Long-Range Planning and Behavioral Biases: A Computational Approach

Jon Kleinberg

Including joint work with Manish Raghavan and Sigal Oren.

Cornell University

Long-Range Planning

WIKIPEDIA
The Free Encyclopedia

stackoverflow

Growth in on-line systems where users and groups have long visible careers and set long-range goals.

- Reputation, promotion, status, individual achievement.
- On-line groups that create multi-step tasks and set timelines and deadlines.

Badges on Stack Overflow

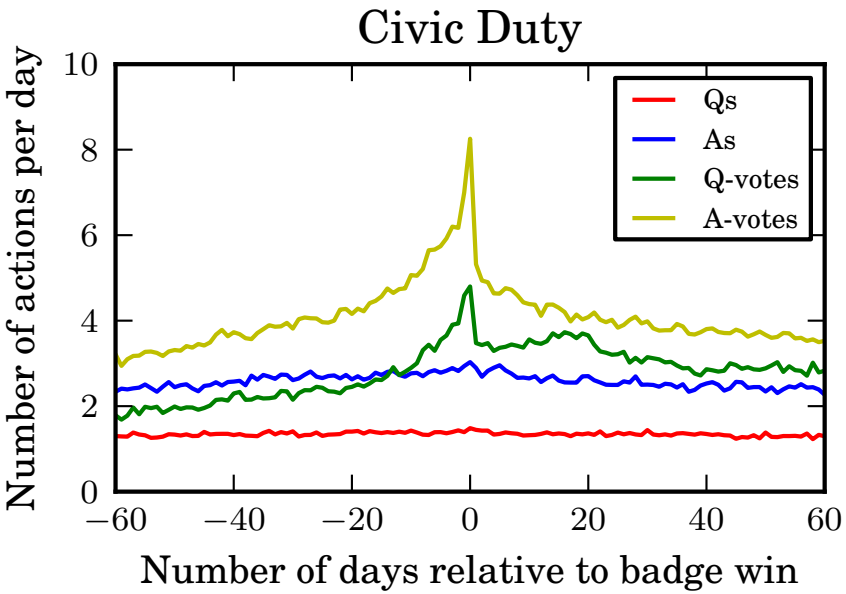

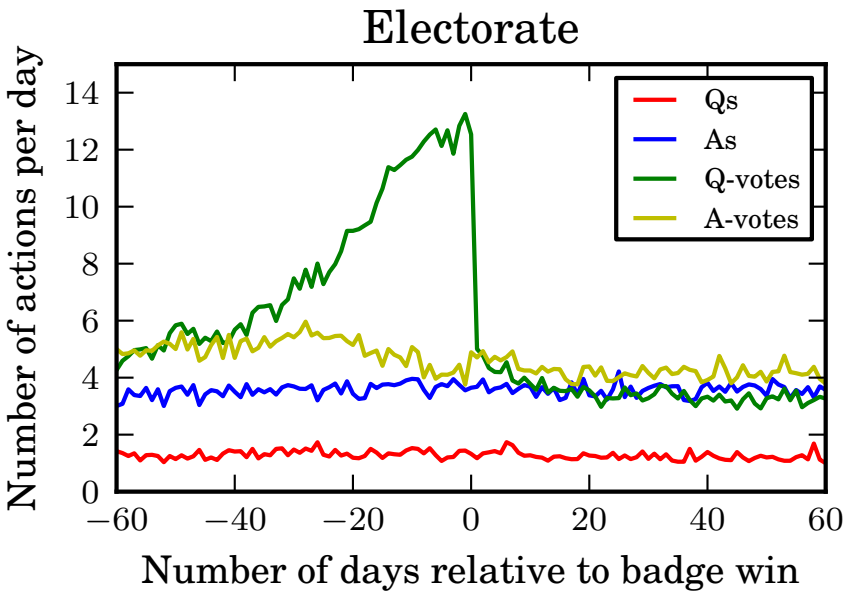

Badges, Milestones, and Incentives

- The Placement Problem:
  Given a desired mixture of actions, how should one define milestones to (approximately) induce these actions?
- How do badges and milestones derive their value?
  Social / Motivational / Transactional?

Antin-Churchill 2011, Deterding et al 2011, Chawla-Hartline-Sivan 2012,

Easley-Ghosh 2013, Anderson-Huttenlocher-Kleinberg-Leskovec 2013

Planning and Time-Inconsistency

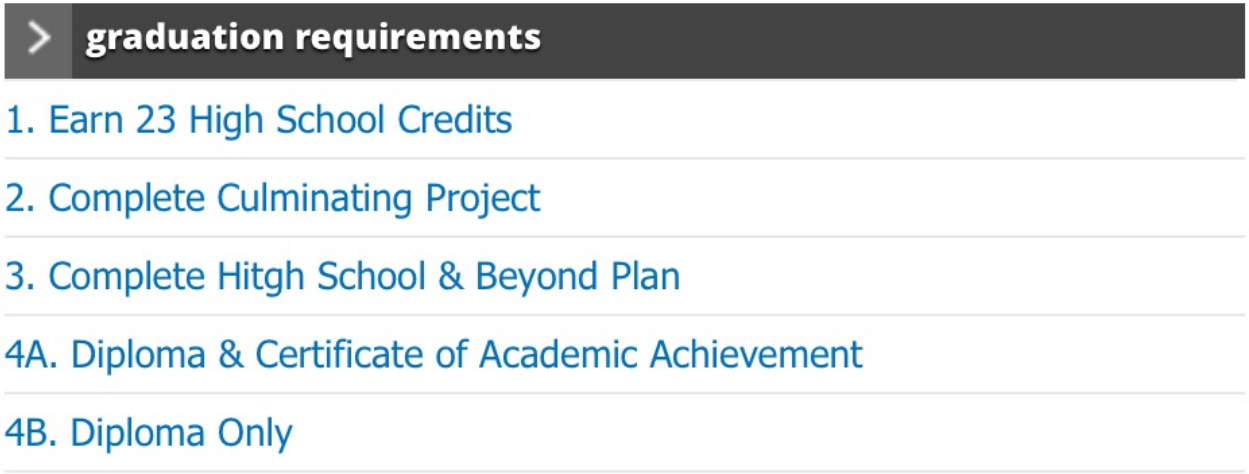

Tacoma Public School System

Fundamental behavioral process: Making plans for the future.

- Plans can be multi-step.
- Natural model: agents chooses optimal sequence given costs and benefits.

What could go wrong?

- Costs and benefits are unknown, and/or genuinely changing over time.
- Time-inconsistency.

Planning and Time-Inconsistency

Fundamental behavioral process: Making plans for the future.

- Plans can be multi-step.
- Natural model: agents chooses optimal sequence given costs and benefits.

What could go wrong?

- Costs and benefits are unknown, and/or genuinely changing over time.
- Time-inconsistency.

Planning and Time-Inconsistency

Fundamental behavioral process: Making plans for the future.

- Plans can be multi-step.
- Natural model: agents chooses optimal sequence given costs and benefits.

What could go wrong?

- Costs and benefits are unknown, and/or genuinely changing over time.
- Time-inconsistency.

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

Agent must ship a package sometime in next n days.

- One-time effort cost c to ship it.
- Loss-of-use cost x each day hasn't been shipped.

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

Agent must ship a package sometime in next n days.

- One-time effort cost c to ship it.
- Loss-of-use cost x each day hasn't been shipped.

An optimization problem:

- If shipped on day t, cost is $c + tx$.
- Goal: $\min_{1 \le t \le n} c + tx$.
- Optimized at $t = 1$.

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

Agent must ship a package sometime in next n days.

- One-time effort cost c to ship it.
- Loss-of-use cost x each day hasn't been shipped.

An optimization problem:

- If shipped on day t, cost is $c + tx$.
- Goal: $\min_{1 \le t \le n} c + tx$.
- Optimized at $t = 1$.

In Akerlof's story, he was the agent, and he *procrastinated*:

- Each day he planned that he'd do it tomorrow.
- Effect: waiting until day n, when it must be shipped, and doing it then, at a significantly higher cumulative cost.

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

Agent must ship a package sometime in next n days.

- One-time effort cost c to ship it.
- Loss-of-use cost x each day hasn't been shipped.

A model based on present bias [Akerlof 91; cf. Strotz 55, Pollak 68]

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

On day t:

- Remaining cost if sent today is bc.
- Remaining cost if sent tomorrow is $bx + c$.
- Tomorrow is preferable if $(b-1)c > bx$.

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

Agent must ship a package sometime in next n days.

- One-time effort cost c to ship it.
- Loss-of-use cost x each day hasn't been shipped.

A model based on present bias [Akerlof 91; cf. Strotz 55, Pollak 68]

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

On day t:

- Remaining cost if sent today is bc.
- Remaining cost if sent tomorrow is $bx + c$.
- Tomorrow is preferable if $(b-1)c > bx$.

General framework: quasi-hyperbolic discounting [Laibson 1997]

- Cost/reward c realized t units in future has present value $\beta\delta^t c$
- Special case: $\delta = 1$, $b = \beta^{-1}$, and agent is naive about bias.
- Can model procrastination, task abandonment [O'Donoghue-Rabin08], and benefits of choice reduction [Ariely and Wertenbroch 02, Kaur-Kremer-Mullainathan 10]

Cost Ratio

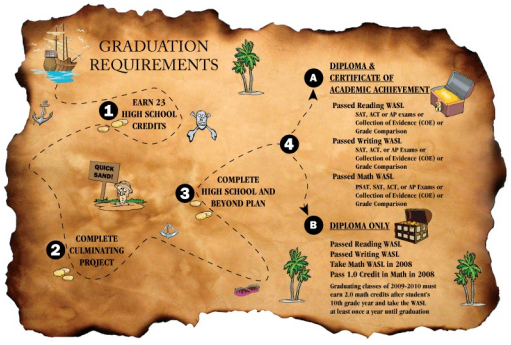

Cost ratio:

$$\frac{\text{Cost incurred by present-biased agent}}{\text{Minimum cost achievable}}$$

Across all stories in which present bias has an effect, what's the worst cost ratio?

$$\max_{\text{stories } S} \text{cost ratio}(S).$$

Cost Ratio

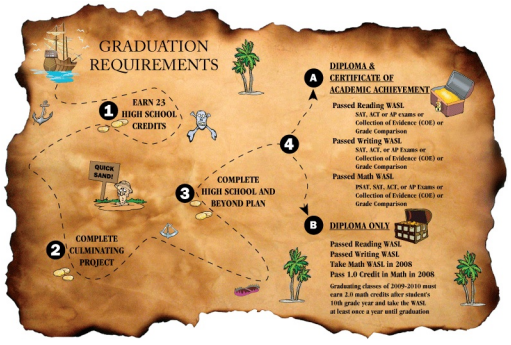

Cost ratio:

$$\frac{\text{Cost incurred by present-biased agent}}{\text{Minimum cost achievable}}$$

Across all stories in which present bias has an effect, what's the worst cost ratio?

$$\max_{\text{stories } S} \text{cost ratio}(S).$$

???

