A tutorial on: Dynamic Mechanism Design

Ruggiero Cavallo
University of Pennsylvania
Department of Computer and Information Science

July 7, 2009
ACM EC

The setting

- <u>Sequence</u> of decisions to be made impacting utility experienced by a group of agents
- Social planner wants to make optimal choices
  - agents hold private valuation information
  - knowledge of private information required to determine optimal decision at every point in time
  - new private information potentially arrives after each decision

Example: a resource – say a government-owned super-computer – is to be allocated *repeatedly* for 1 week intervals.

Problem: Agents' goals differ from center's goals, but agent cooperation is essential.

The solution framework

- **Dynamic mechanism design**: specification of payment schemes such that optimal outcomes are achieved *in equilibrium*.
  - An extension/generalization of "static" mechanism design.

Does static MD really fall short?

- **Yes**. Rare is the decision scenario that is completely independent of future decisions.
- E.g., allocating a resource. What future opportunities will there be for procuring the resource? What opportunities for reselling the resource?

Tutorial plan

1. Rudimentary review of static mechanism design.
2. Dynamic MD basics: modeling of “types” in dynamic settings, dynamic equilibrium notions.
3. **Key solutions so far.**
4. Extensions.

Static mechanism design

Just social-welfare maximizing, here

- Many of the marquee static MD results have (more complicated) analogs in the dynamic setting.
- So start with quick review of static case...

