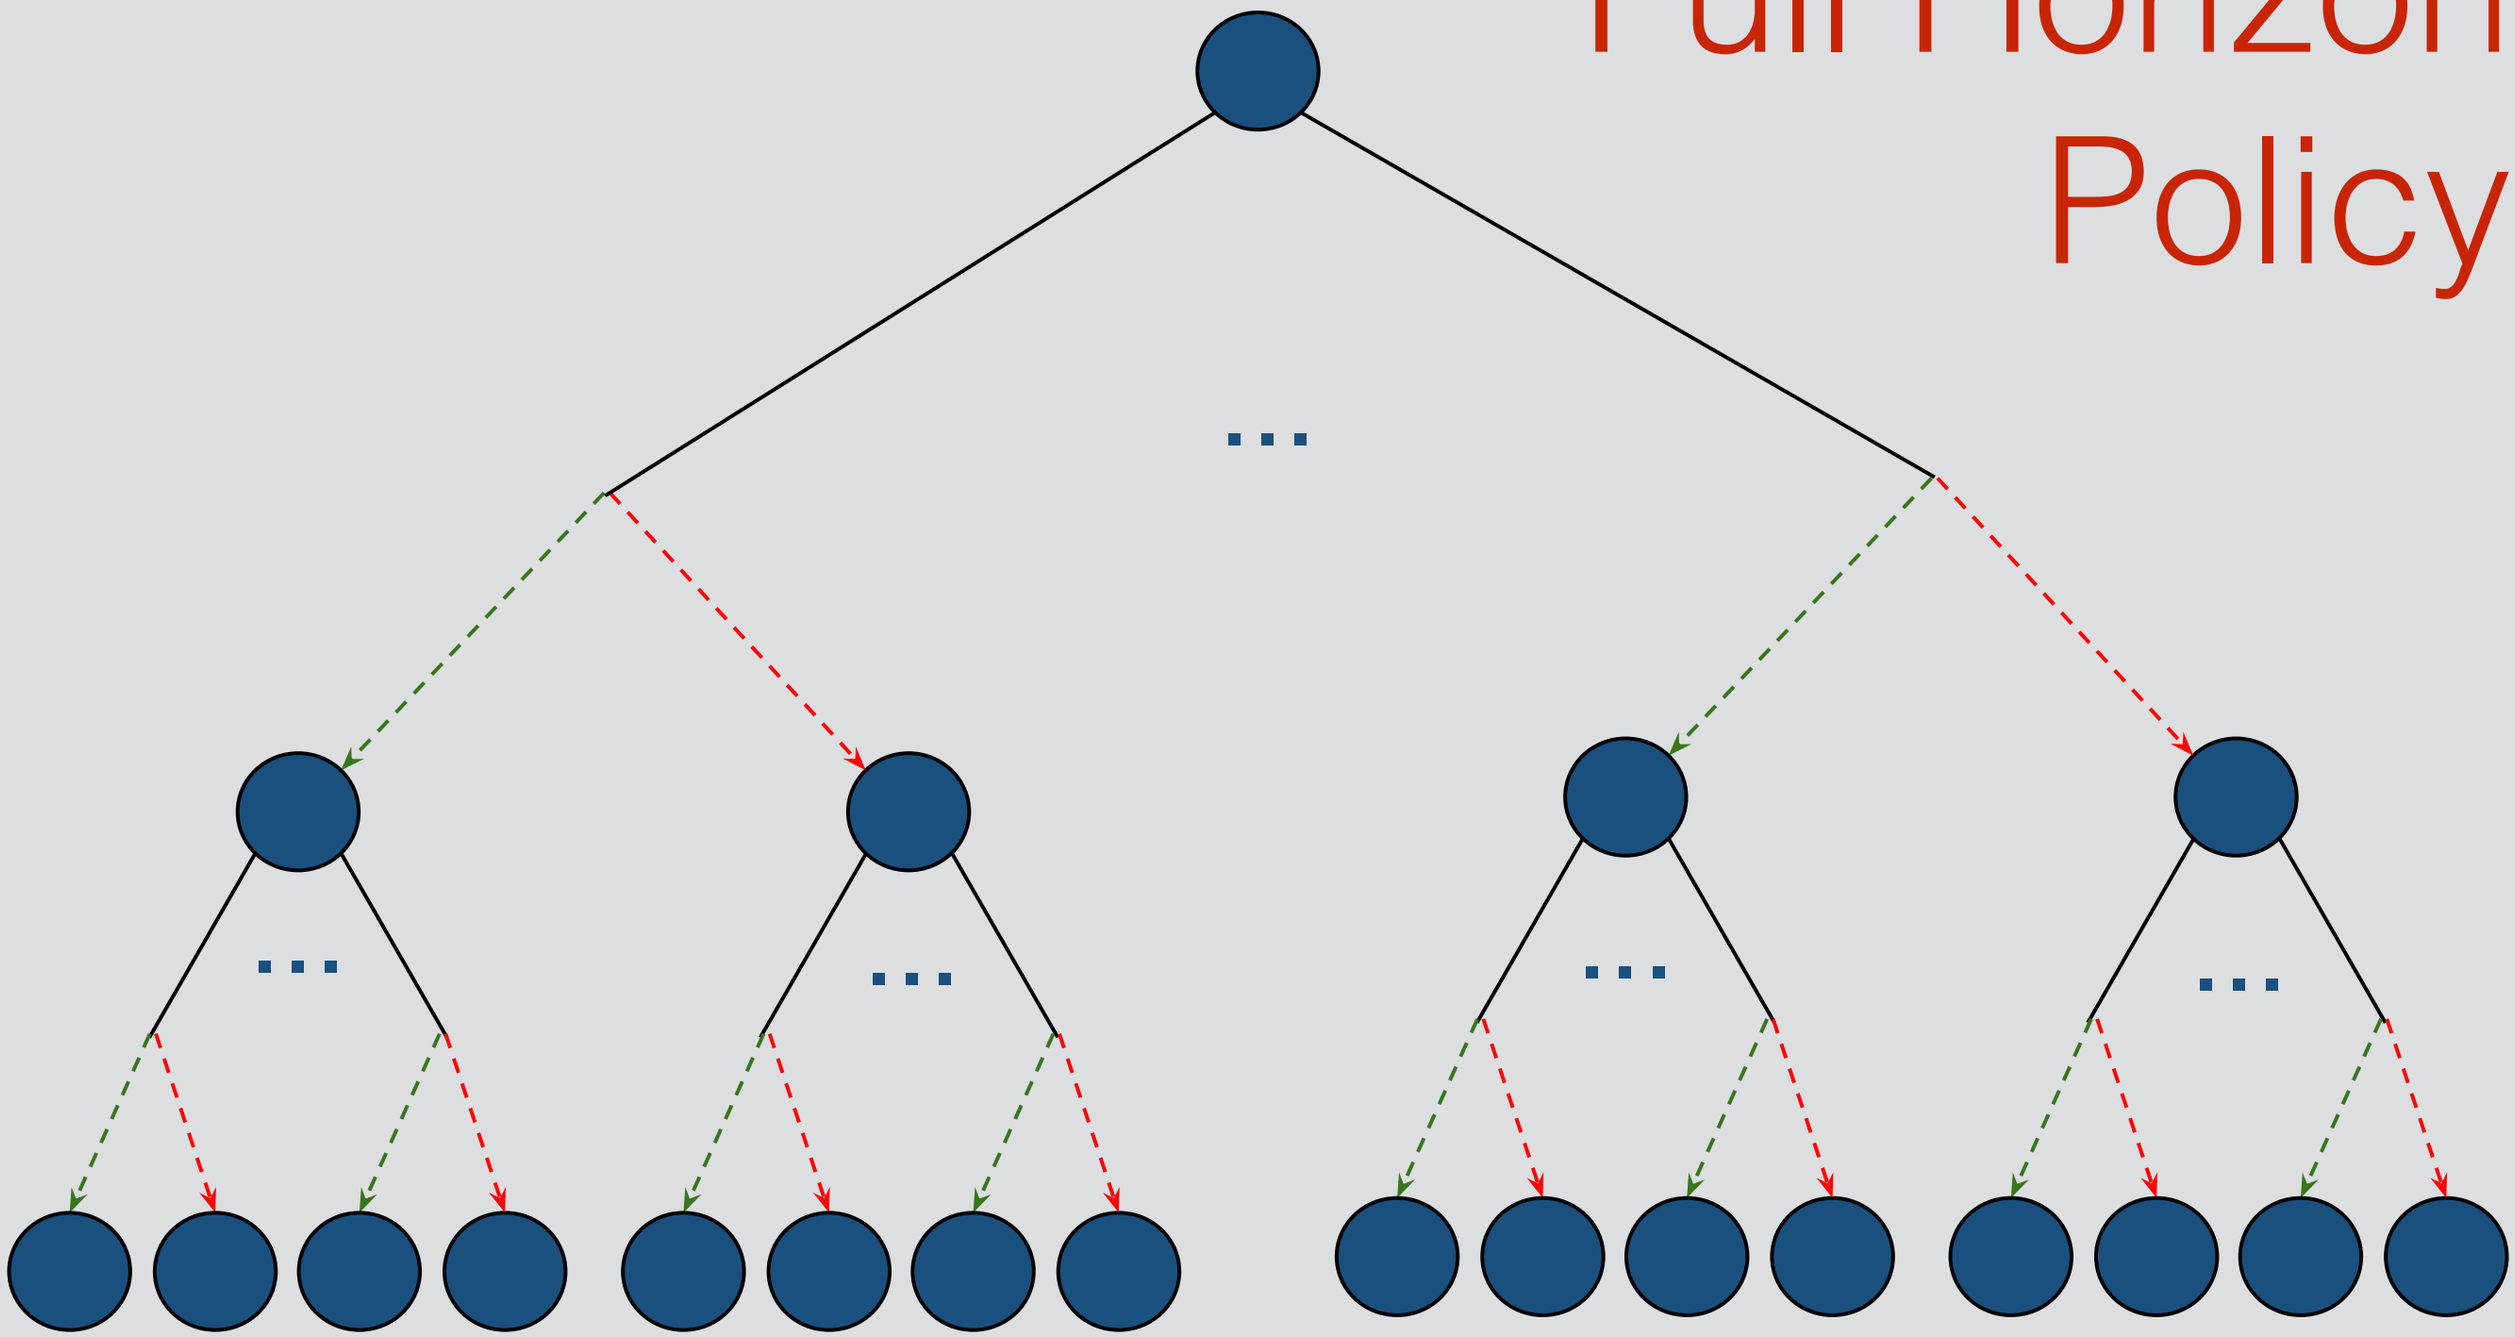

$$O\left((2|K|)^T\right)$$