A Graph-Theoretic Framework

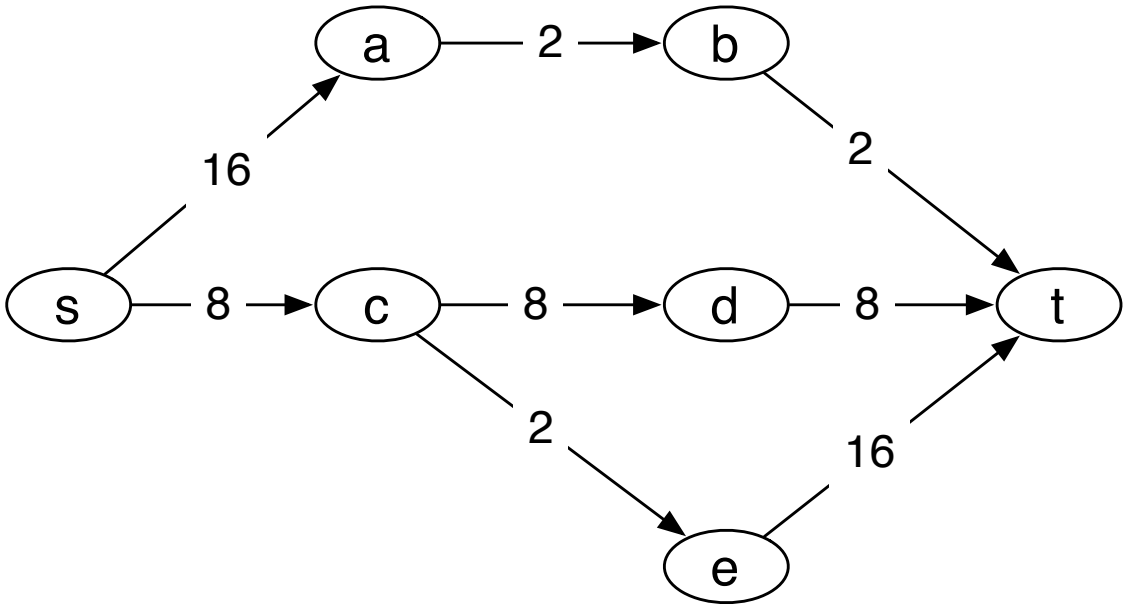

Use graphs as basic structure to represent scenarios

[Kleinberg-Oren 2014]

- Agent plans to follow cheapest path from s to t.
- From a given node, immediately outgoing edges have costs multplied by $b > 1$.

A Graph-Theoretic Framework

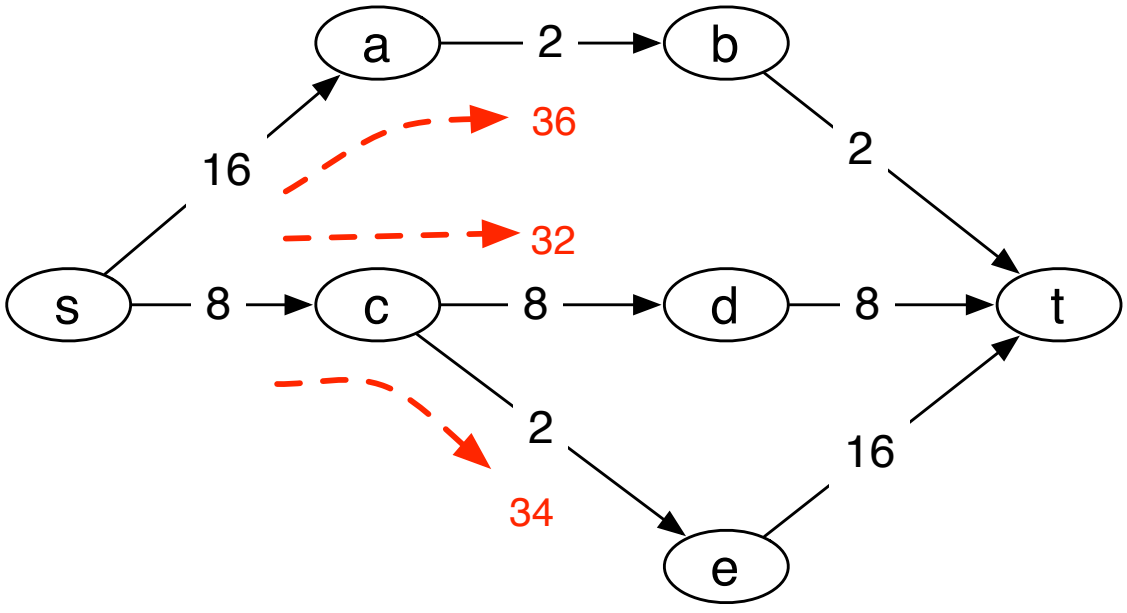

Use graphs as basic structure to represent scenarios

[Kleinberg-Oren 2014]

- Agent plans to follow cheapest path from s to t.
- From a given node, immediately outgoing edges have costs multplied by $b > 1$.

A Graph-Theoretic Framework

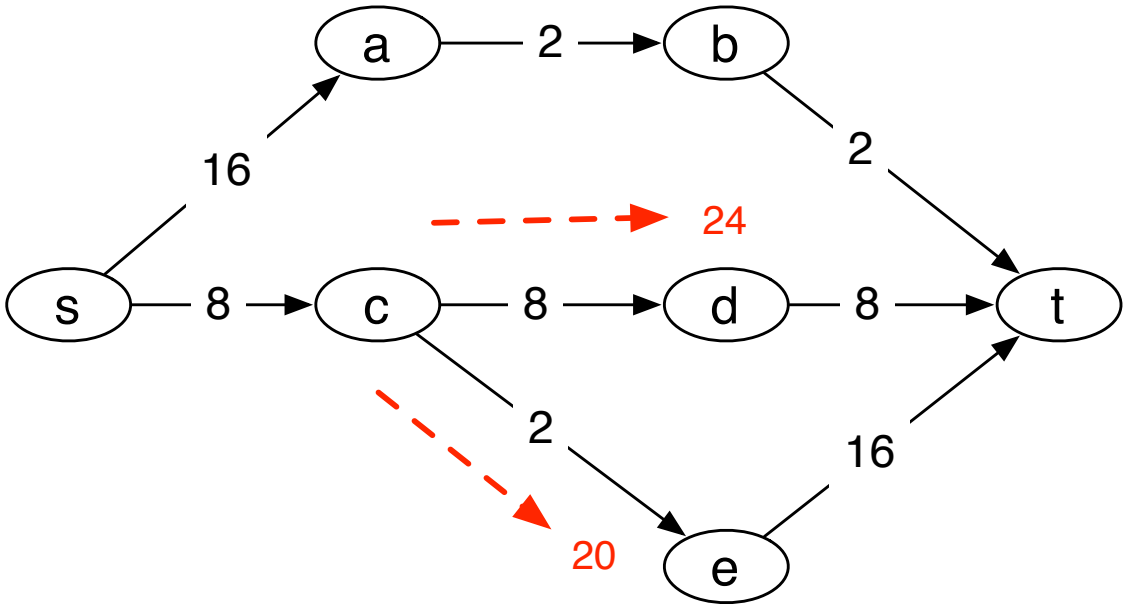

Use graphs as basic structure to represent scenarios

[Kleinberg-Oren 2014]

- Agent plans to follow cheapest path from s to t.
- From a given node, immediately outgoing edges have costs multplied by $b > 1$.

Example: Akerlof's Story as a Graph

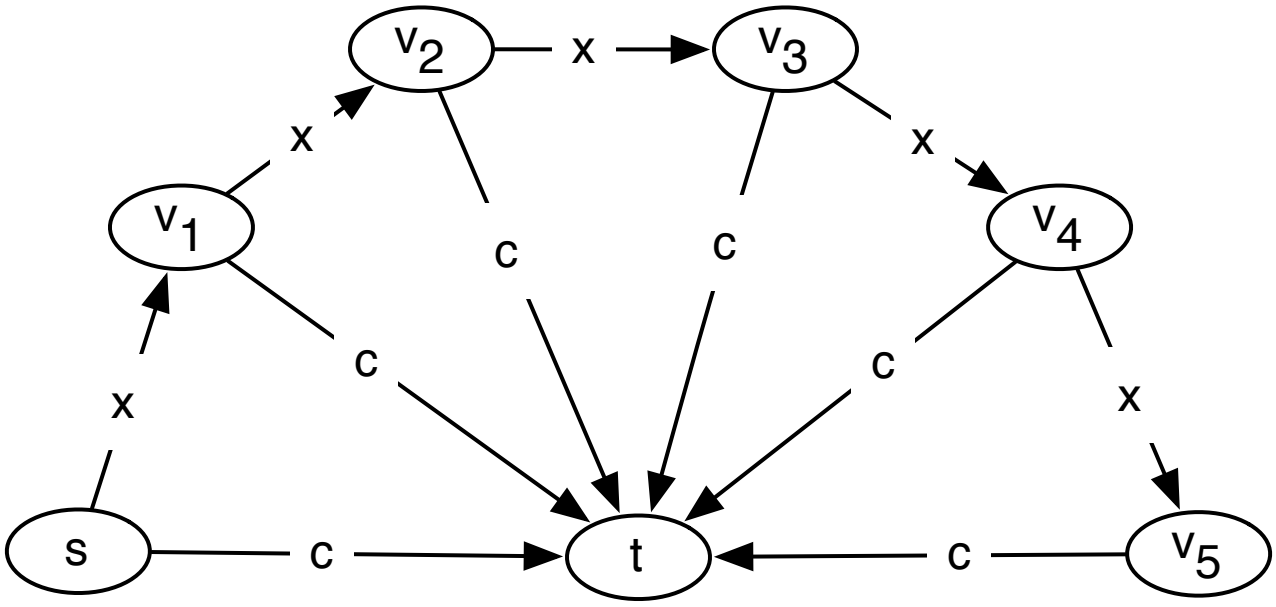

Node v_i = reaching day i without sending the package.

