On Planning, Cognitive Biases and Prophet Inequalities

Sigal Oren

Ben-Gurion University

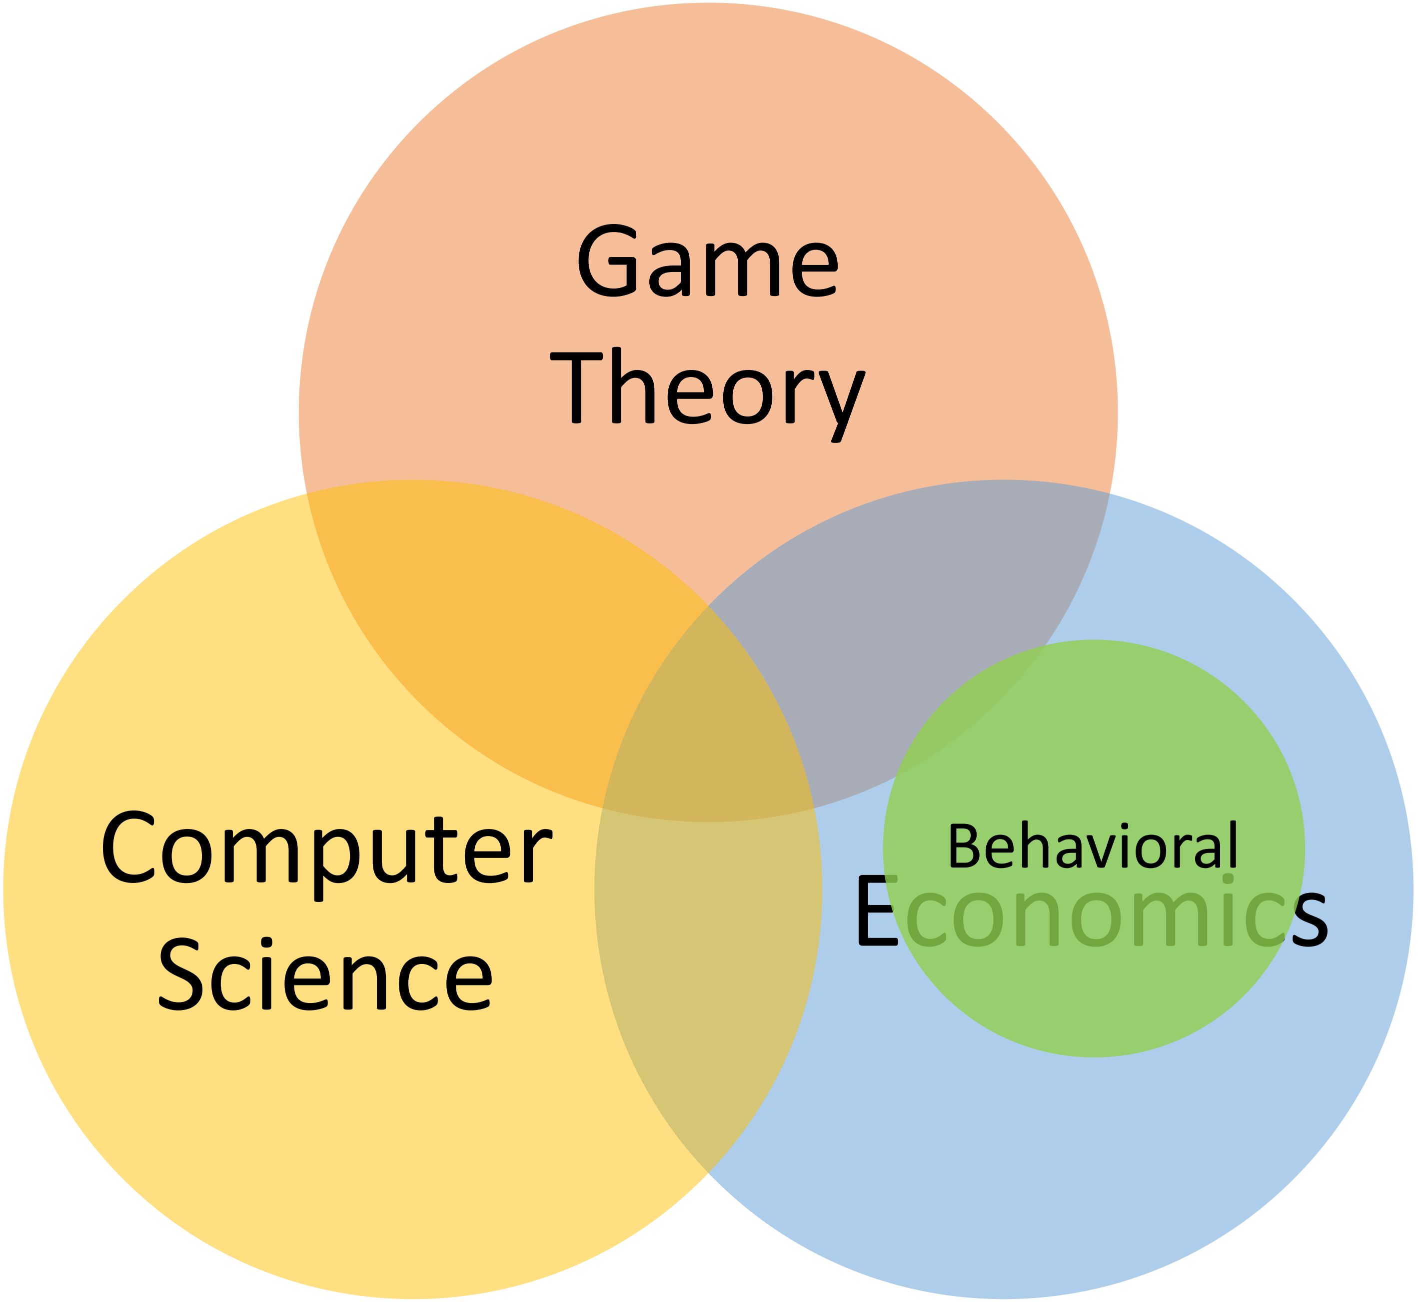
Game Theory
Computer Science
Behavioral
Economics