Greedy Policy

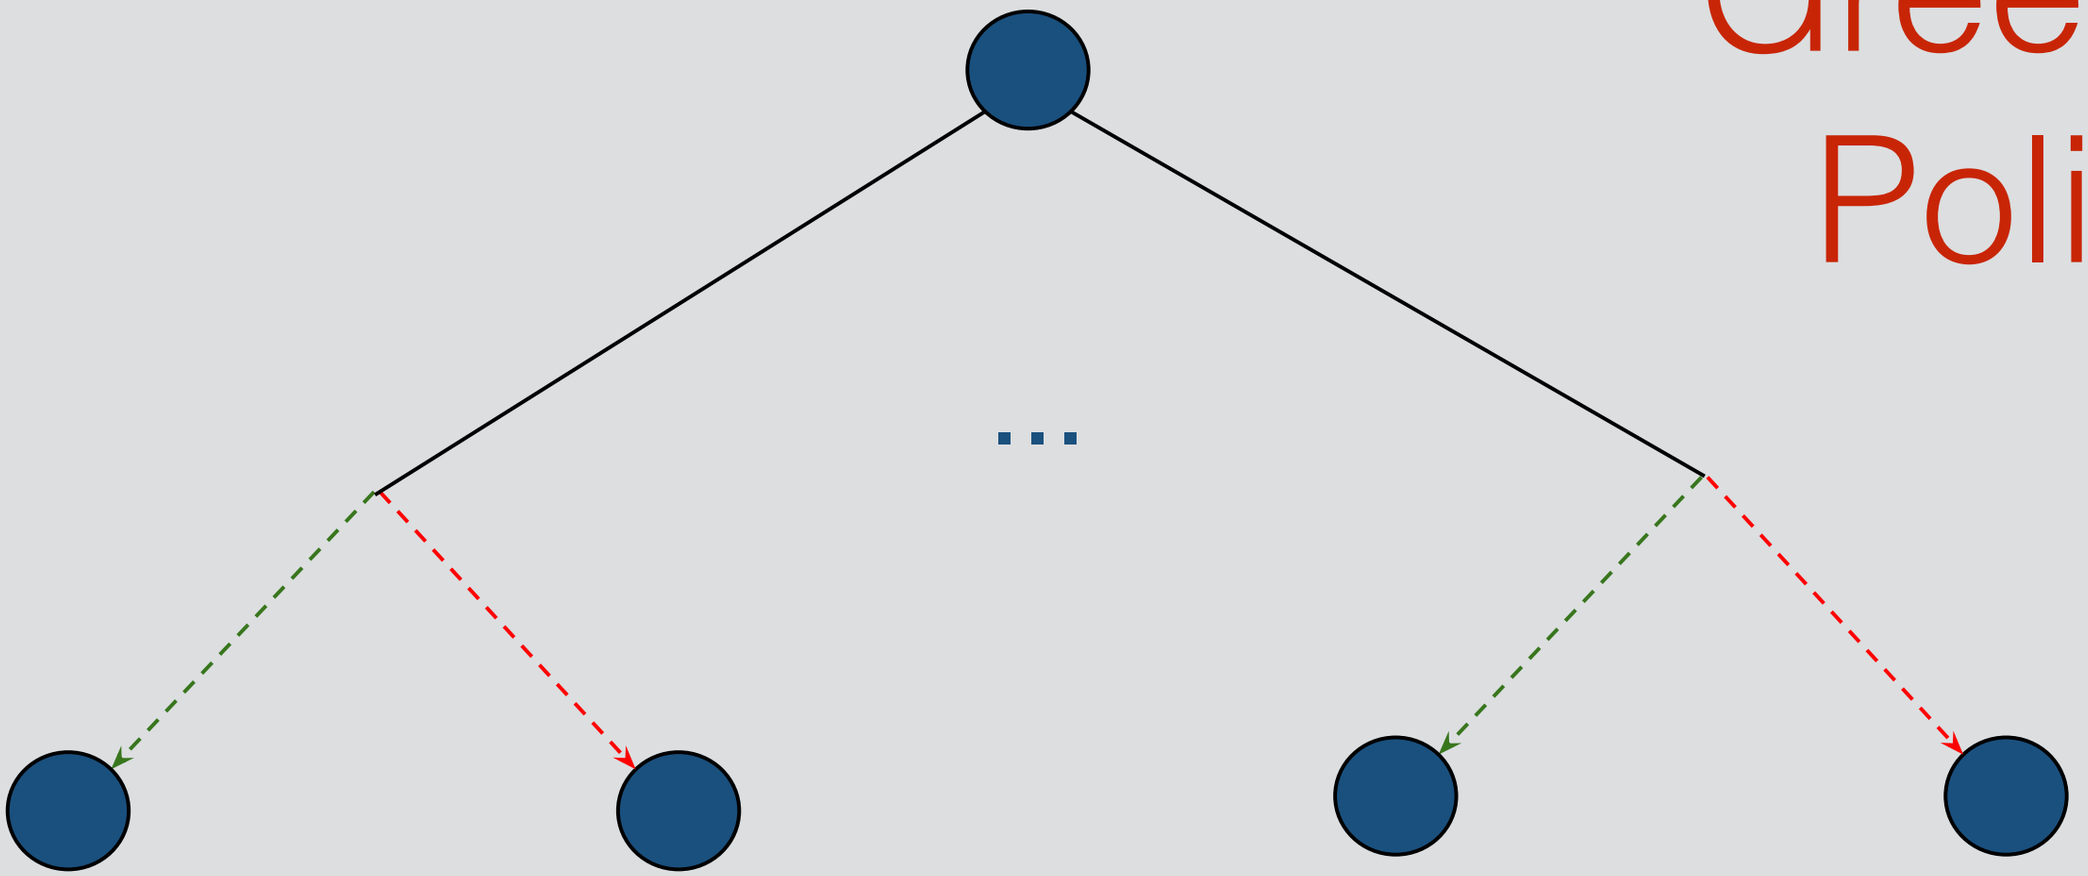

$O(2|K|))$

How does the Greedy policy compare to the Full Horizon Policy?