Paths with Rewards

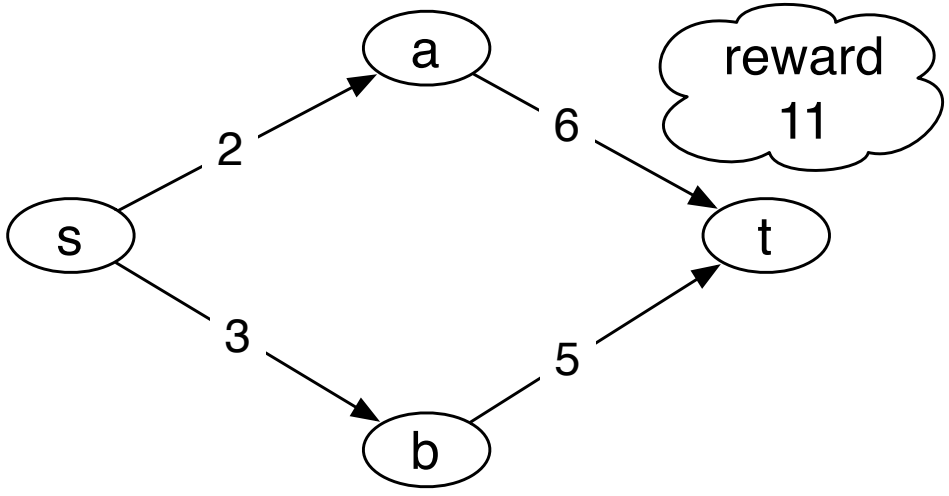

Variation: agent only continues on path if cost $\leq$ reward at t.

- Can model abandonment: agent stops partway through a completed path.
- Can model benefits of choice reduction: deleting nodes can sometimes make graph become traversable.

Paths with Rewards

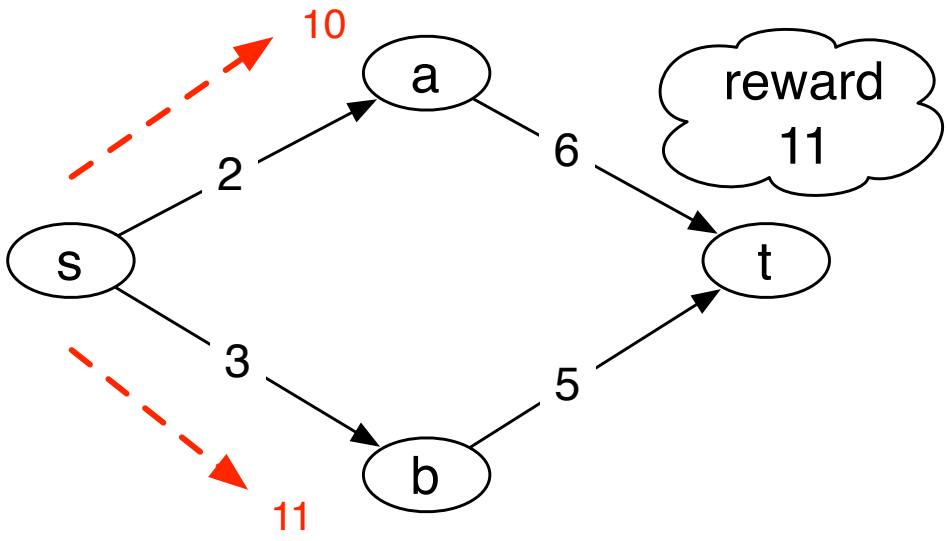

Variation: agent only continues on path if cost $\leq$ reward at t.

- Can model abandonment: agent stops partway through a completed path.
- Can model benefits of choice reduction: deleting nodes can sometimes make graph become traversable.

Paths with Rewards

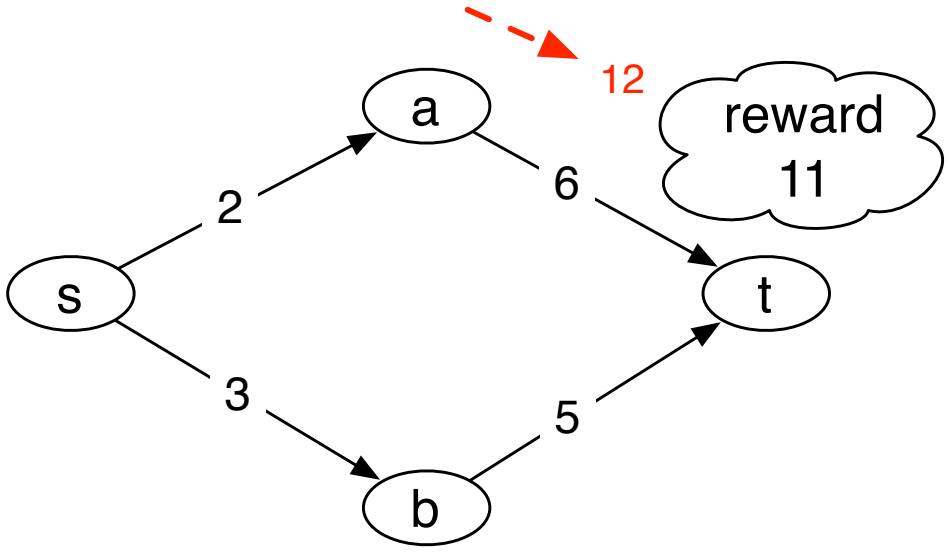

Variation: agent only continues on path if cost $\leq$ reward at t.

- Can model abandonment: agent stops partway through a completed path.
- Can model benefits of choice reduction: deleting nodes can sometimes make graph become traversable.

Paths with Rewards

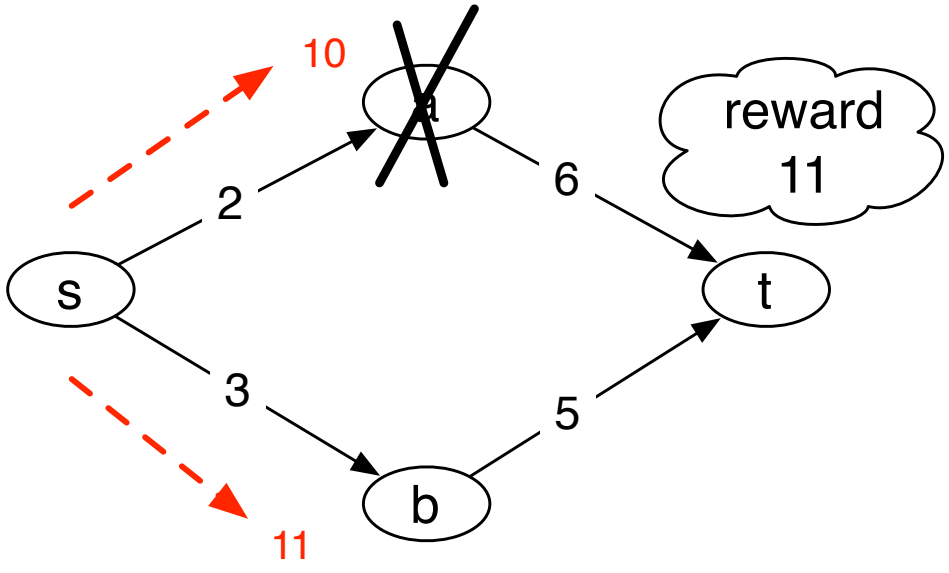

Variation: agent only continues on path if cost $\leq$ reward at t.

- Can model abandonment: agent stops partway through a completed path.
- Can model benefits of choice reduction: deleting nodes can sometimes make graph become traversable.

A More Elaborate Example

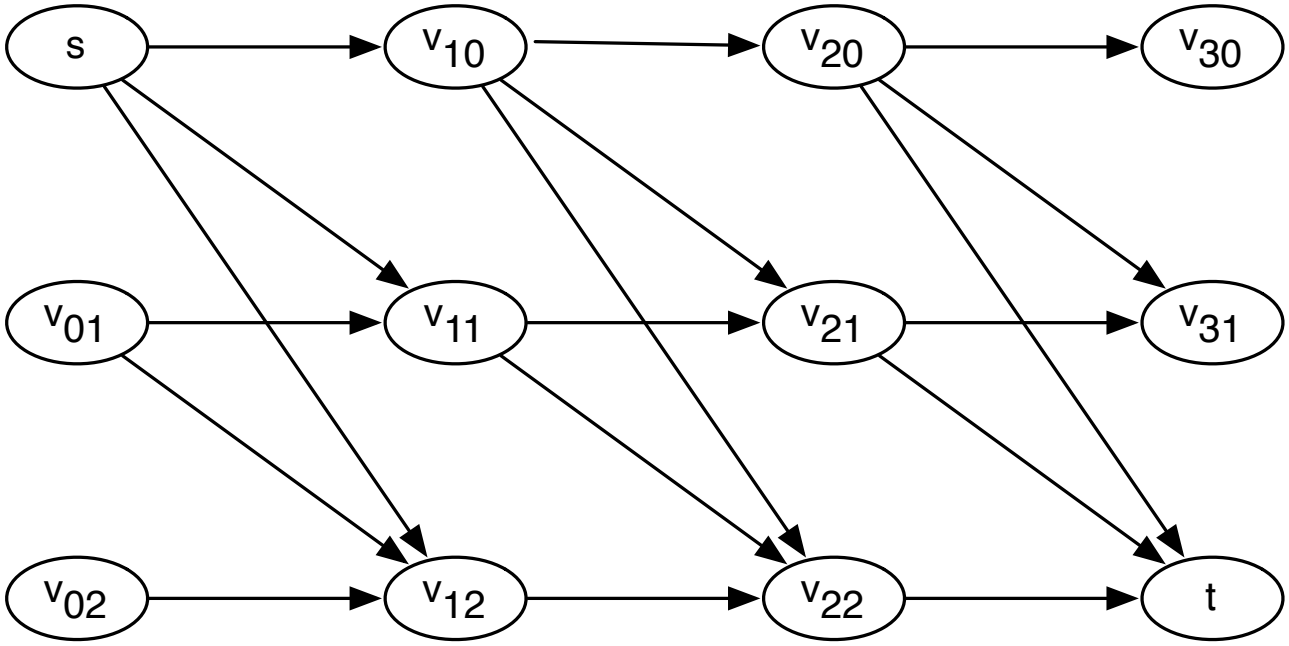

Three-week short course with two projects.