PUBLICIS
FLORA DANICA


PATISSERIE
BOULANGERIE
LA PETITE ALSACIE
MAISON MAEDER


La Fournée d'Augustine
Boulangerie Pâtisserie à l'Ancienne
bar brasserie


BOULANGERIE
FABRICATION TRADITIONNELLE
Le Point


by The Baker Upstairs
for Somewhat Simple

Biases in Decision Making

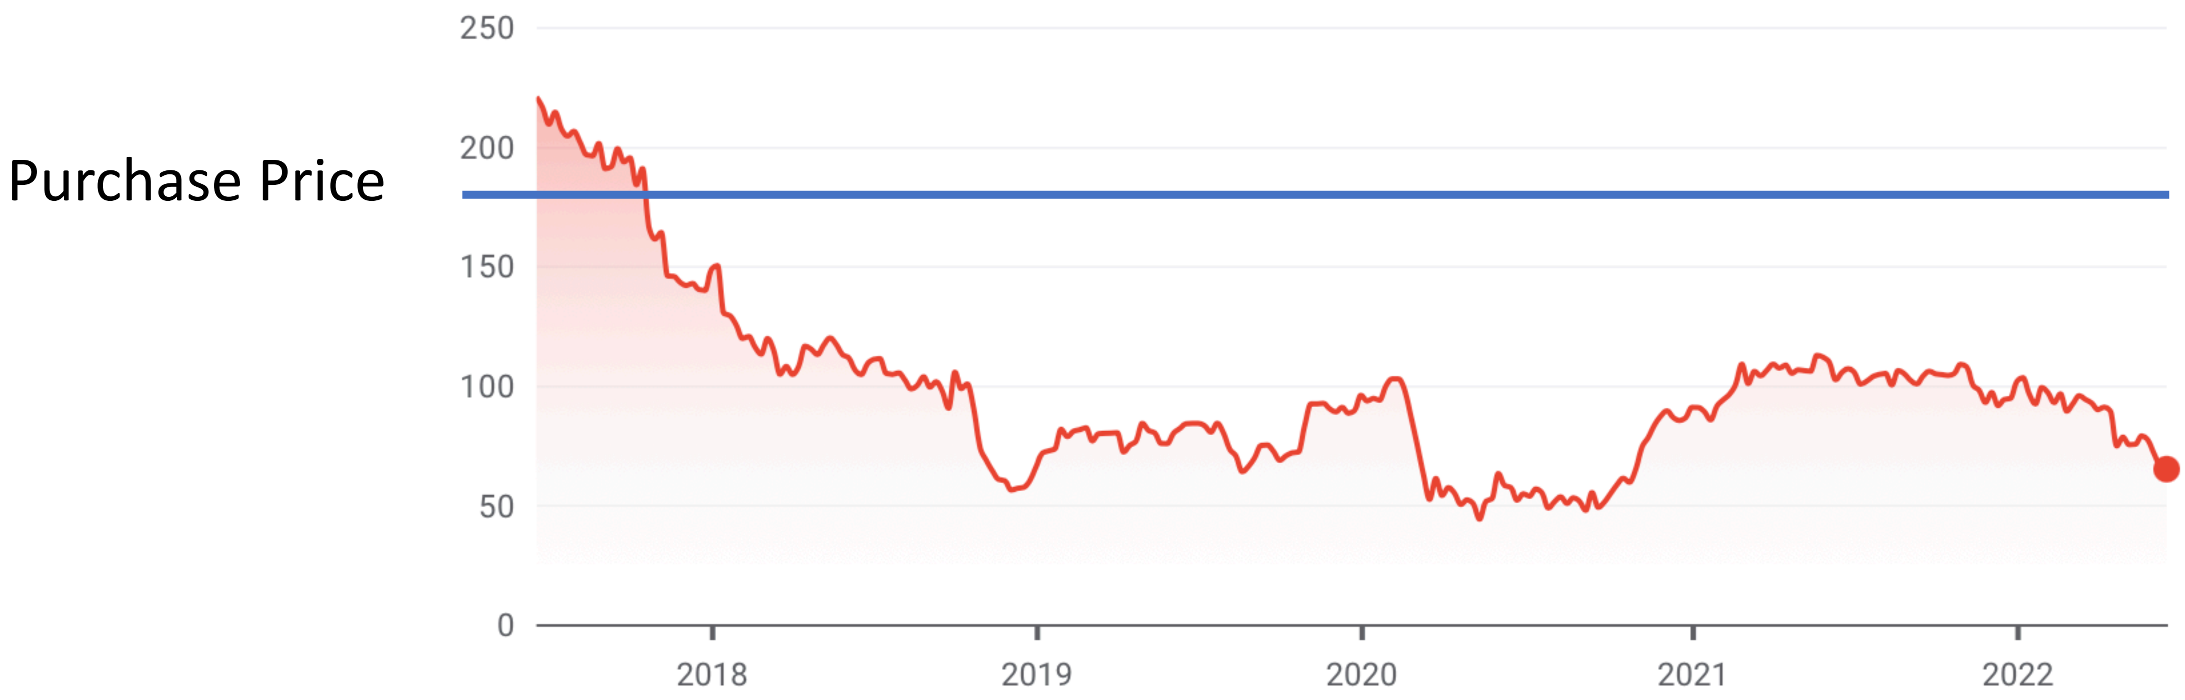

Investors tend to avoid selling stocks below the purchase price (Odean'98)

This behavior is related to two behavioral biases (Kahneman and Tversky'79):

Reference point – *gains and losses are computed with respect to some reference point* - the purchase price.

Loss aversion – *losses are weighed higher than gains.*

Biases in Decision Making

Is the reference point necessarily fixed?


Reference Point: Highest quality of boulangerie observed so far.

Subjects performing a sequential search task in an experiment stopped searching too early (Schunk and Winter'09)

Consistent with a *changing reference* point: minimal price observed.

Optimal Stopping Problems

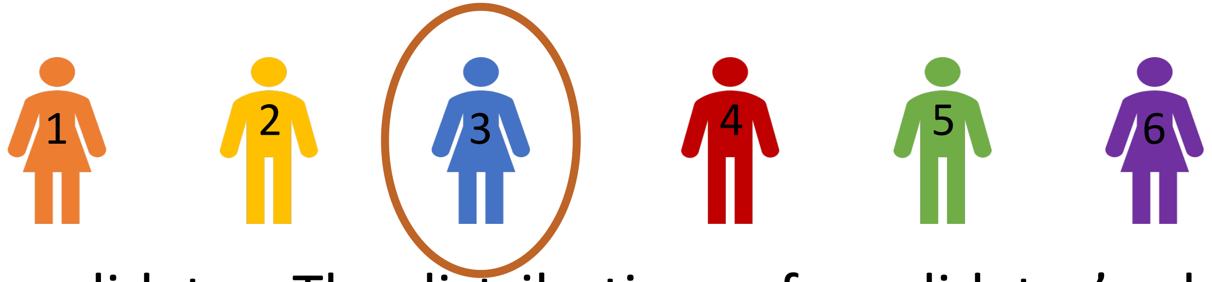

$v_t \sim F_t$

- **Setup:** n candidates. The distributions of candidates' values are known.
- **Process:** The agent interviews the candidates by order to reveal their value. After interviewing the agent decides whether to hire or not.
- **Objective:** maximize the quality of the candidate that is hired.

Prophet inequality: There exists a stopping rule selecting a candidate with expected value at least ½ of the best candidate in hindsight (Krengel and Sucheston'78).