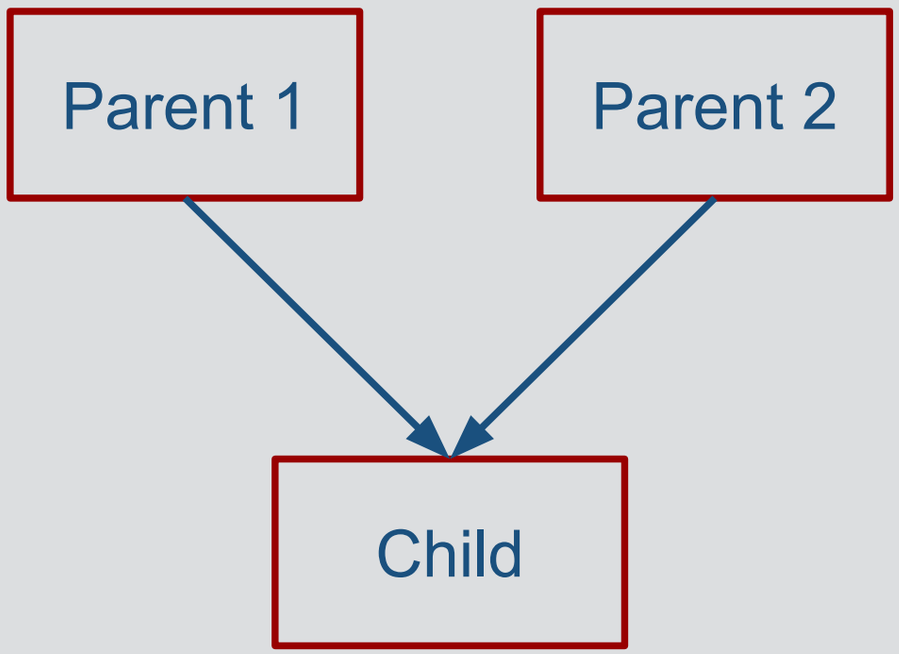

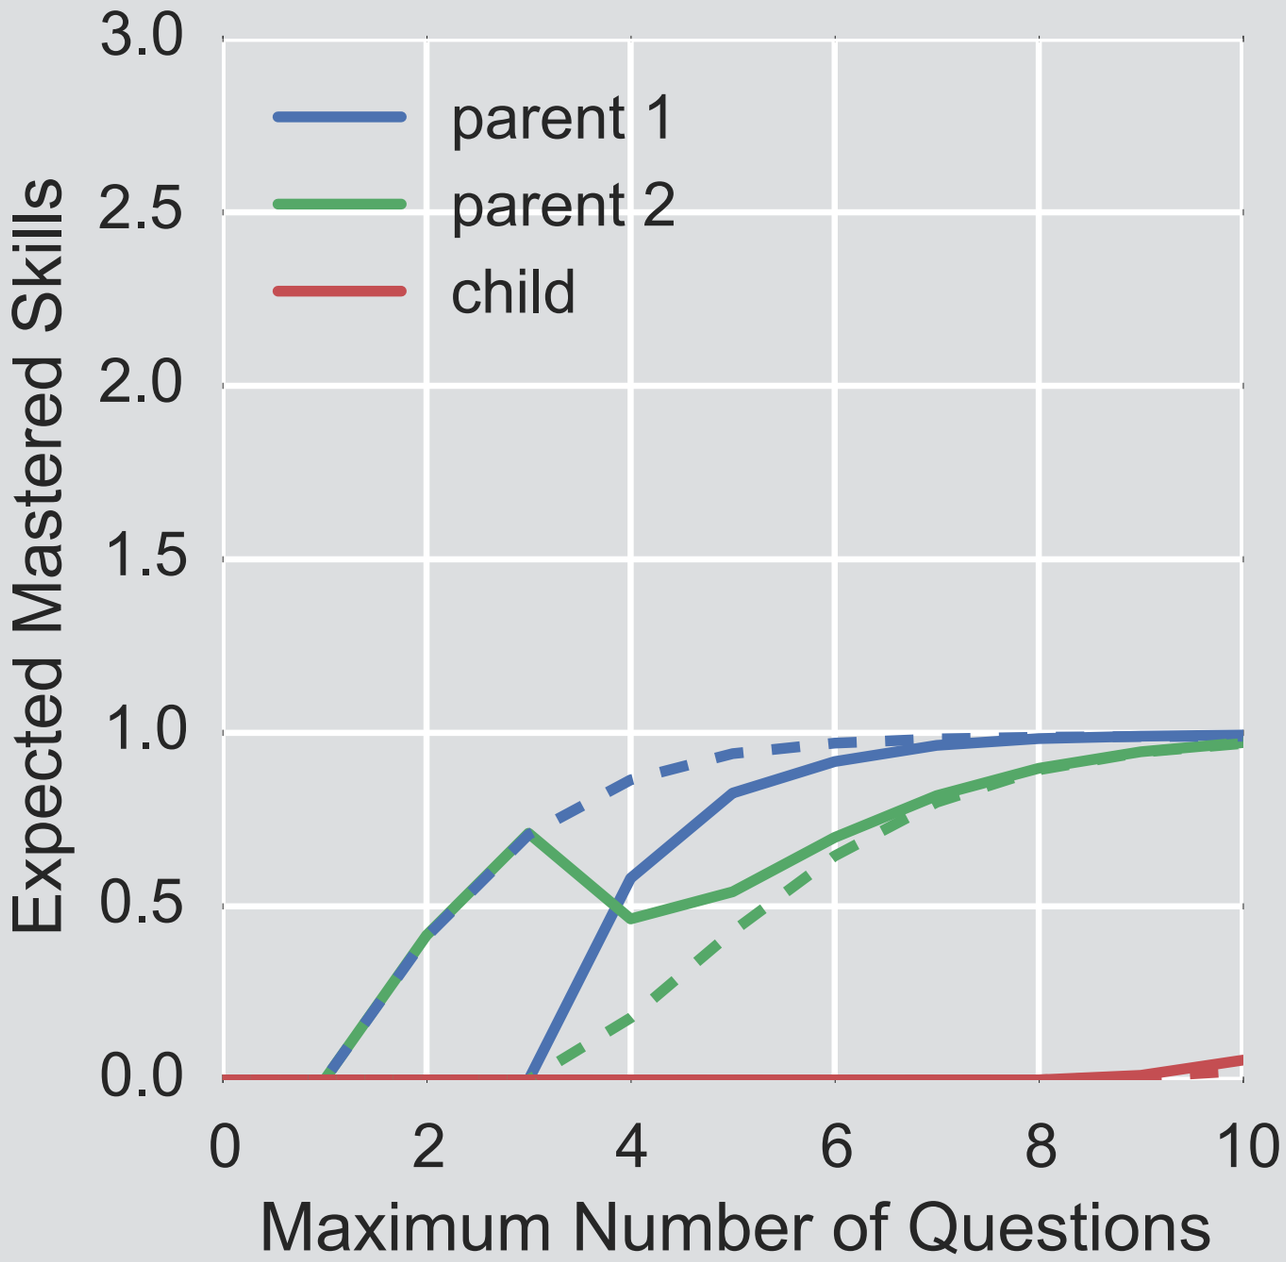

How does the Greedy policy compare to the Full Horizon Policy?