- Reward of 16 from finishing the course.
- Effort cost in a given week: 1 from doing no project, 4 from doing one, 9 from doing both.
- v_{ij} = the state in which i weeks of the course are done and the student has completed j projects.

A More Elaborate Example

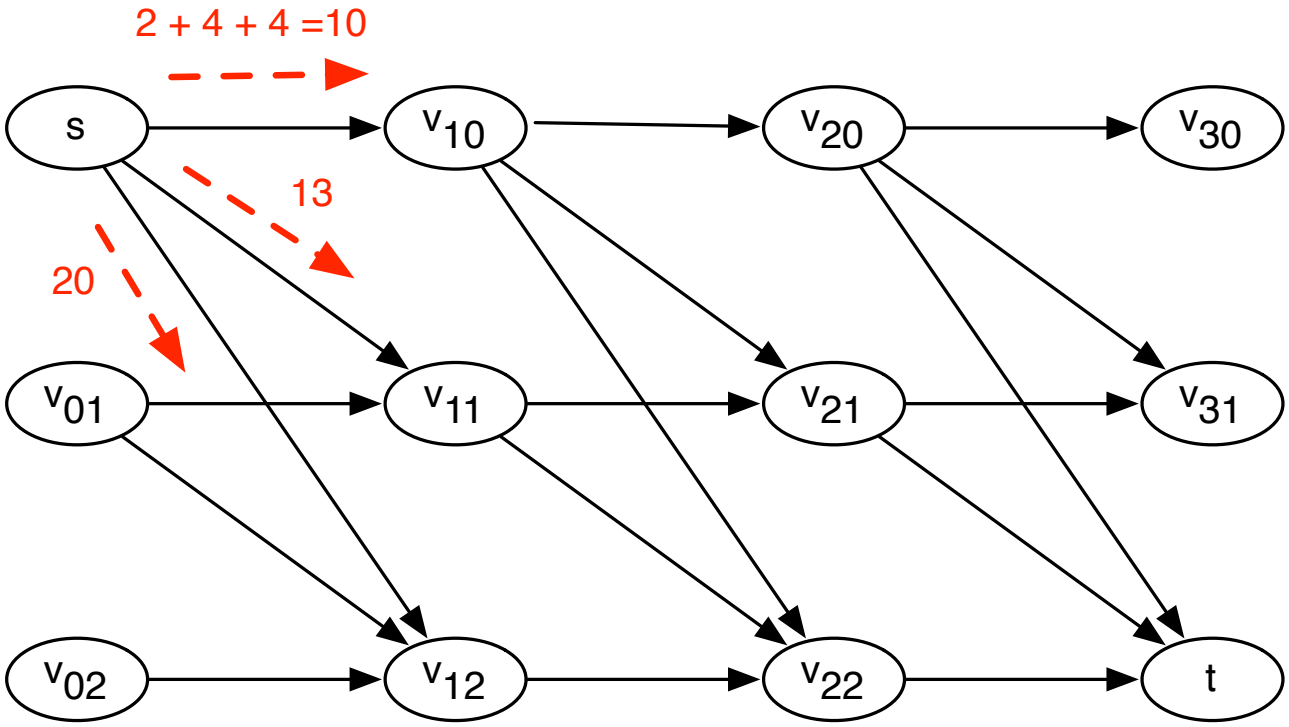

Three-week short course with two projects.

- Reward of 16 from finishing the course.
- Effort cost in a given week: 1 from doing no project, 4 from doing one, 9 from doing both.
- v_{ij} = the state in which i weeks of the course are done and the student has completed j projects.

A More Elaborate Example

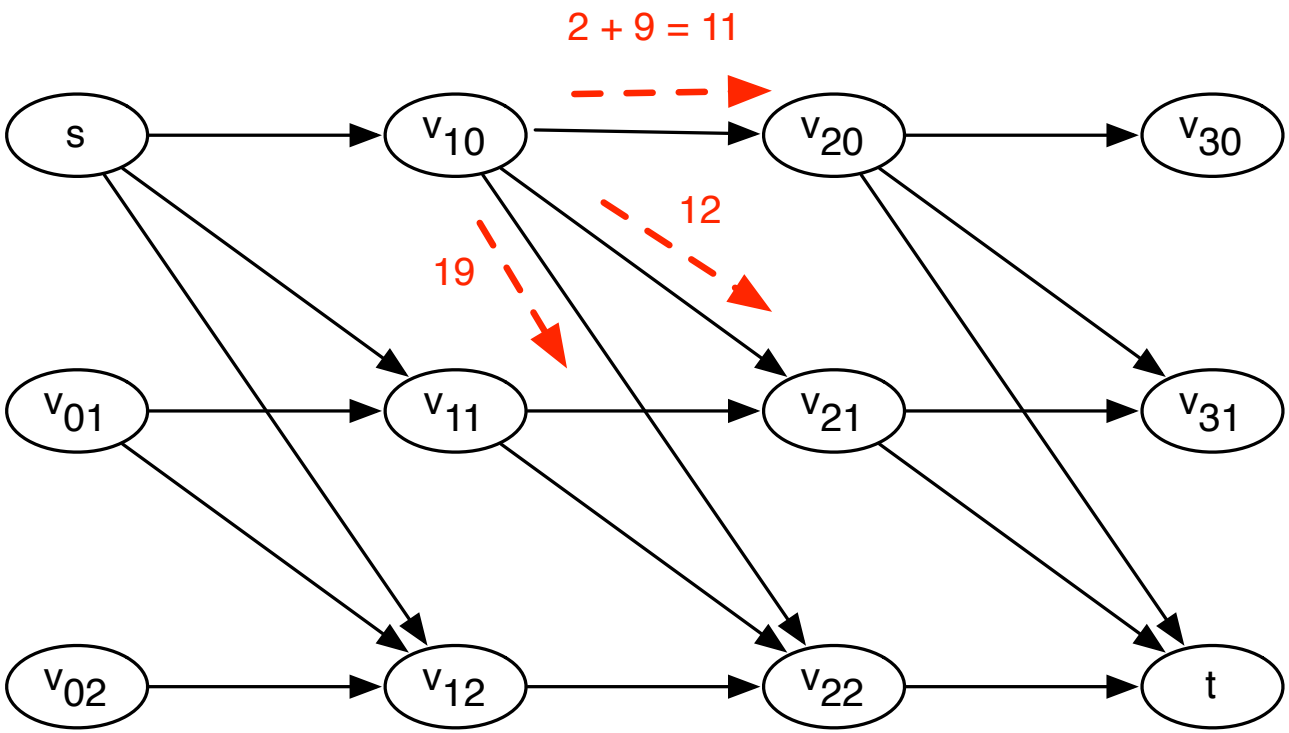

Three-week short course with two projects.

- Reward of 16 from finishing the course.
- Effort cost in a given week: 1 from doing no project, 4 from doing one, 9 from doing both.
- v_{ij} = the state in which i weeks of the course are done and the student has completed j projects.

A More Elaborate Example

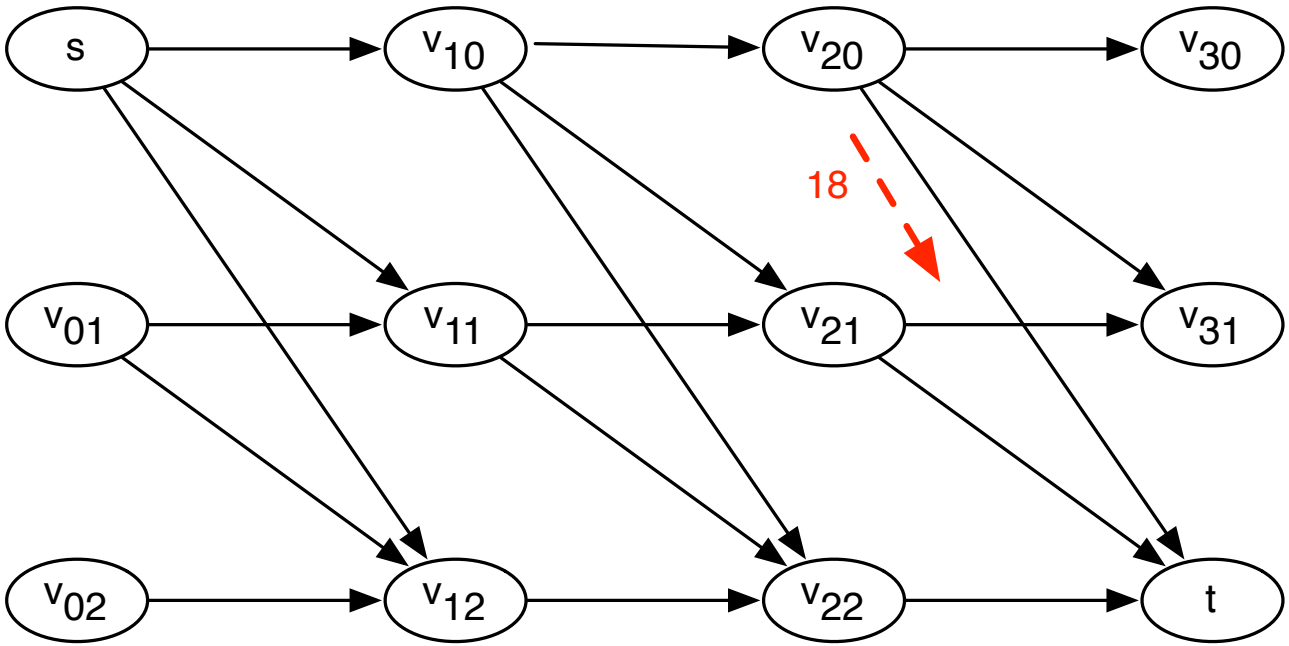

Three-week short course with two projects.