Behavioral Adaptation (Kleinberg, Kleinberg and Oren'21):

Consider a reference-dependent agent that has loss aversion.

- Let $\lambda \geq 0$ be a parameter denoting the extent of the loss aversion.
- Let v denote the value of the best candidate so far (=the reference point).
- If the agent hires a candidate of value $u < v$ its utility is $u - \lambda(v - u)$.
- The biased agent aims to maximize its expected utility.

An Example

Behavioral Adaptation: *Consider a reference-dependent agent that has loss aversion.*

- Let $\lambda \geq 0$ be a parameter denoting the extent of the loss aversion.
- Let v denote the value of the best candidate so far (=the reference point).
- If the agent hires a candidate of value $u < v$ its utility is $u - \lambda(v - u)$
- The biased agent aims to maximize its expected utility.

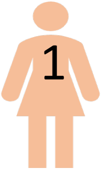

$$v_1 = \frac{1}{2}$$

$$v_2 = \begin{cases} 4 & w.p\ 1/4 \\ 0 & w.p\ 3/4 \end{cases}$$

Expected value of a prophet: $\frac{3}{4} \cdot \frac{1}{2} + \frac{1}{4} \cdot 4 = 1.375$

Expected value of an unbiased agent : 1

Expected value of a 2-biased agent : 0.5

Performance ratio: $\frac{\text{prophet}}{2\text{–biased}}$ =2.75, $\frac{\text{unbiased}}{2\text{–biased}}$= 2

2-biased agent:

$$\underbrace{1 - \frac{3}{4} \cdot 2\left(\frac{1}{2} - 0\right) = \frac{1}{4}}_{\substack{\text{Expected utility} \\ \text{for taking } v_2}} \leq \underbrace{\frac{1}{2}}_{\substack{\text{Utility} \\ \text{for taking } v_1}}$$

Results: "Prophet Inequality"

V^*- value of the best candidate in hindsight.

V_λ - value of the candidate chosen by an optimal λ-biased stopping rule.

Theorem: $\dfrac{\text{prophet}}{\lambda\text{–biased}} = \dfrac{E[V^*]}{E[V_\lambda]} \leq \lambda + 2$ and this is tight.

Proof Sketch (Samuel-Cahn'84) – show that the following stopping rule achieves the desired bound: hire the first candidate with value greater than θ where $\Pr\left(V^* > \theta\right) = \dfrac{\lambda+1}{\lambda+2}$.

Detour: Prophet Inequality for optimal agents - $\frac{E[V^*]}{E[V_0]} \leq 2$

Accept the first candidate with value above a threshold θ such that $Pr[V^* > \theta] = 1/2$

Donate this stopping rule by τ. We will show $\frac{E[V^*]}{E[V(\tau)]} \leq 2$

$$E[V(\tau)] = \int_0^\infty Pr(V(\tau) > y)dy = \underbrace{\boxed{\int_0^\theta Pr(V(\tau) > y)dy}}_{\geq \frac{1}{2}\theta} + \int_\theta^\infty Pr(V(\tau) > y)dy$$

$$Pr(V(\tau) > y) = \sum_{t=1}^n Pr(V_t > y, V_{-t} < \theta) = \underbrace{\boxed{\sum_{t=1}^n Pr(V_t > y)}}_{\geq Pr(V^* > y)} \cdot \underbrace{Pr(V_{-t} < \theta)}_{\geq 1/2}$$

$$\int_\theta^\infty Pr(V^* > y)dy = \int_0^\infty Pr(V^* - \theta > z)dz \geq E[(V^* - \theta)]$$

$$\boxed{E[V(\tau)] \geq \frac{1}{2}\theta + \frac{1}{2}E[V^* - \theta] = \frac{1}{2}E[V^*]}$$

Results: “Prophet Inequality”

V^*- value of the best candidate in hindsight.

V_λ - value of the candidate chosen by an optimal λ-biased stopping rule.

Theorem: $\dfrac{\text{prophet}}{\lambda\text{–biased}} = \dfrac{E[V^*]}{E[V_\lambda]} \leq \lambda + 2$ and this is tight.

Proof Sketch (Samuel-Cahn’84) – show that the following stopping rule achieves the desired bound: hire the first candidate with value greater than θ where $\Pr\left(V^* > \theta\right) = \dfrac{\lambda+1}{\lambda+2}$.

Tightness: 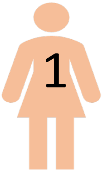

$$v_1 = \frac{1}{1+(1-\varepsilon)\lambda}$$

$$v_2 = \begin{cases} 1/\varepsilon & w.p\ \varepsilon \\ 0 & w.p\ 1-\varepsilon \end{cases}$$

λ-biased always selects candidate 1:

$$\frac{1}{1+(1-\varepsilon)\lambda} = 1 - \frac{\lambda(1-\varepsilon)}{1+(1-\varepsilon)\lambda}$$

Expected value of a prophet:

$$(1-\varepsilon)\cdot\frac{1}{1+(1-\varepsilon)\lambda} + 1 = \frac{2-\varepsilon+\lambda(1-\varepsilon)}{1+(1-\varepsilon)\lambda}$$

Results: Comparison to unbiased

Theorem: $\frac{\text{unbiased}}{\lambda\text{–biased}} = \frac{E[V_0]}{E[V_\lambda]} \leq \lambda + 1$ and this is tight.

π_λ - optimal stopping rule for a λ-biased agent.

π_0 - optimal stopping rule for an unbiased agent.

V^*- value of the best candidate in hindsight.

V_λ - value of the candidate chosen by π_λ.

V_0 - value of the candidate chosen by π_0.

Proof Idea 1: Lower bound the expected utility of π_λ by π_0.

Expected value of π_0: $= E[V_0]$ Expected loss of π_0 : $\leq E[V^* - V_0] = E[V^*] - E[V_0]$

Results: Comparison to unbiased

Theorem: $\frac{\text{unbiased}}{\lambda\text{–biased}} = \frac{E[V_0]}{E[V_\lambda]} \leq \lambda + 1$ and this is tight.

π_λ - optimal stopping rule for a λ-biased agent.

π_0 - optimal stopping rule for an unbiased agent.

V^*- value of the best candidate in hindsight.

V_λ - value of the candidate chosen by π_λ.

V_0 - value of the candidate chosen by π_0.

Proof Idea 1: Lower bound the expected utility of π_λ by π_0.

Expected utility of π_λ $\geq E[V_0] - \lambda(E[V^*] - E[V_0]) = (\lambda+1)E[V_0] - \lambda E[V^*]$

Proof Idea 2: Apply the bounds that we have: Let

$\beta = \frac{E[V^*]}{E[V_\lambda]}$, $\Delta = \frac{E[V^*]}{E[V_0]}$. Then, $\frac{E[V_0]}{E[V_\lambda]} = \frac{\beta}{\Delta}$.