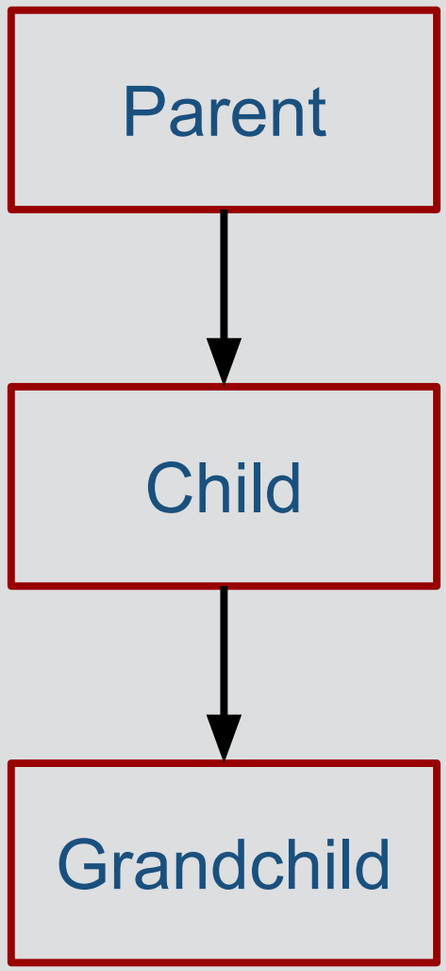

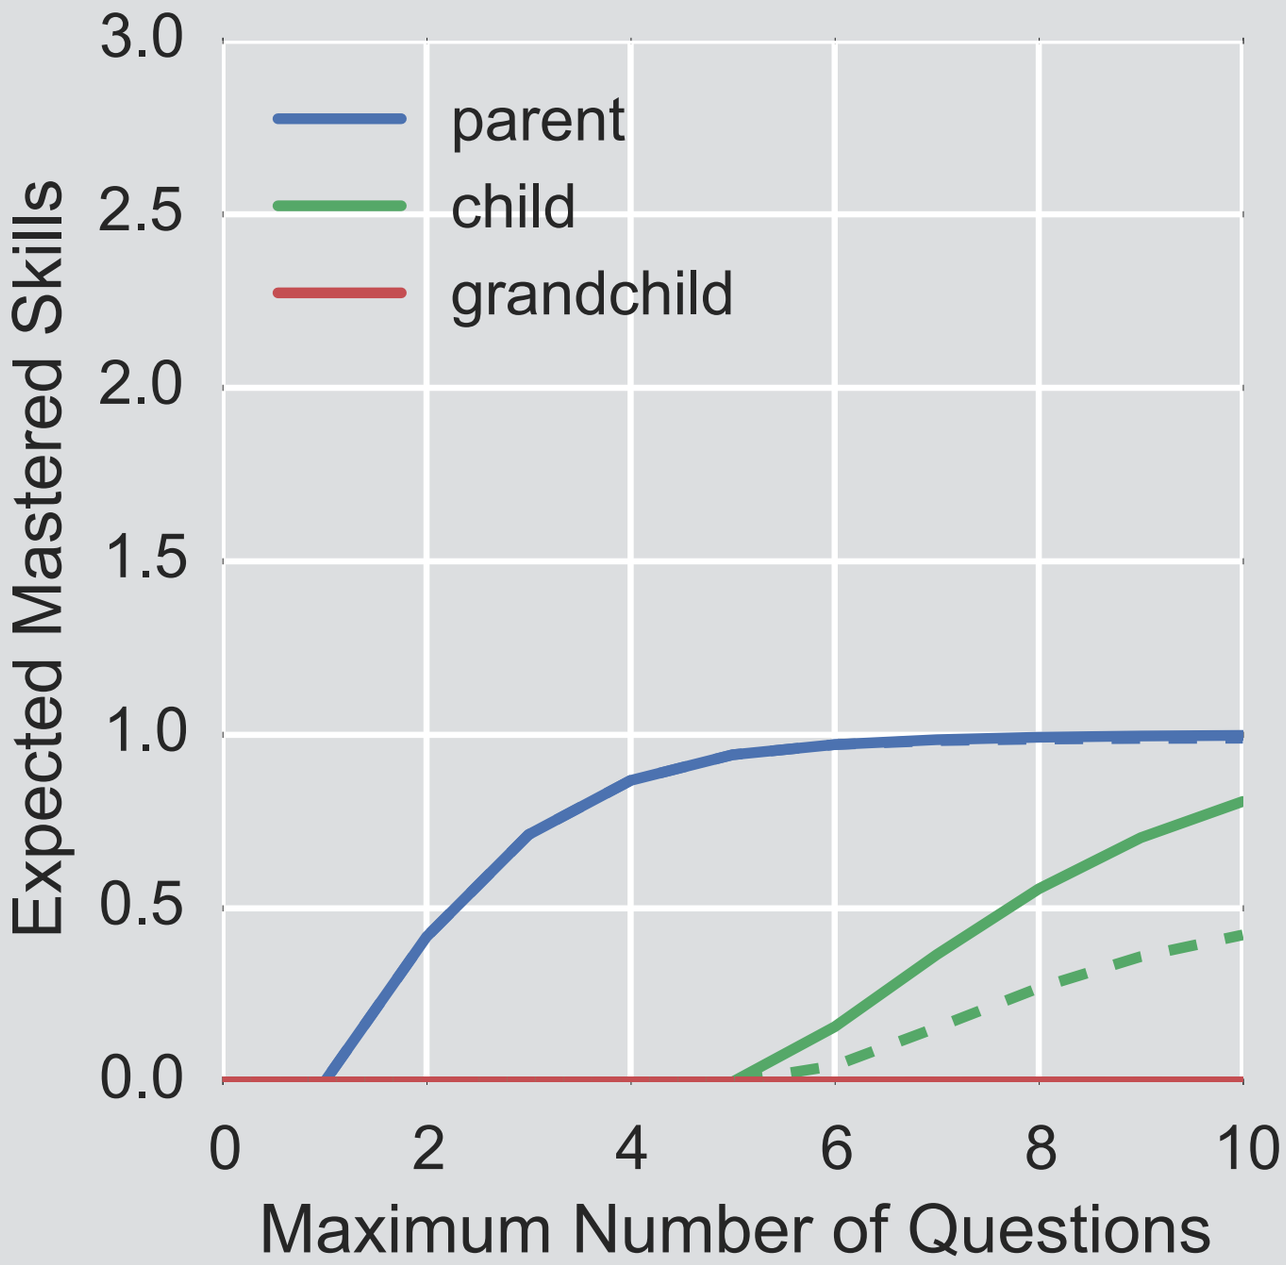

Looking ahead makes the student learn more

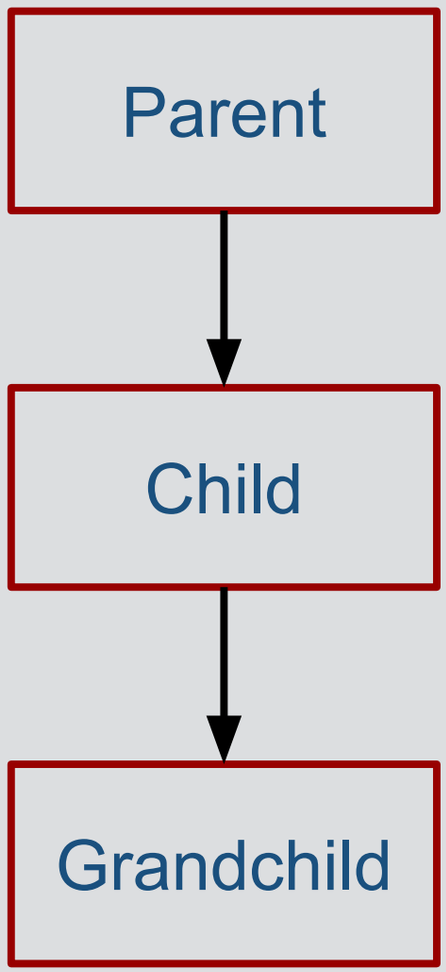

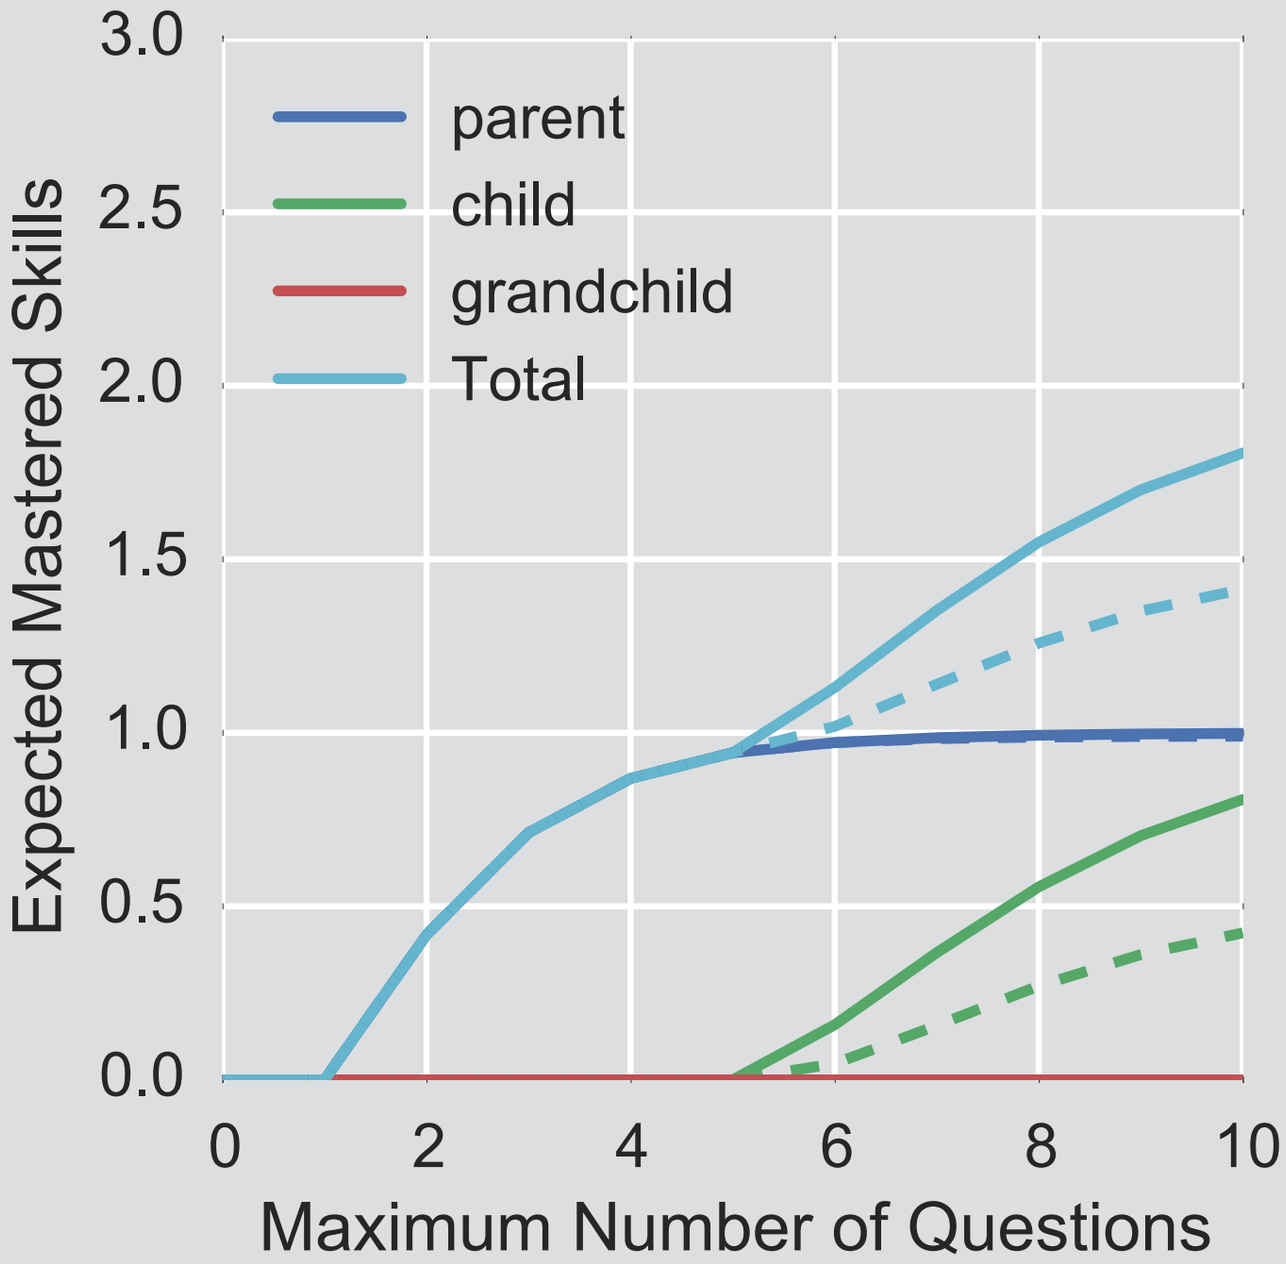