- Reward of 16 from finishing the course.
- Effort cost in a given week: 1 from doing no project, 4 from doing one, 9 from doing both.
- v_{ij} = the state in which i weeks of the course are done and the student has completed j projects.

A More Elaborate Example

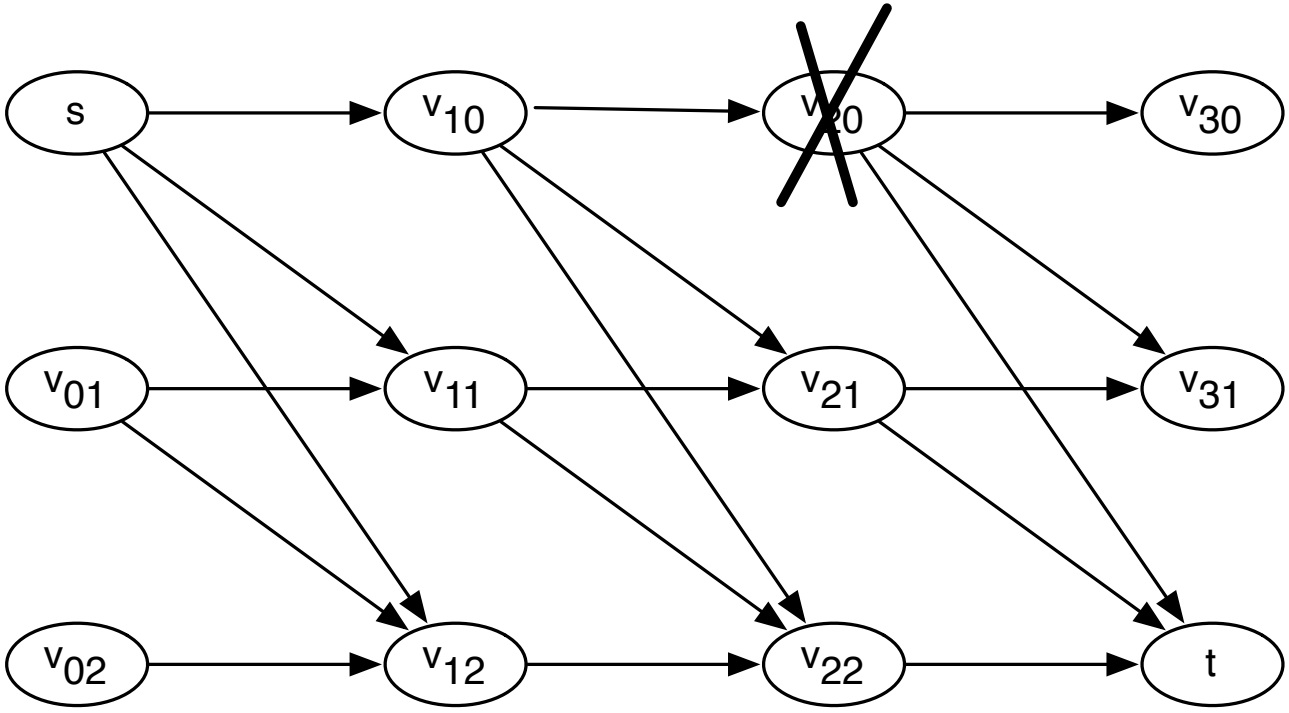

Three-week short course with two projects.

- Reward of 16 from finishing the course.
- Effort cost in a given week: 1 from doing no project, 4 from doing one, 9 from doing both.
- v_{ij} = the state in which i weeks of the course are done and the student has completed j projects.

A More Elaborate Example

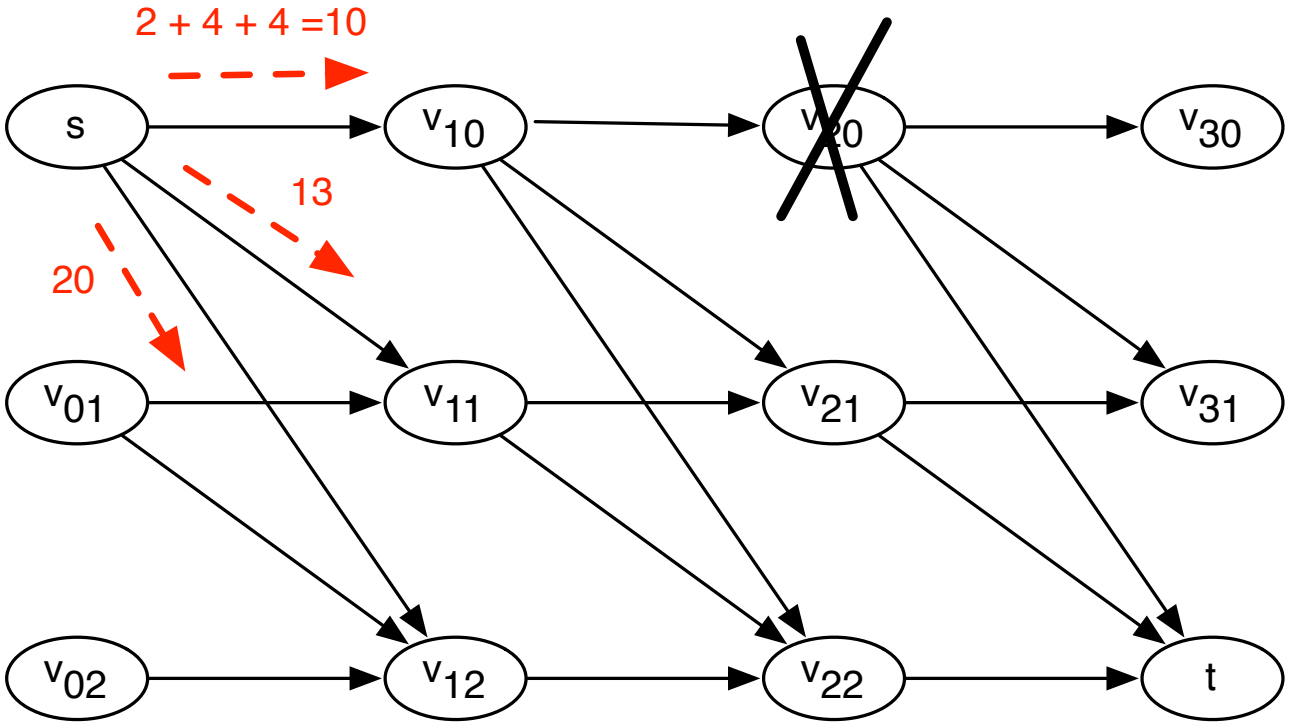

Three-week short course with two projects.

- Reward of 16 from finishing the course.
- Effort cost in a given week: 1 from doing no project, 4 from doing one, 9 from doing both.
- v_{ij} = the state in which i weeks of the course are done and the student has completed j projects.

A More Elaborate Example

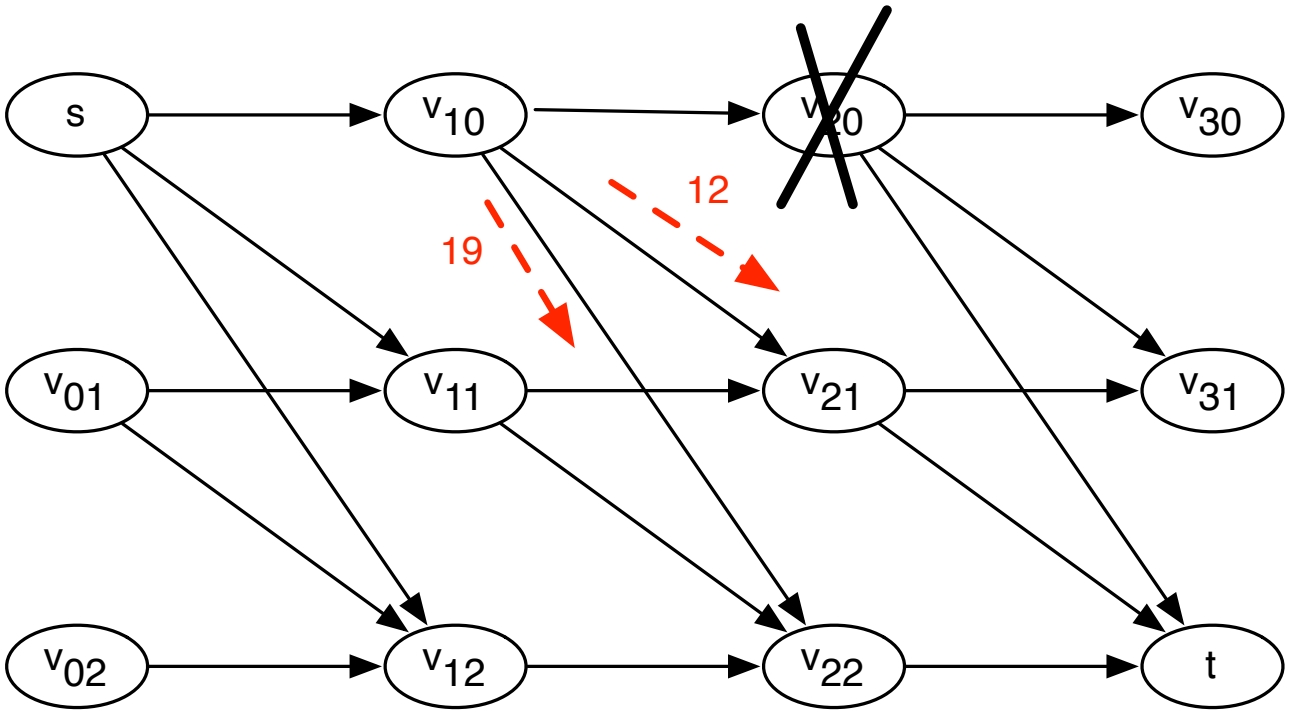

Three-week short course with two projects.