$E[V_\lambda] = \frac{E[V^*]}{\beta} \geq$ Expected utility of π_λ $\geq (\lambda+1)\frac{E[V^*]}{\Delta} - \lambda E[V^*] \Rightarrow \beta \leq \frac{\Delta}{1-\lambda(\Delta-1)}$

We conclude: $\frac{\beta}{\Delta} \leq \min\left\{\frac{1}{1-\lambda(\Delta-1)}, \frac{\lambda+2}{\Delta}\right\} \Rightarrow \frac{E[V_0]}{E[V_\lambda]} \leq \lambda + 1$

Recall $\beta \leq \lambda + 2$

Some Monotonicity Results

	Expected utility	Expected value of candidate
Reference value increases	↘	↗ or ↘
λ increases	↘	↘
Adding a candidate at the end	↗	↗
Adding a candidate at the beginning	↗ or ↘ If ↘ at most factor $\lambda + 1$	↗ or ↘ If ↘ at most factor $\lambda + 1$

Ordering Problems

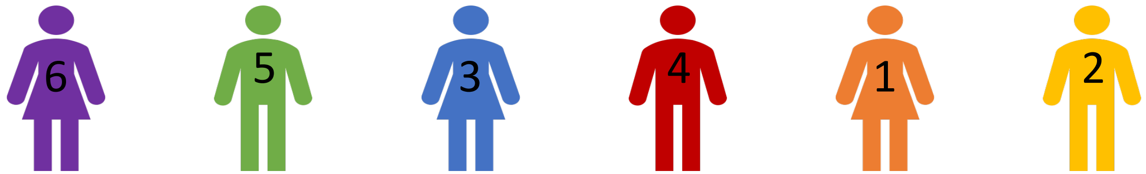

Candidates are ordered uniformly at random:

- $\frac{\text{prophet}}{\lambda\text{–biased}} = \frac{E[V^*]}{E[V_\lambda]} \leq n$ (tight).
- $\frac{\text{prophet}}{\lambda\text{–biased}} = \frac{E[V^*]}{E[V_\lambda]} \leq \Theta(\log \lambda)$ (essesntially tight).

Exponential improvement

The order is chosen to maximize the expected value:

- $\frac{\text{prophet}}{\lambda\text{–biased}} = \frac{E[V^*]}{E[V_\lambda]} \leq n$ (tight) –distributions that have at least 3 values in their support.
- $\frac{\text{prophet}}{\lambda\text{–biased}} = \frac{E[V^*]}{E[V_\lambda]} \leq 2$ (tight) –distributions that have 2 values in their support.