2 models

similar accuracy

different instruction

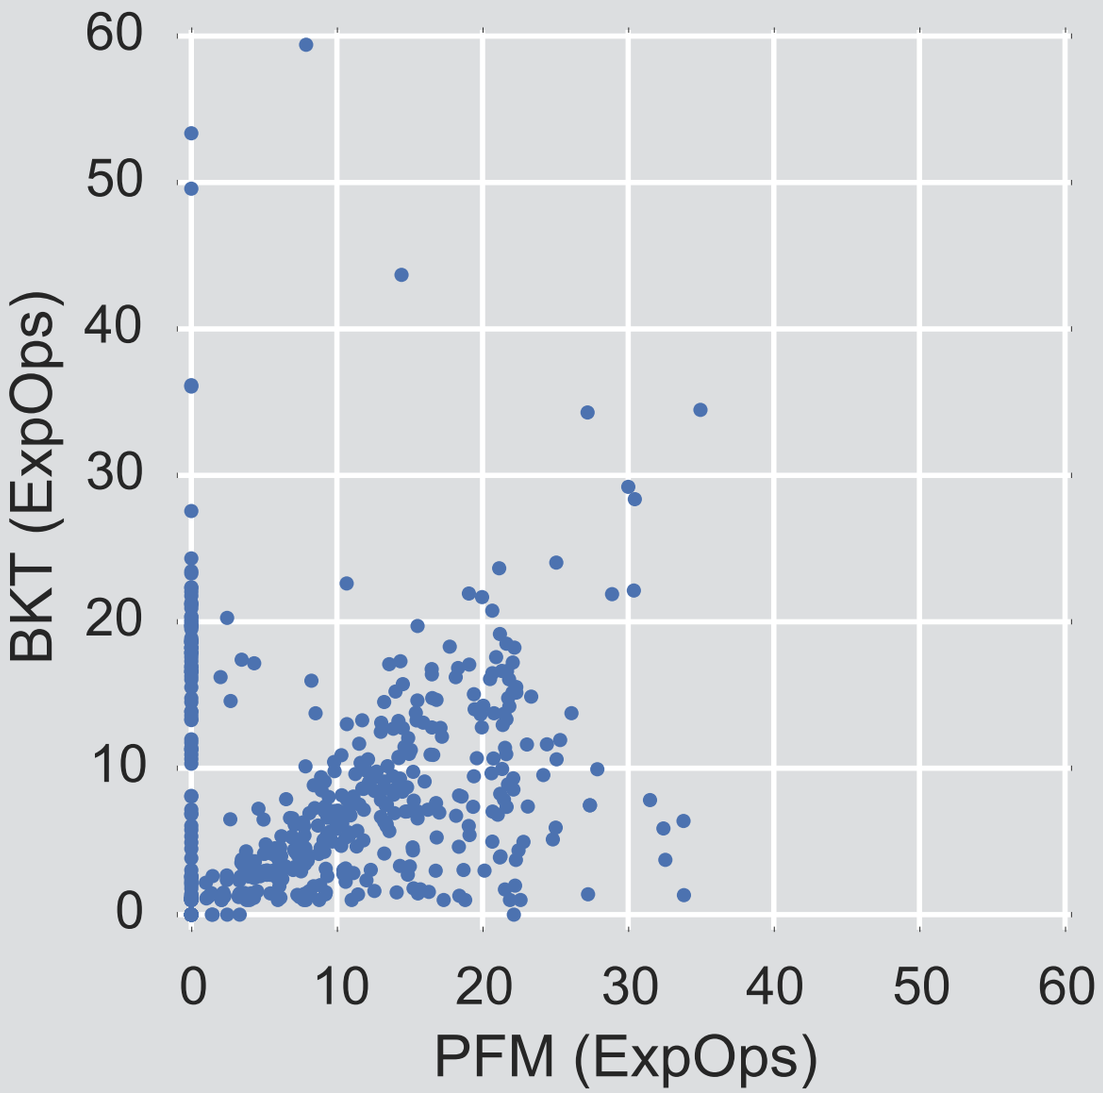

choosing multiple skills

looking ahead helps student learn more