- Reward of 16 from finishing the course.
- Effort cost in a given week: 1 from doing no project, 4 from doing one, 9 from doing both.
- v_{ij} = the state in which i weeks of the course are done and the student has completed j projects.

A Bad Example for the Cost Ratio

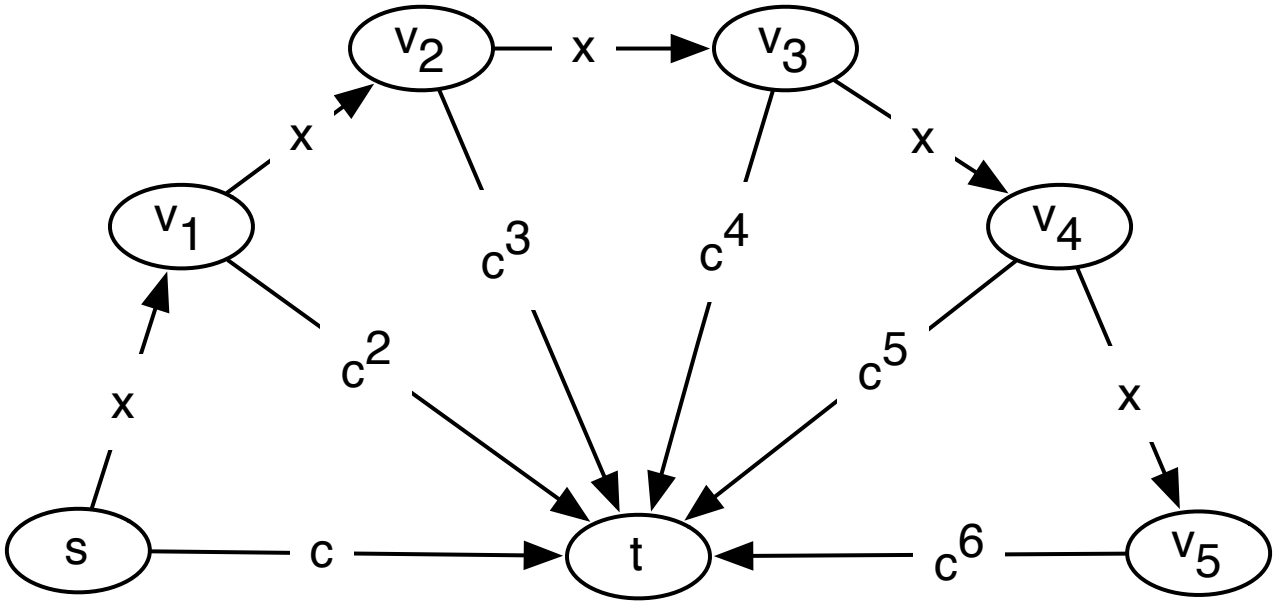

Cost ratio can be roughly b^n, and this is essentially tight.

Can we characterize the instances with exponential cost ratio?

- Goal, informally stated: Must any instance with large cost ratio contain Akerlof's story as a sub-structure?

Characterizing Bad Instances via Graph Minors

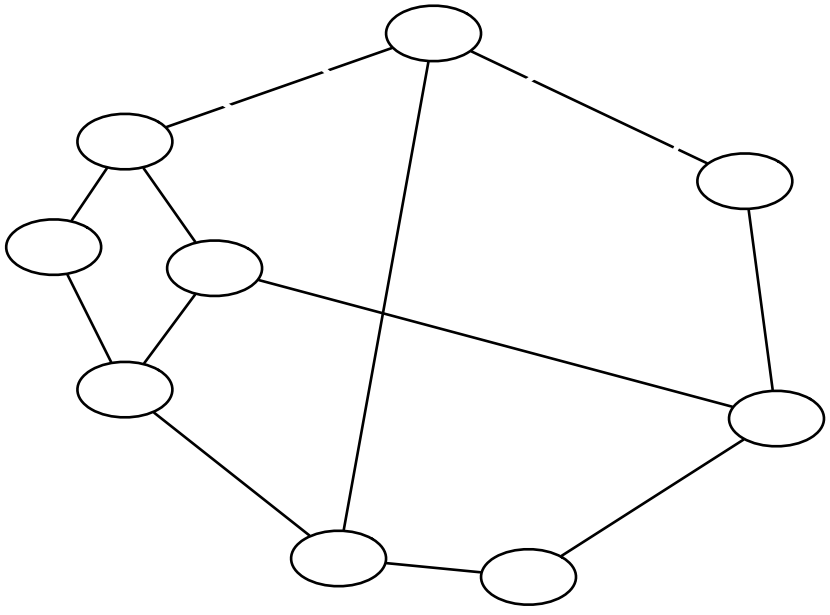

Graph H is a *minor* of graph G if
we can contract connected subsets of G into "super-nodes"
so as to produce a copy of H.

- In the example: G has a K_4-minor.

Characterizing Bad Instances via Graph Minors

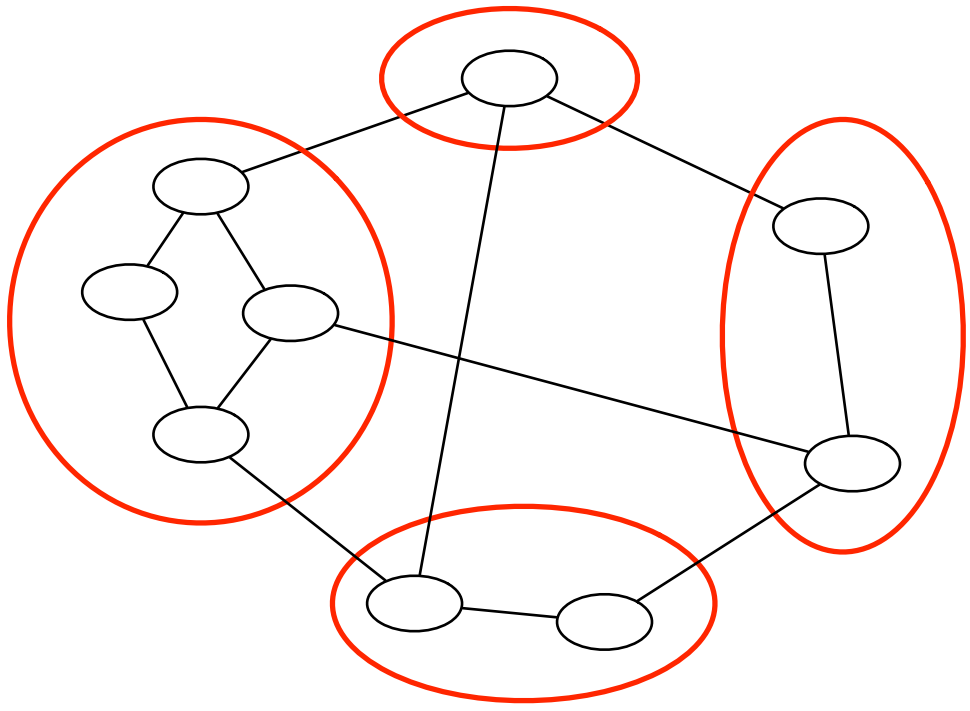

Graph H is a *minor* of graph G if
we can contract connected subsets of G into "super-nodes"
so as to produce a copy of H.

- In the example: G has a K_4-minor.

Characterizing Bad Instances via Graph Minors

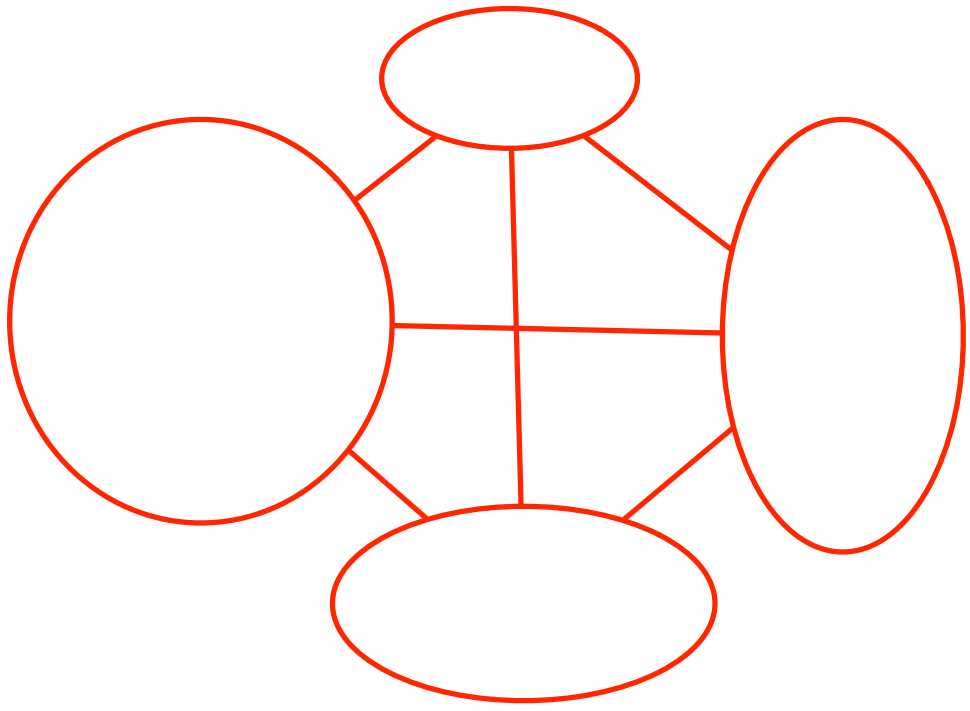

Graph H is a *minor* of graph G if
we can contract connected subsets of G into "super-nodes"
so as to produce a copy of H.