Exponential improvement

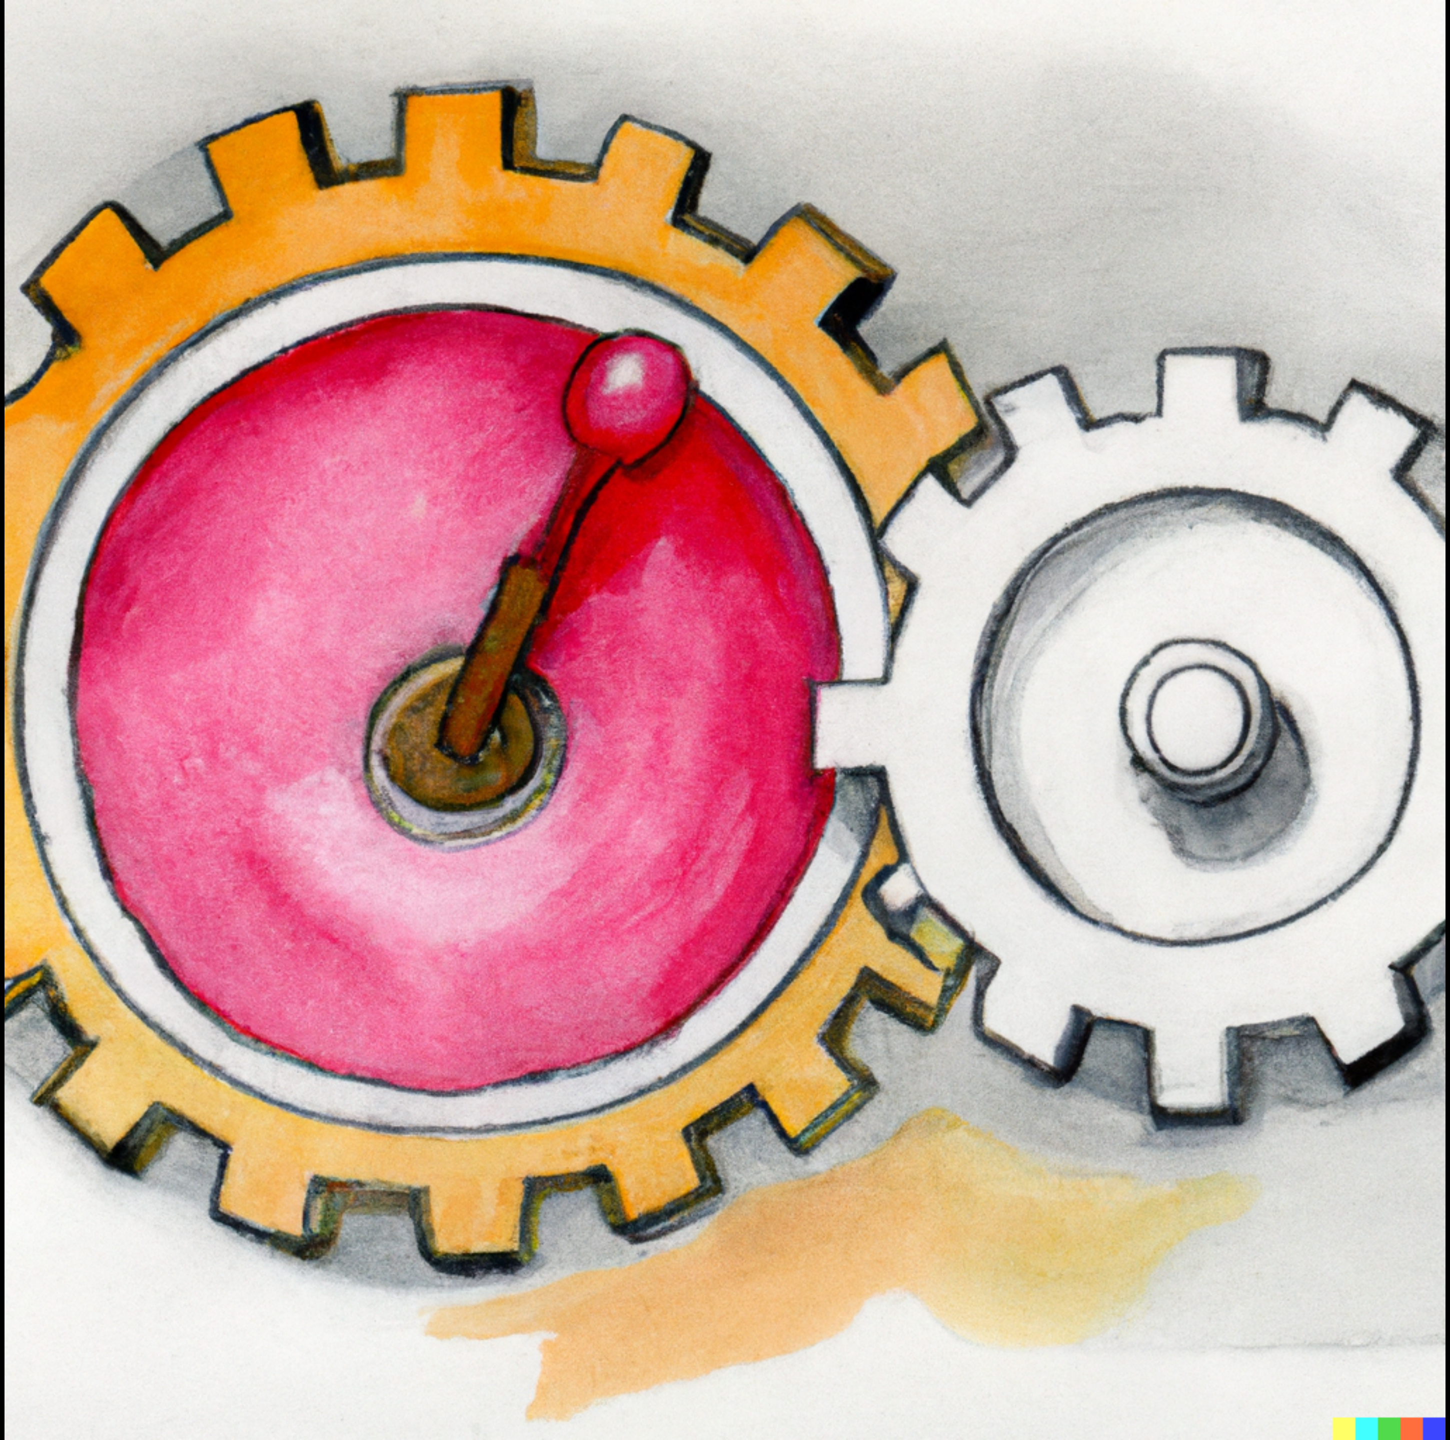

Why did George Akerlof not make it to the post office?

* An agent has to ship a package till day n.
* One-time effort cost c to ship the package.
* Loss-of-use cost x each day it hasn't been shipped.

Cost for shipping the package on day t: c+tx.

=> Package should be sent on the first day.

Instead of sending the package on the first day,
Akerlof procrastinated...

Present Bias/focus

A model of present bias (Akerlof'91, Strotz'55, Pollak'68):

Costs incurred today are more salient: raised by factor $b > 1$.

- The cost for sending the package today is bc
- The cost for sending it tomorrow is $bx + c$.
- Tomorrow is preferable if $bc > bx + c$.

General framework: quasi-hyperbolic discounting (Laibson'97)

Can model procrastination, task abandonment (O'Donoghue-Rabin'08) and benefits of choice reduction (Ariely and Wertenbroch'02, Kaur-Kremer-Mullainathan'10).

Graph Theoretic Framework for Planning

* Previous (theoretical) work mainly focused on the question of when to complete a single task (e.g., Akerlof'91, O'Donoghue-Rabin'99) or when to execute steps of a long term project (O'Donoghue-Rabin'08).

* What happens if the task structure is more complicated?

* For example, a student should decide which elective classes to take.

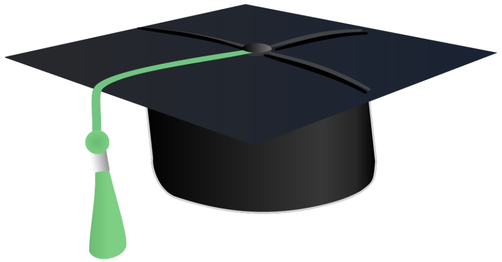

Graph Theoretic Framework for Present Bias

Kleinberg and Oren'14

* An agent has to achieve some goal (get from s to t).
* Nodes represent progress points towards the goal.
* Edges represent the tasks that the agent should complete to advance forward.
* The graph is a directed acyclic graph.