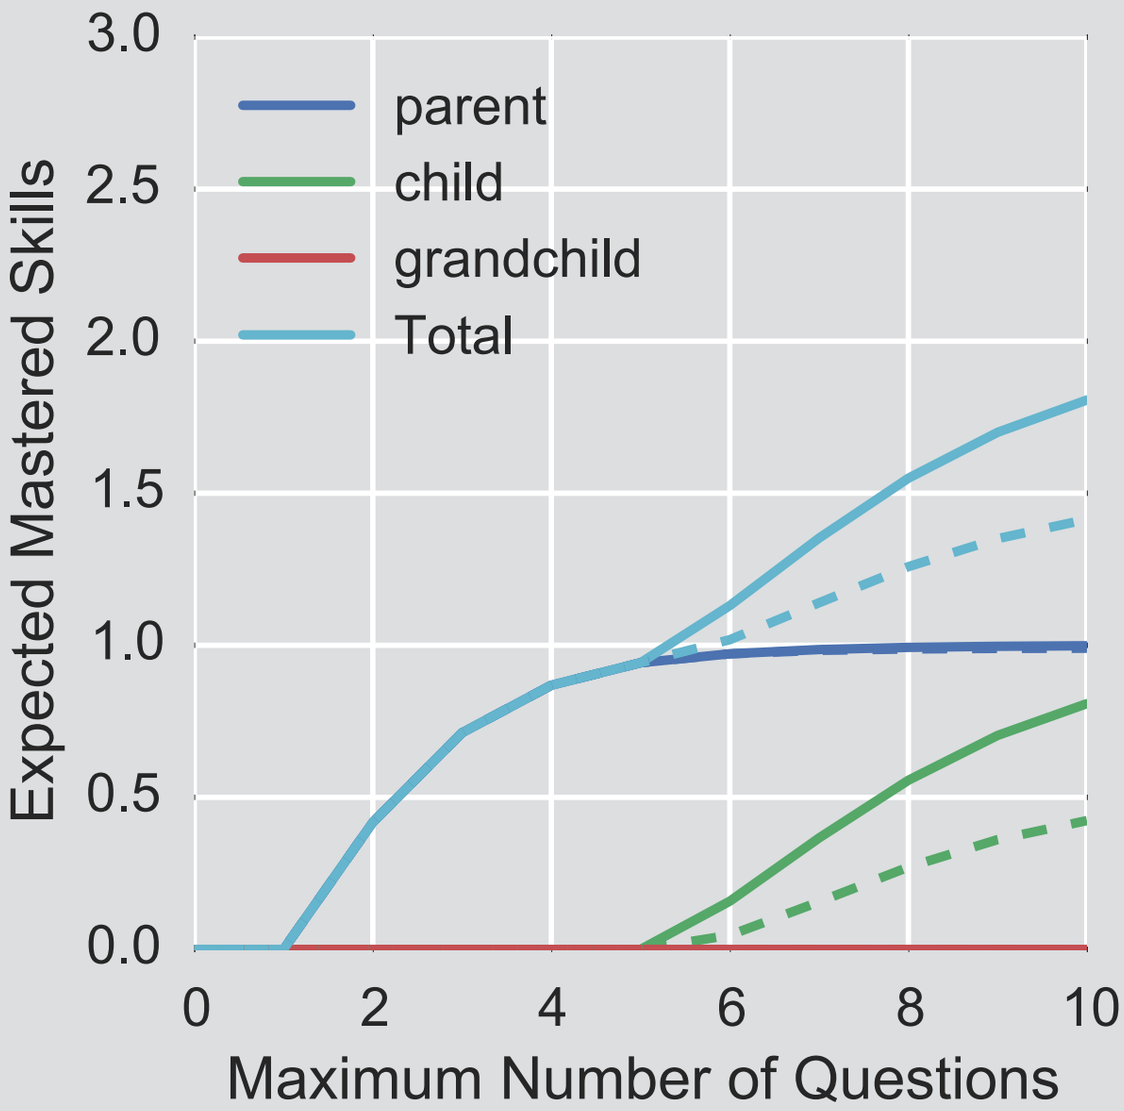

Future Work

- Multi-skill experiments with real data.
- Test more complicated skill hierarchies.
- Evaluating instructional policies.

Contributions this year

- Model agnostic when-to-stop policy.
- Windowed PFM for preventing extreme asymptotes.
- Hierarchical PFM for capturing skill hierarchies.
- Model agnostic skill-choice policies.

Questions?