- In the example: G has a K_4-minor.

Characterizing Bad Instances via Graph Minors

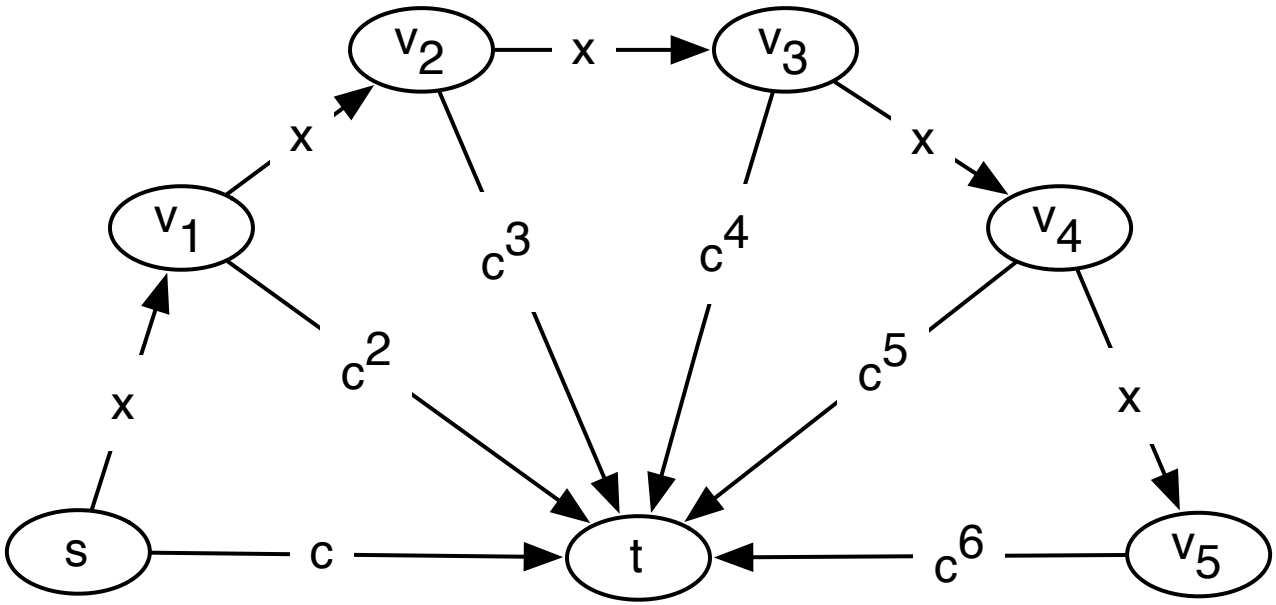

Characterizing Bad Instances via Graph Minors

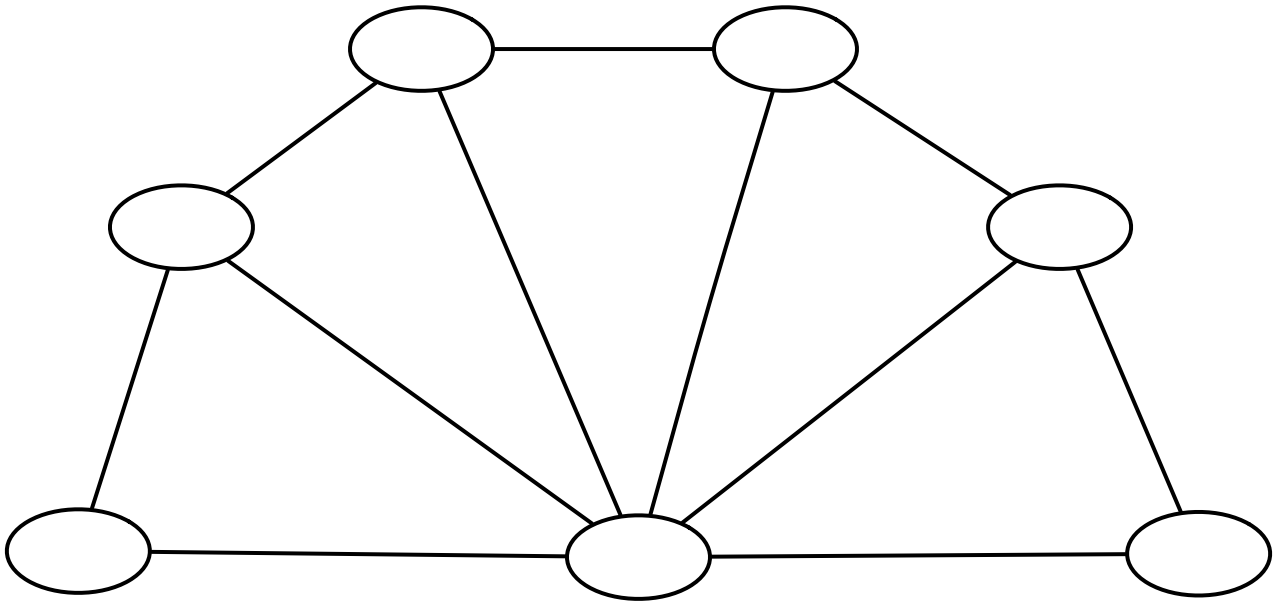

Characterizing Bad Instances via Graph Minors

The k-fan $\mathcal{F}_k$: the graph consisting of a k-node path, and one more node that all others link to.

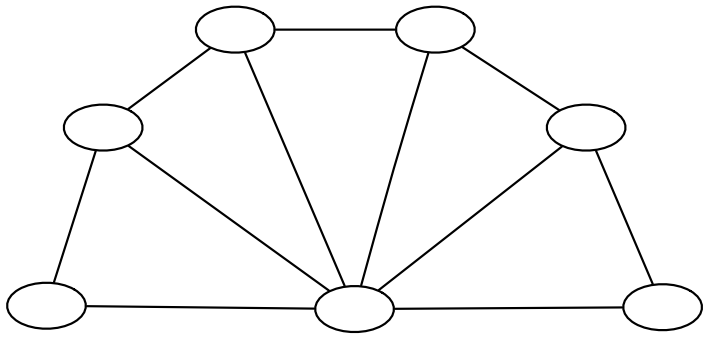

Theorem

For every $\lambda > 1$ there exists $\varepsilon > 0$ such that
if the cost ratio is $> \lambda^n$,
then the underlying undirected graph of the instance contains an $\mathcal{F}_k$-minor for $k = \varepsilon n$.

Choice Reduction

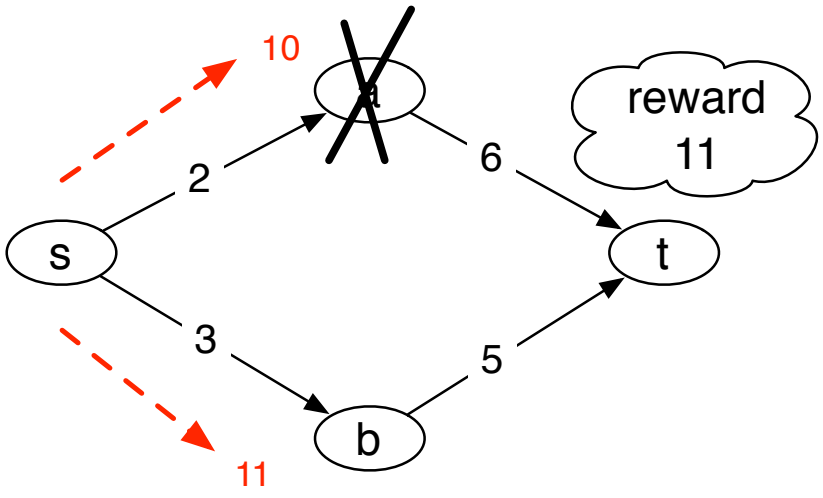

Choice reduction problem: Given G, not traversable by an agent, is there a subgraph of G that is traversable?

- Our initial idea: if there is a traversable subgraph in G, then there is a traversable subgraph that is a path.
- But this is not the case.

Results:

- A characterization of the structure of minimal traversable subgraphs.
- NP-completeness [Feige 2014, Tang et al 2015]

Choice Reduction

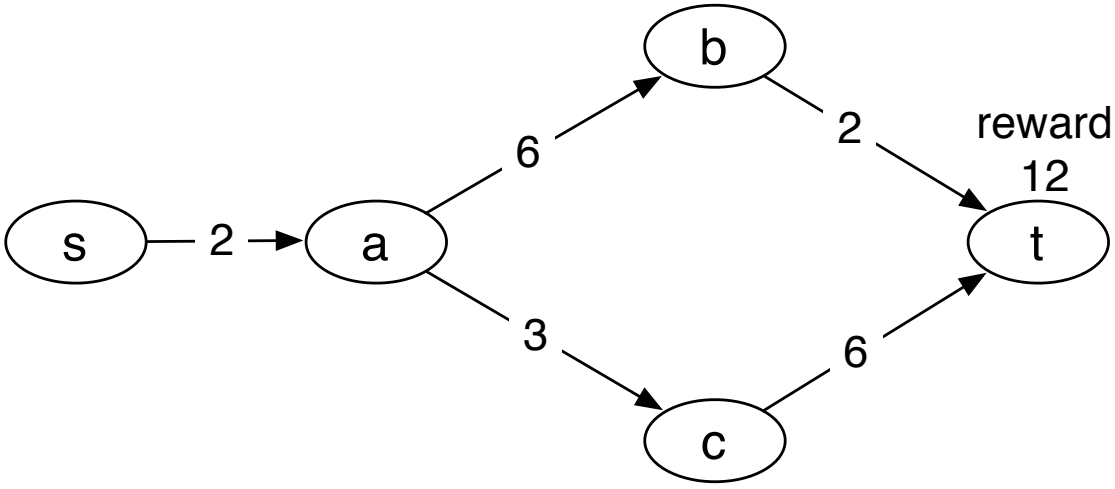

Choice reduction problem: Given G, not traversable by an agent, is there a subgraph of G that is traversable?