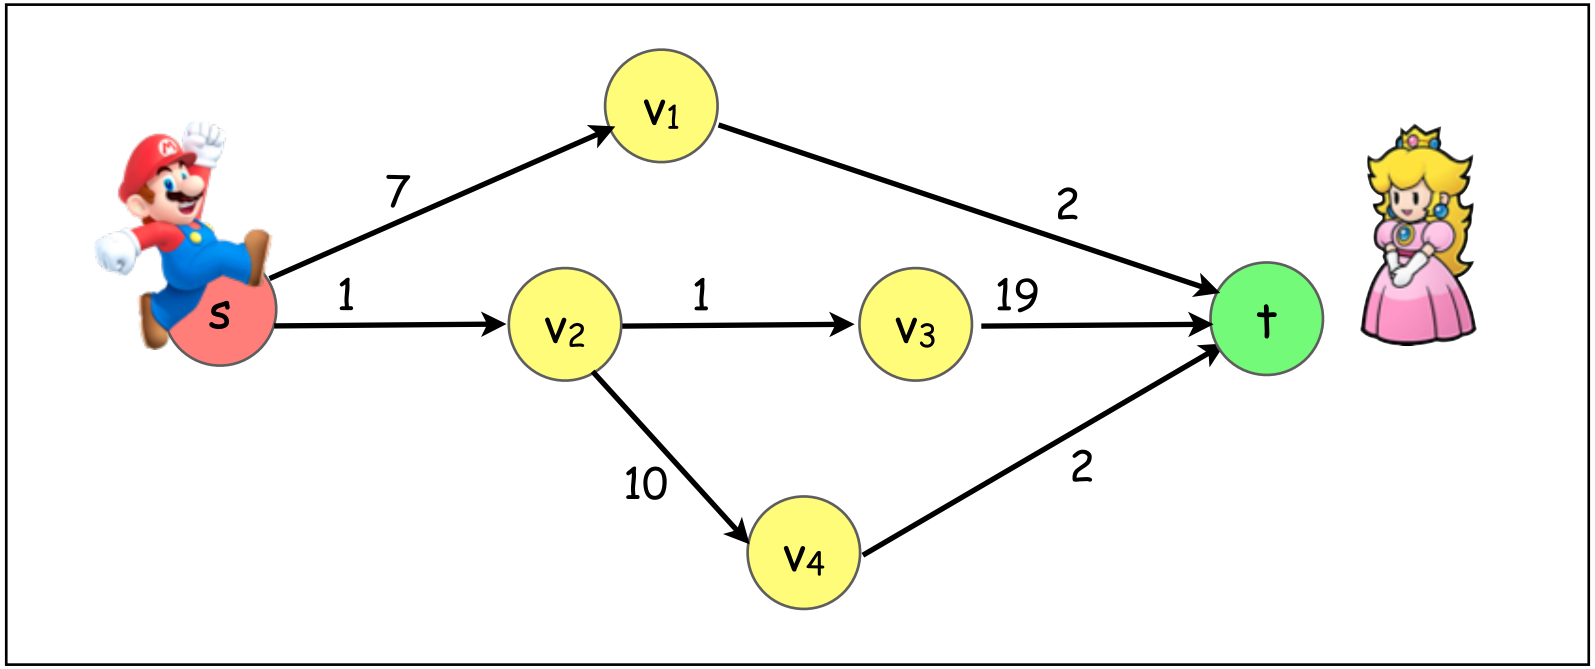

Graph Theoretic Framework for Present Bias

Kleinberg and Oren'14

* Agent has to achieve some goal (get from s to t).
* A *naive* agent plans to follow the shortest path from s to t.
* From a given node, immediately outgoing edges have costs multiplied by b > 1.

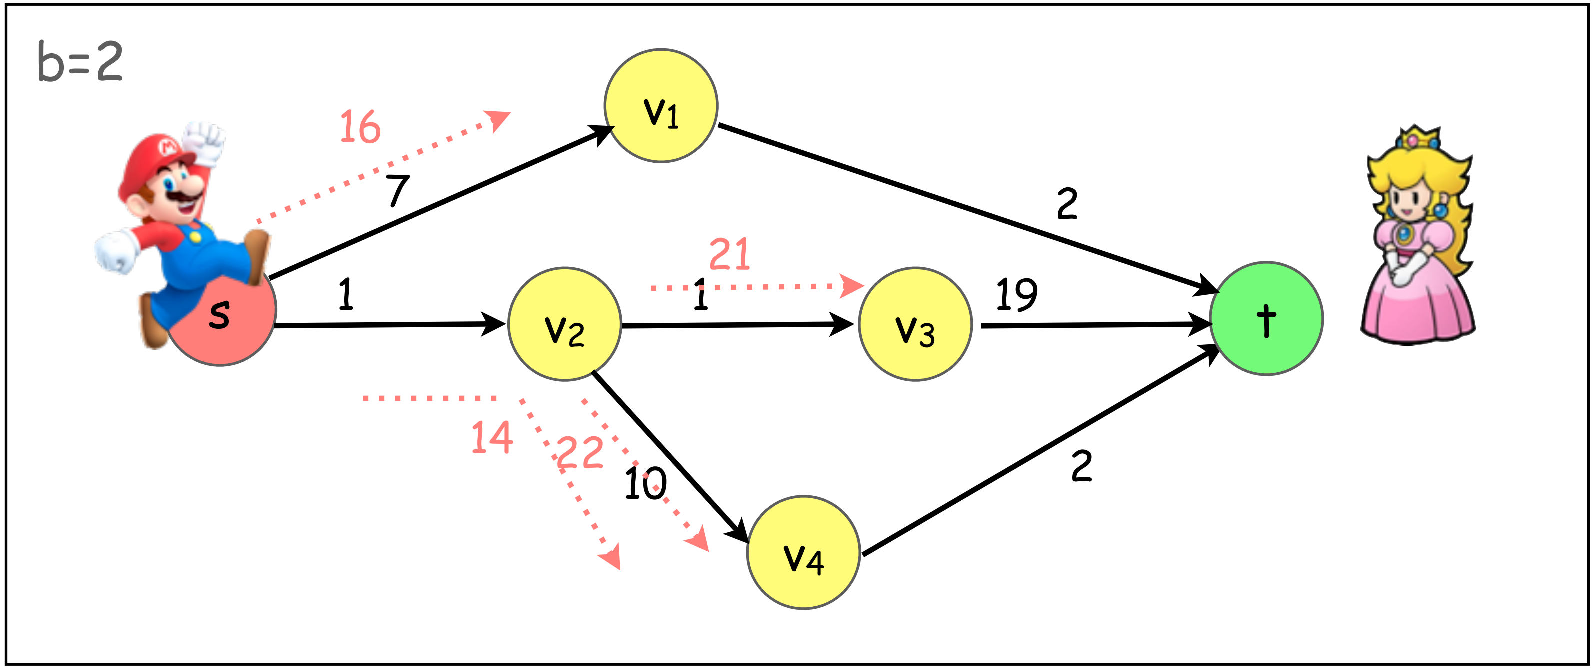

A naive agent constantly changes its plan

Graph Theoretic Framework for Present Bias

Kleinberg, Oren and Raghavan'16

* Agent has to achieve some goal (get from s to t).
* From a given node, immediately outgoing edges have costs multiplied by b > 1.
* **Sophisticated** agent takes into account his bias when planning a path.