- Our initial idea: if there is a traversable subgraph in G, then there is a traversable subgraph that is a path.
- But this is not the case.

Results:

- A characterization of the structure of minimal traversable subgraphs.
- NP-completeness [Feige 2014, Tang et al 2015]

Sophistication

Sophisticated agents [O'Donoghue-Rabin 1999]

- Can successfully anticipate their behavior in the future.
- Plan in the present based on this awareness.

Example: It's Thursday; a progress report must be written and submitted by Saturday at midnight.

- Cost to do it Thursday $= 3$.
- Cost to do it Friday $= 5$.
- Cost to do it Saturday $= 9$.

A struggle between three *selves*: one for each of Thurs, Fri, Sat.

- On Saturday: must be done for cost of 9.
- Your Friday self perceives the cost as $2 \cdot 5 = 10 > 9$.
  Makes the Saturday self do it.
- Your Thursday self perceives the cost as $2 \cdot 3 = 6$.
  But doesn't want to leave the decision to the Friday self (since $6 < 9$).

Sophisticated Planning on a Graph

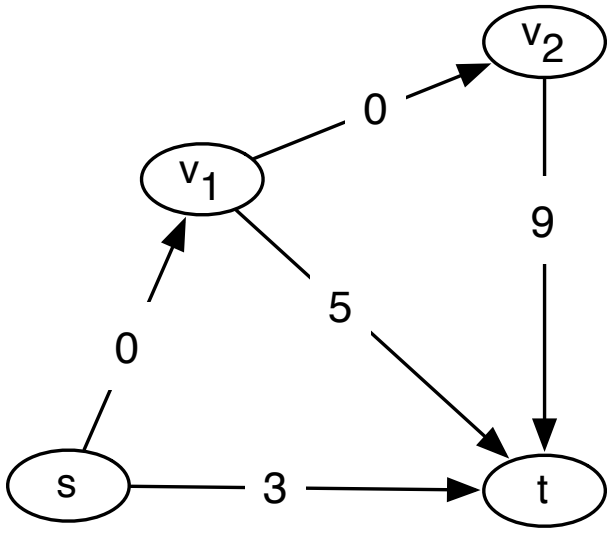

A graph-theoretic model of sophisticated planning
[Kleinberg-Oren-Raghavan 2016]

- There is a "self" for each node.
- Working backward in a topological ordering of the graph, determine what the self at node v will do, given known behaviors at later nodes.

Sophisticated Planning on a Graph

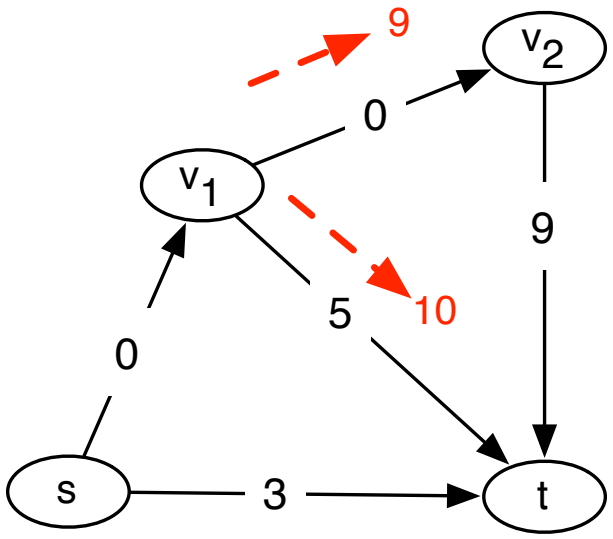

A graph-theoretic model of sophisticated planning
[Kleinberg-Oren-Raghavan 2016]

- There is a "self" for each node.
- Working backward in a topological ordering of the graph, determine what the self at node v will do, given known behaviors at later nodes.

Sophisticated Planning on a Graph

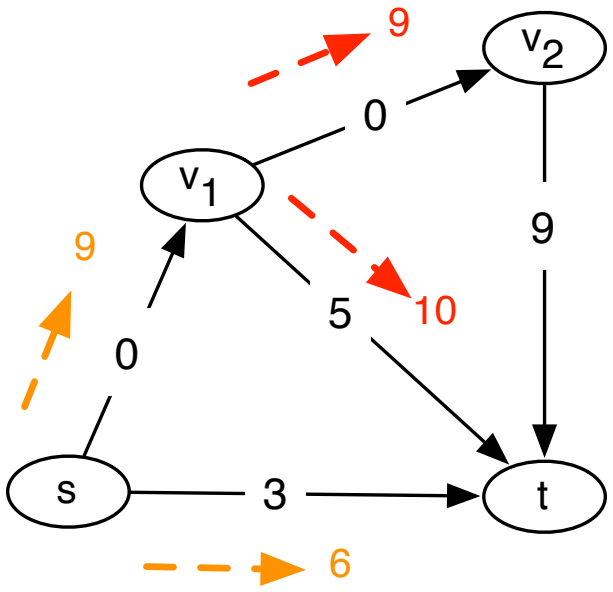

A graph-theoretic model of sophisticated planning
[Kleinberg-Oren-Raghavan 2016]

- There is a "self" for each node.
- Working backward in a topological ordering of the graph, determine what the self at node v will do, given known behaviors at later nodes.

Worst-Case Performance for Sophisticated Agents

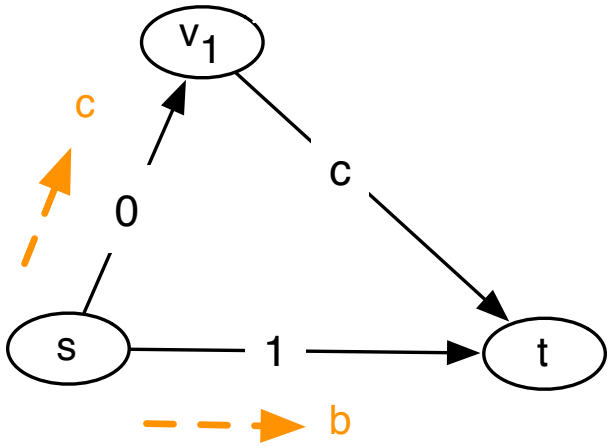

Sophisticated agent can be c times worse than optimal, for any $c \leq b$.

Worst-Case Performance for Sophisticated Agents

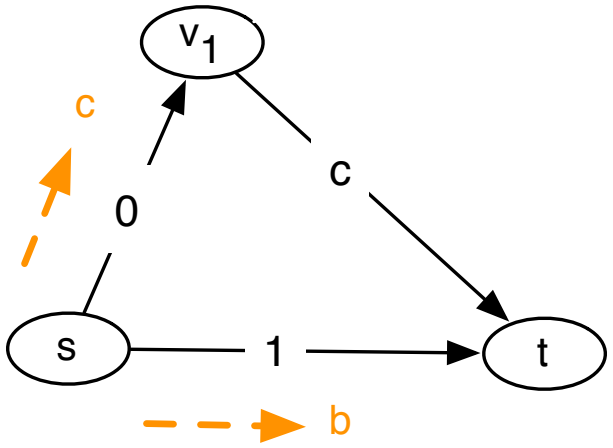

Sophisticated agent can be c times worse than optimal,
for any $c \leq b$.

Theorem [Kleinberg-Oren-Raghavan 2016]: In every instance G, a sophisticated agent incurs at most b times the optimal cost.

- Worst case is exponentially better than in the case of naive agents.

Further Directions

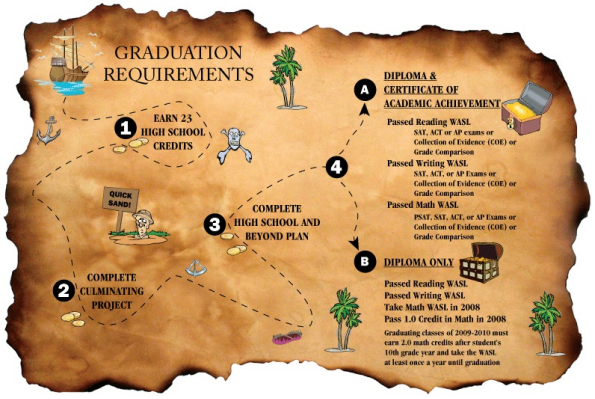

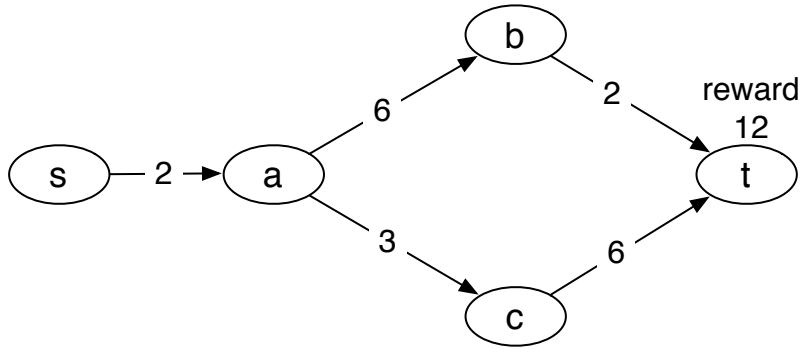

- Reasoning about long-range planning requires a model for decisions.
- Graph-theoretic framework for present bias uncovers new questions and new phenomena.
- Can study the interaction of multiple biases: present bias and sunk-cost bias [Kleinberg-Oren-Raghavan 2017].
- Connecting these ideas back to incentive design.