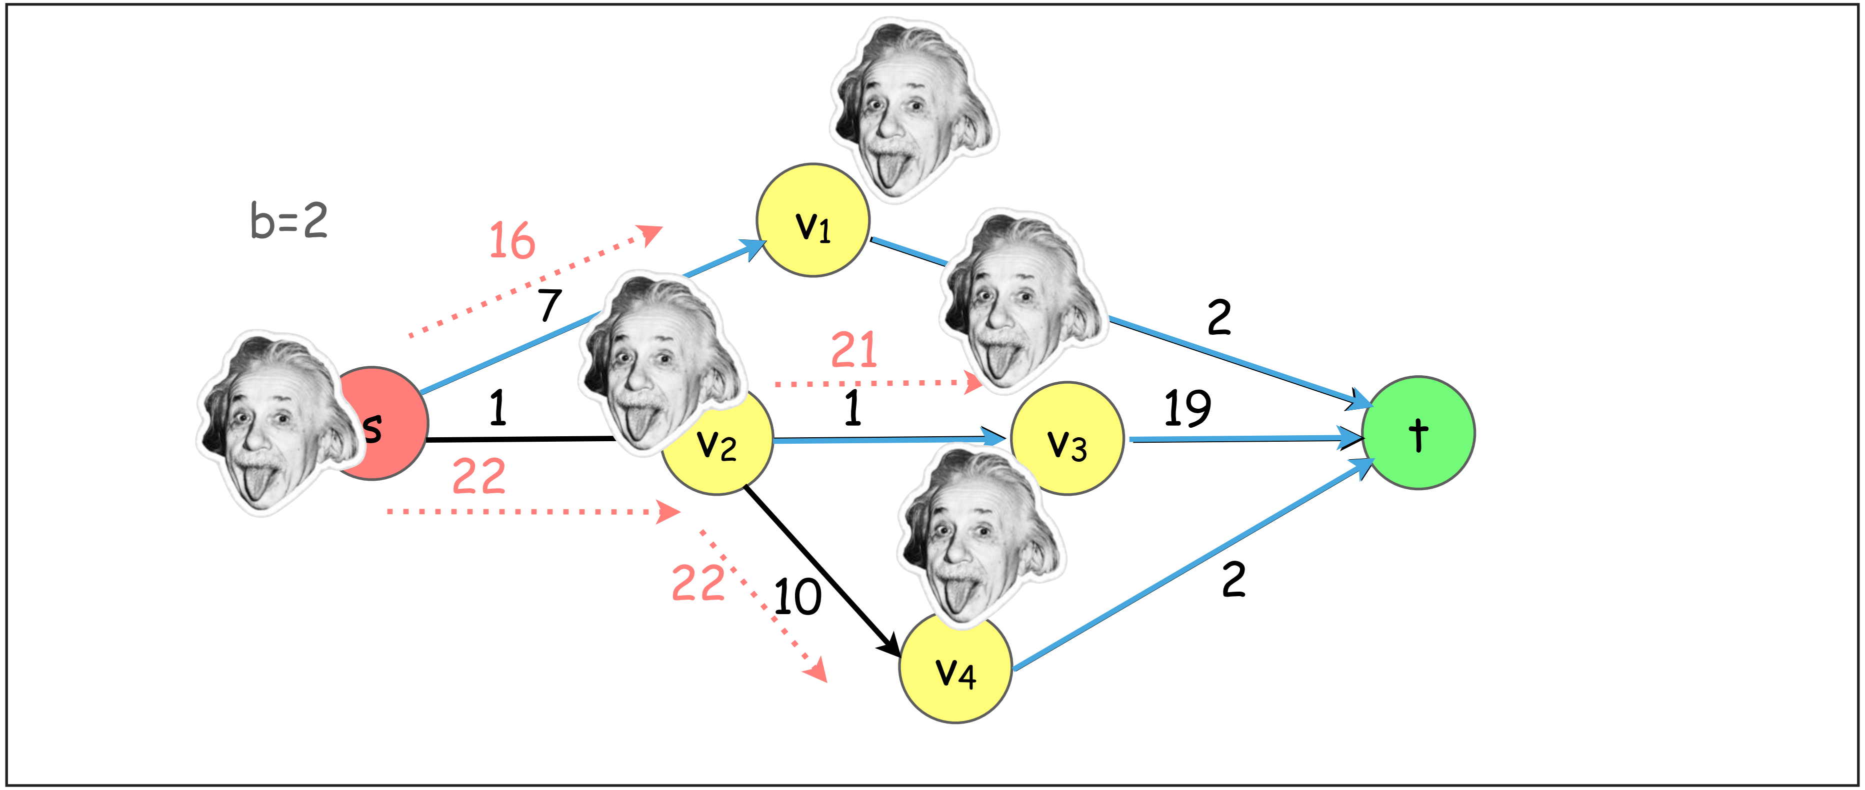

A sophisticated agent makes a plan and sticks to it

Cost Ratio

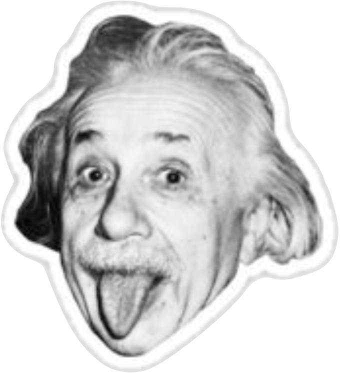

What is the ratio between the cost of a present biased agent and an optimal agent?

Answer: can be as high as b^n

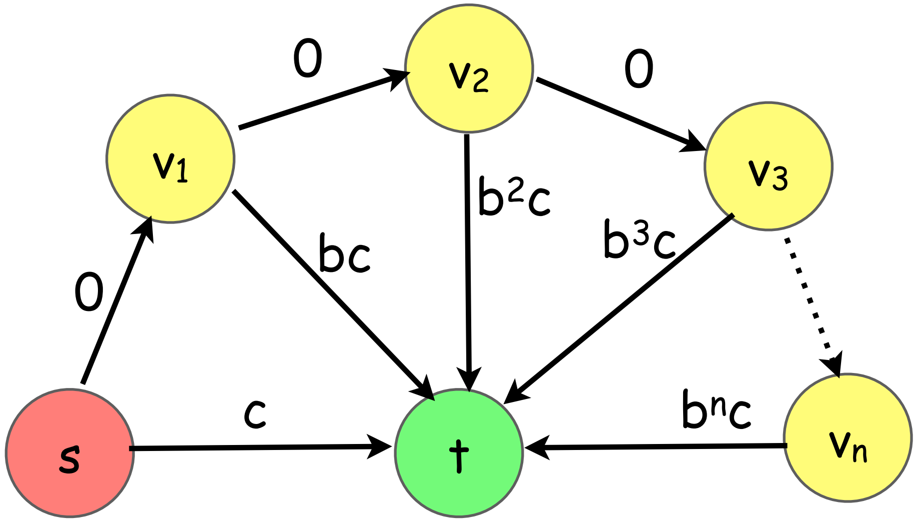

All instances with exponential cost ratio contain this graph as a subgraph (formally minor).

(Kleinberg and Oren'14, Tang et al.'15)

Answer: at most b

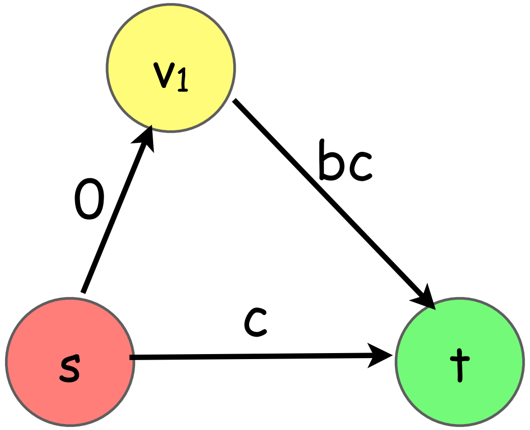

Proof by an inductive argument.

(Kleinberg, Oren and Raghavan'16)

Variation: Paths with Rewards

- An agent that reaches the target receives a reward r.
- A naive agent at node v continues to traverse the graph if

$$min_{u \in N(v)}\ b \cdot c(v,u) + d(u,t) < r$$

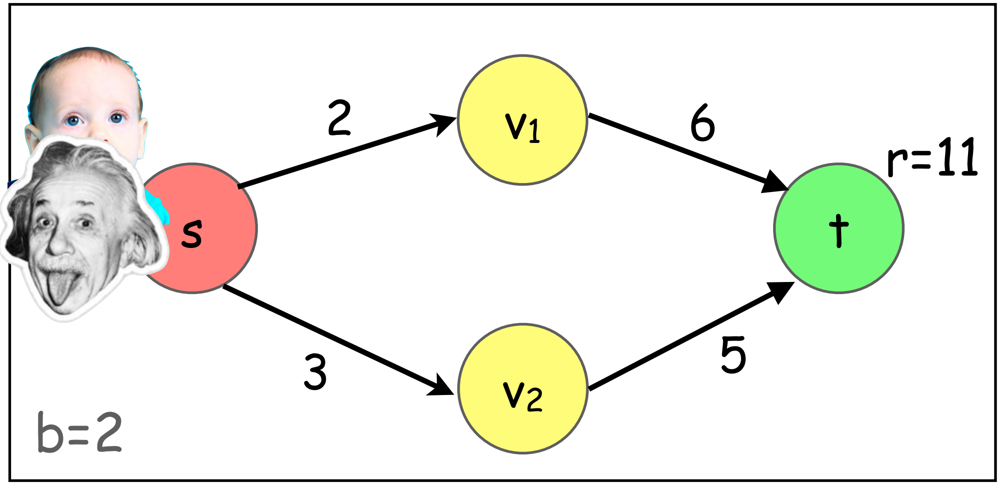

- A sophisticated agent goes over the graph in reverse topological order and prunes "dead ends".

Distinction: A Naive agent might stop traversing the graph. A sophisticated agent will either stay at s or reach t.

Choice Reduction

* Choice reduction problem: given G, not traversable by an agent, is there a subgraph of G that is traversable?
* First attempt: if there is a traversable subgraph in G, then there is a traversable subgraph that is a path.
* Correct for sophisticated agents but wrong for naive agents.
* A characterization of the structure of minimal traversable subgraphs for naive agents.

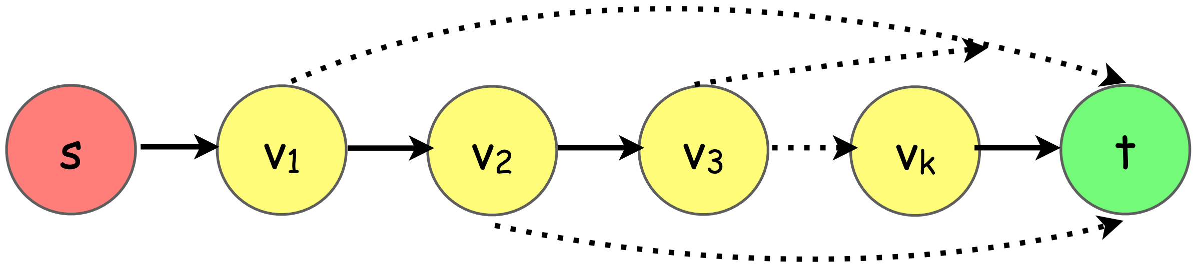

* For naive agents: NP-completeness [Tang et al '15], Hard to approximate by $\sqrt{n}$ [Albers and Kraft '16].

Minimal Reward

Given a graph what is the minimal reward required for motivating the agent to traverse the graph?

Answer: the maximal perceived cost of the path the agent will take without a reward

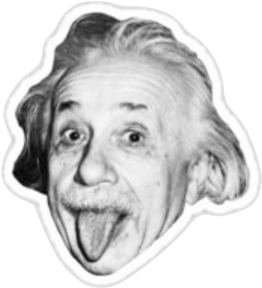

Answer: **Open**

Main Challenge: non-monotone in the reward. The graph might be traversable for a reward r but not traversable for r'>r.

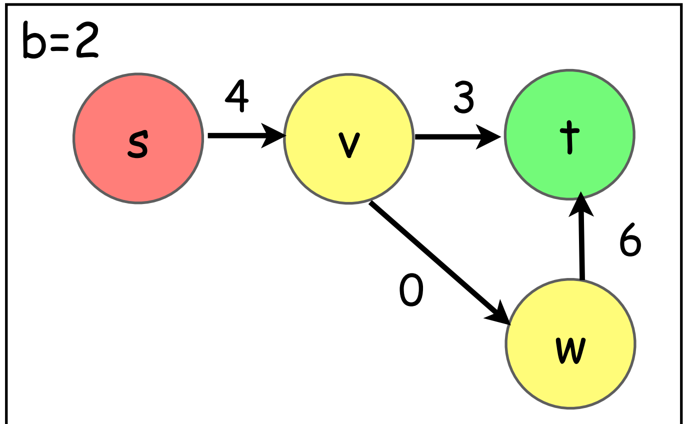

traversable for r=11 but not for r=13.

[KOR'16] As r increases there can be an exponential number of switches between traversable and non traversable.

Sunk Cost and Present Bias

Kleinberg, Oren and Raghavan ’17

* Sunk cost - taking into account past costs even when these are irrelevant. (Arkes and Blumer, 1985, Thaler, 1980, 1999)
* Under our framework an agent exhibiting sunk cost that already exhibited cost C will continue to traverse the graph if λC is greater than the perceived cost for reaching the target minus r.

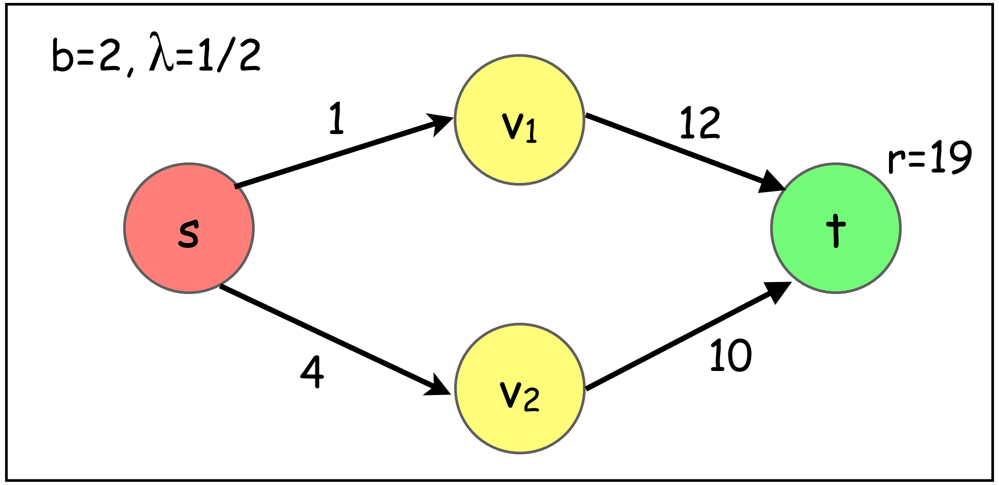

Summary

- Optimal stopping with reference-dependent agents:
  - Agents tend to stop searching prematurely.
  - Random ordering or picking a specific order can help a lot.

- Planning:
  - Graph theoretic framework for planning related biases.
  - Discussed sophisticated and naive agents.
  - Different Phenomena captured by this framework.