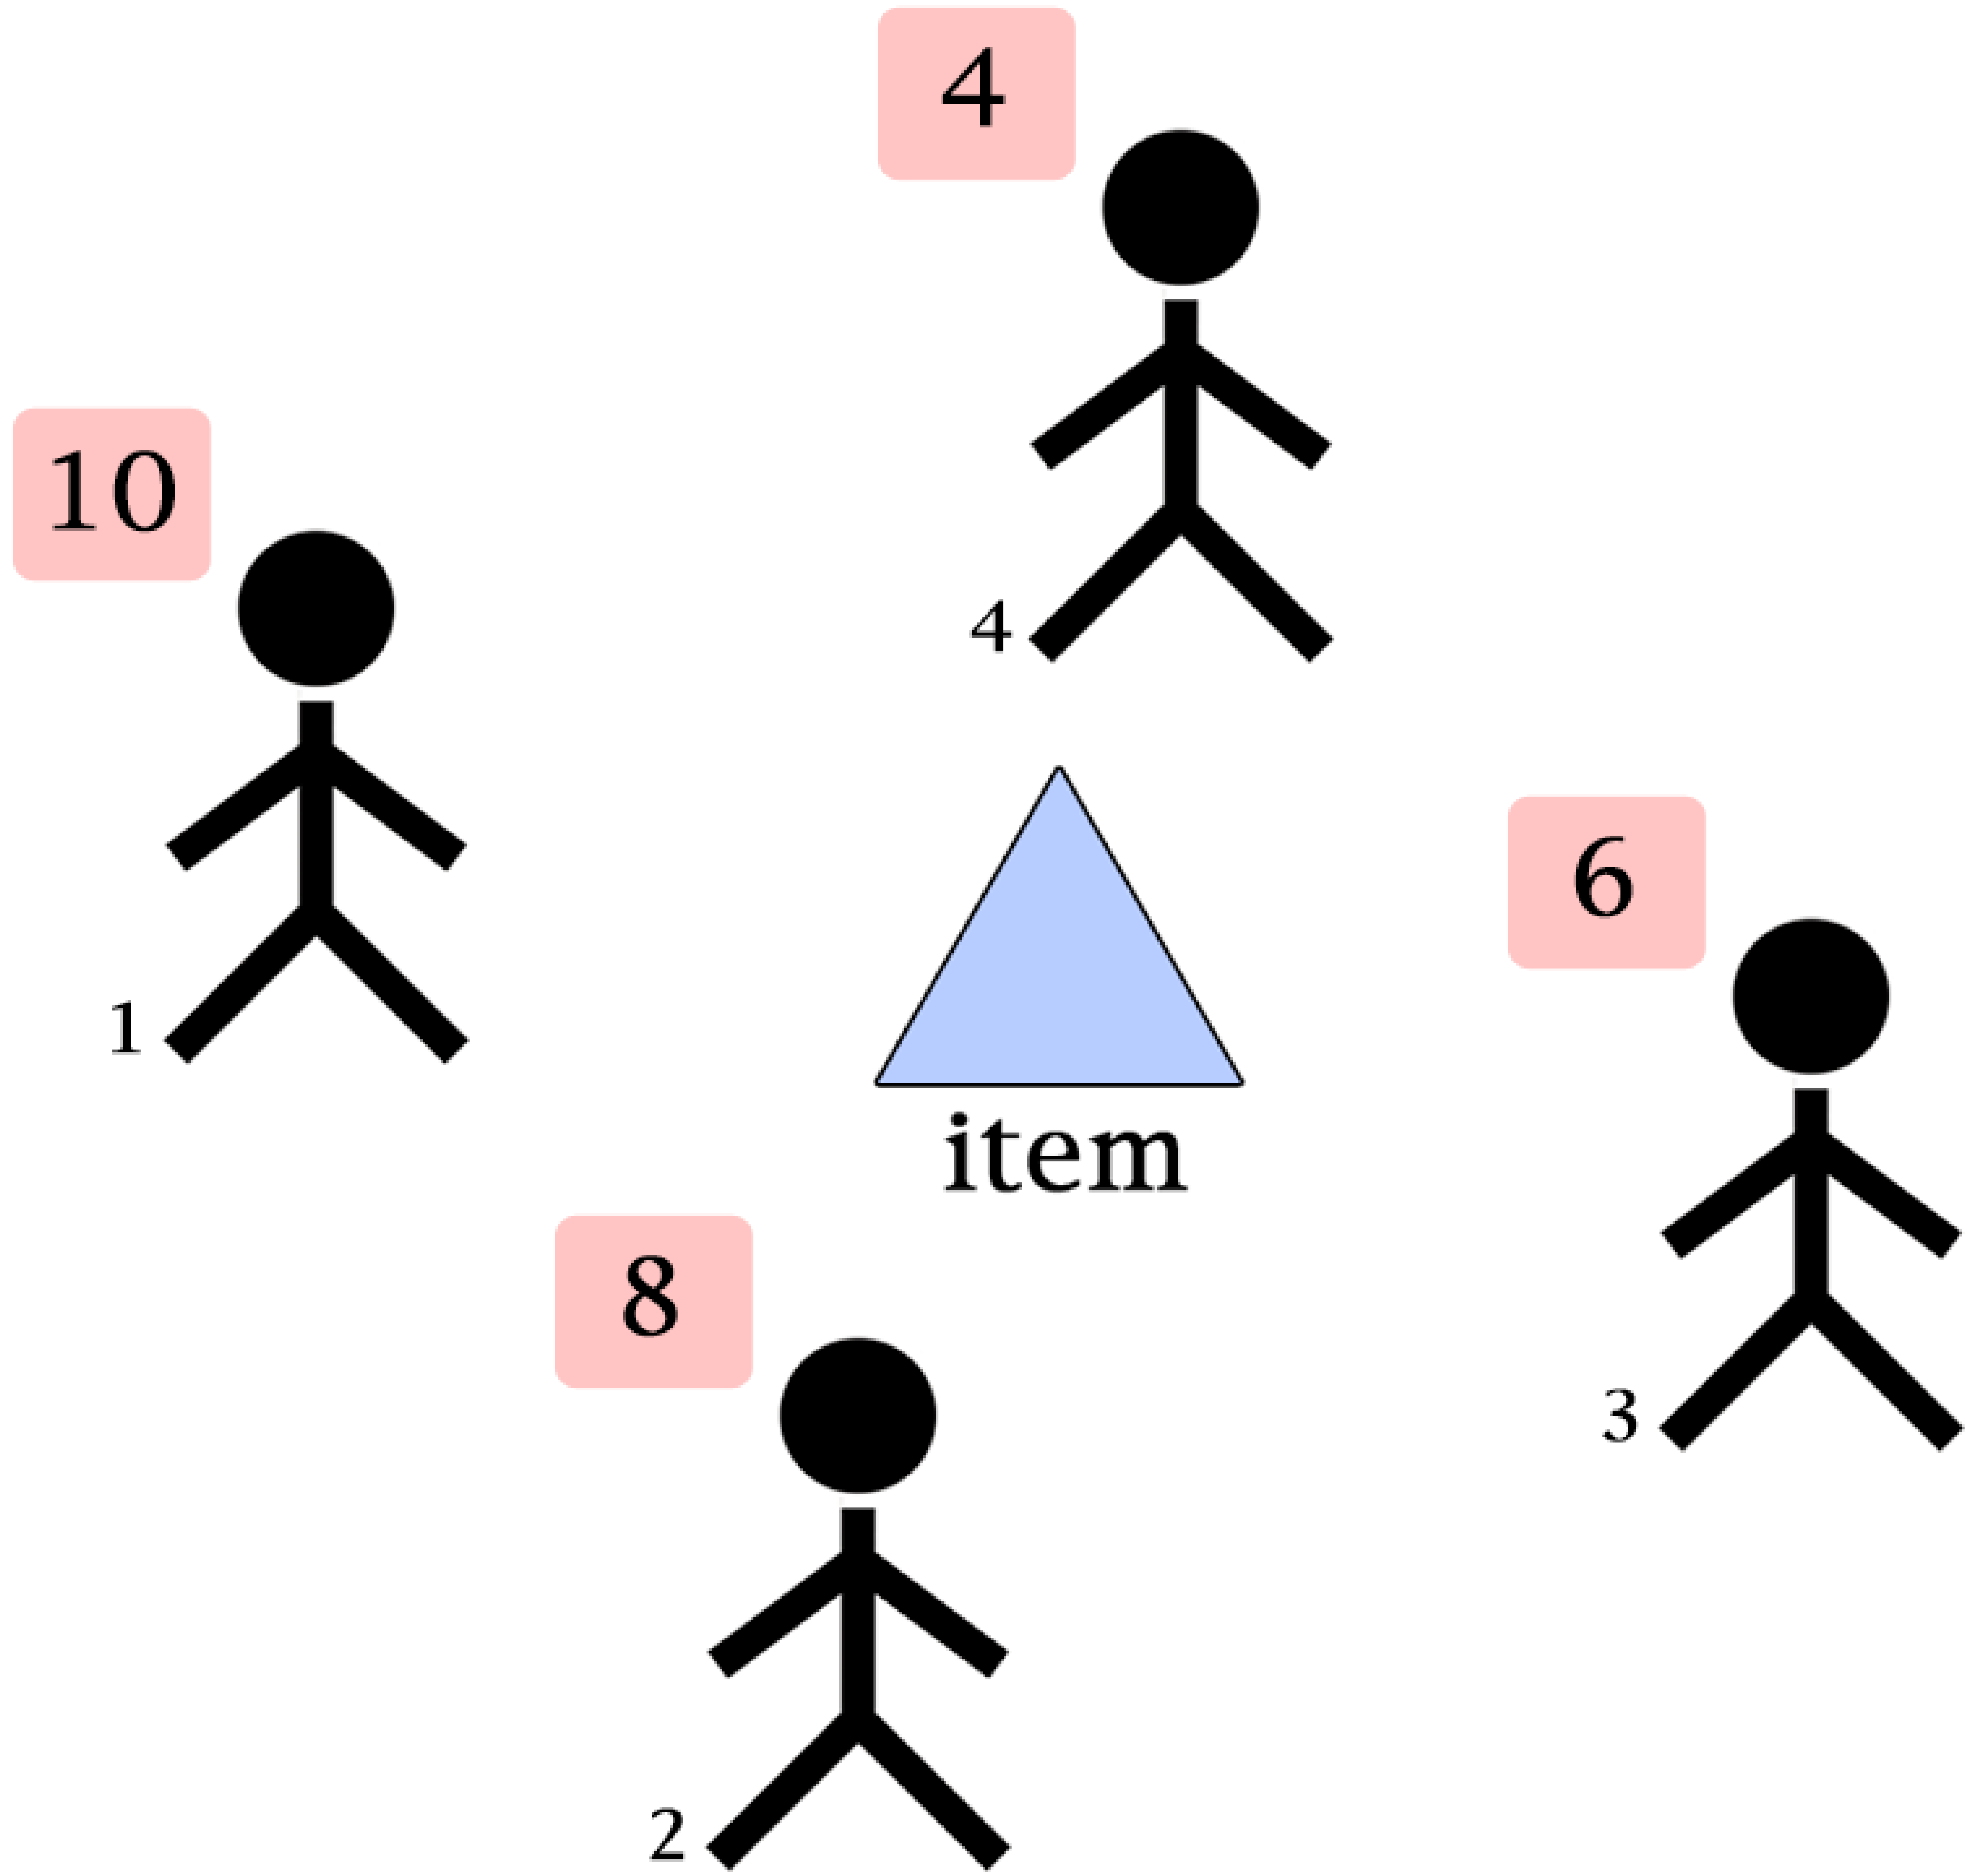
4
10
6
8
4
1
3
2
item

Mechanism Design

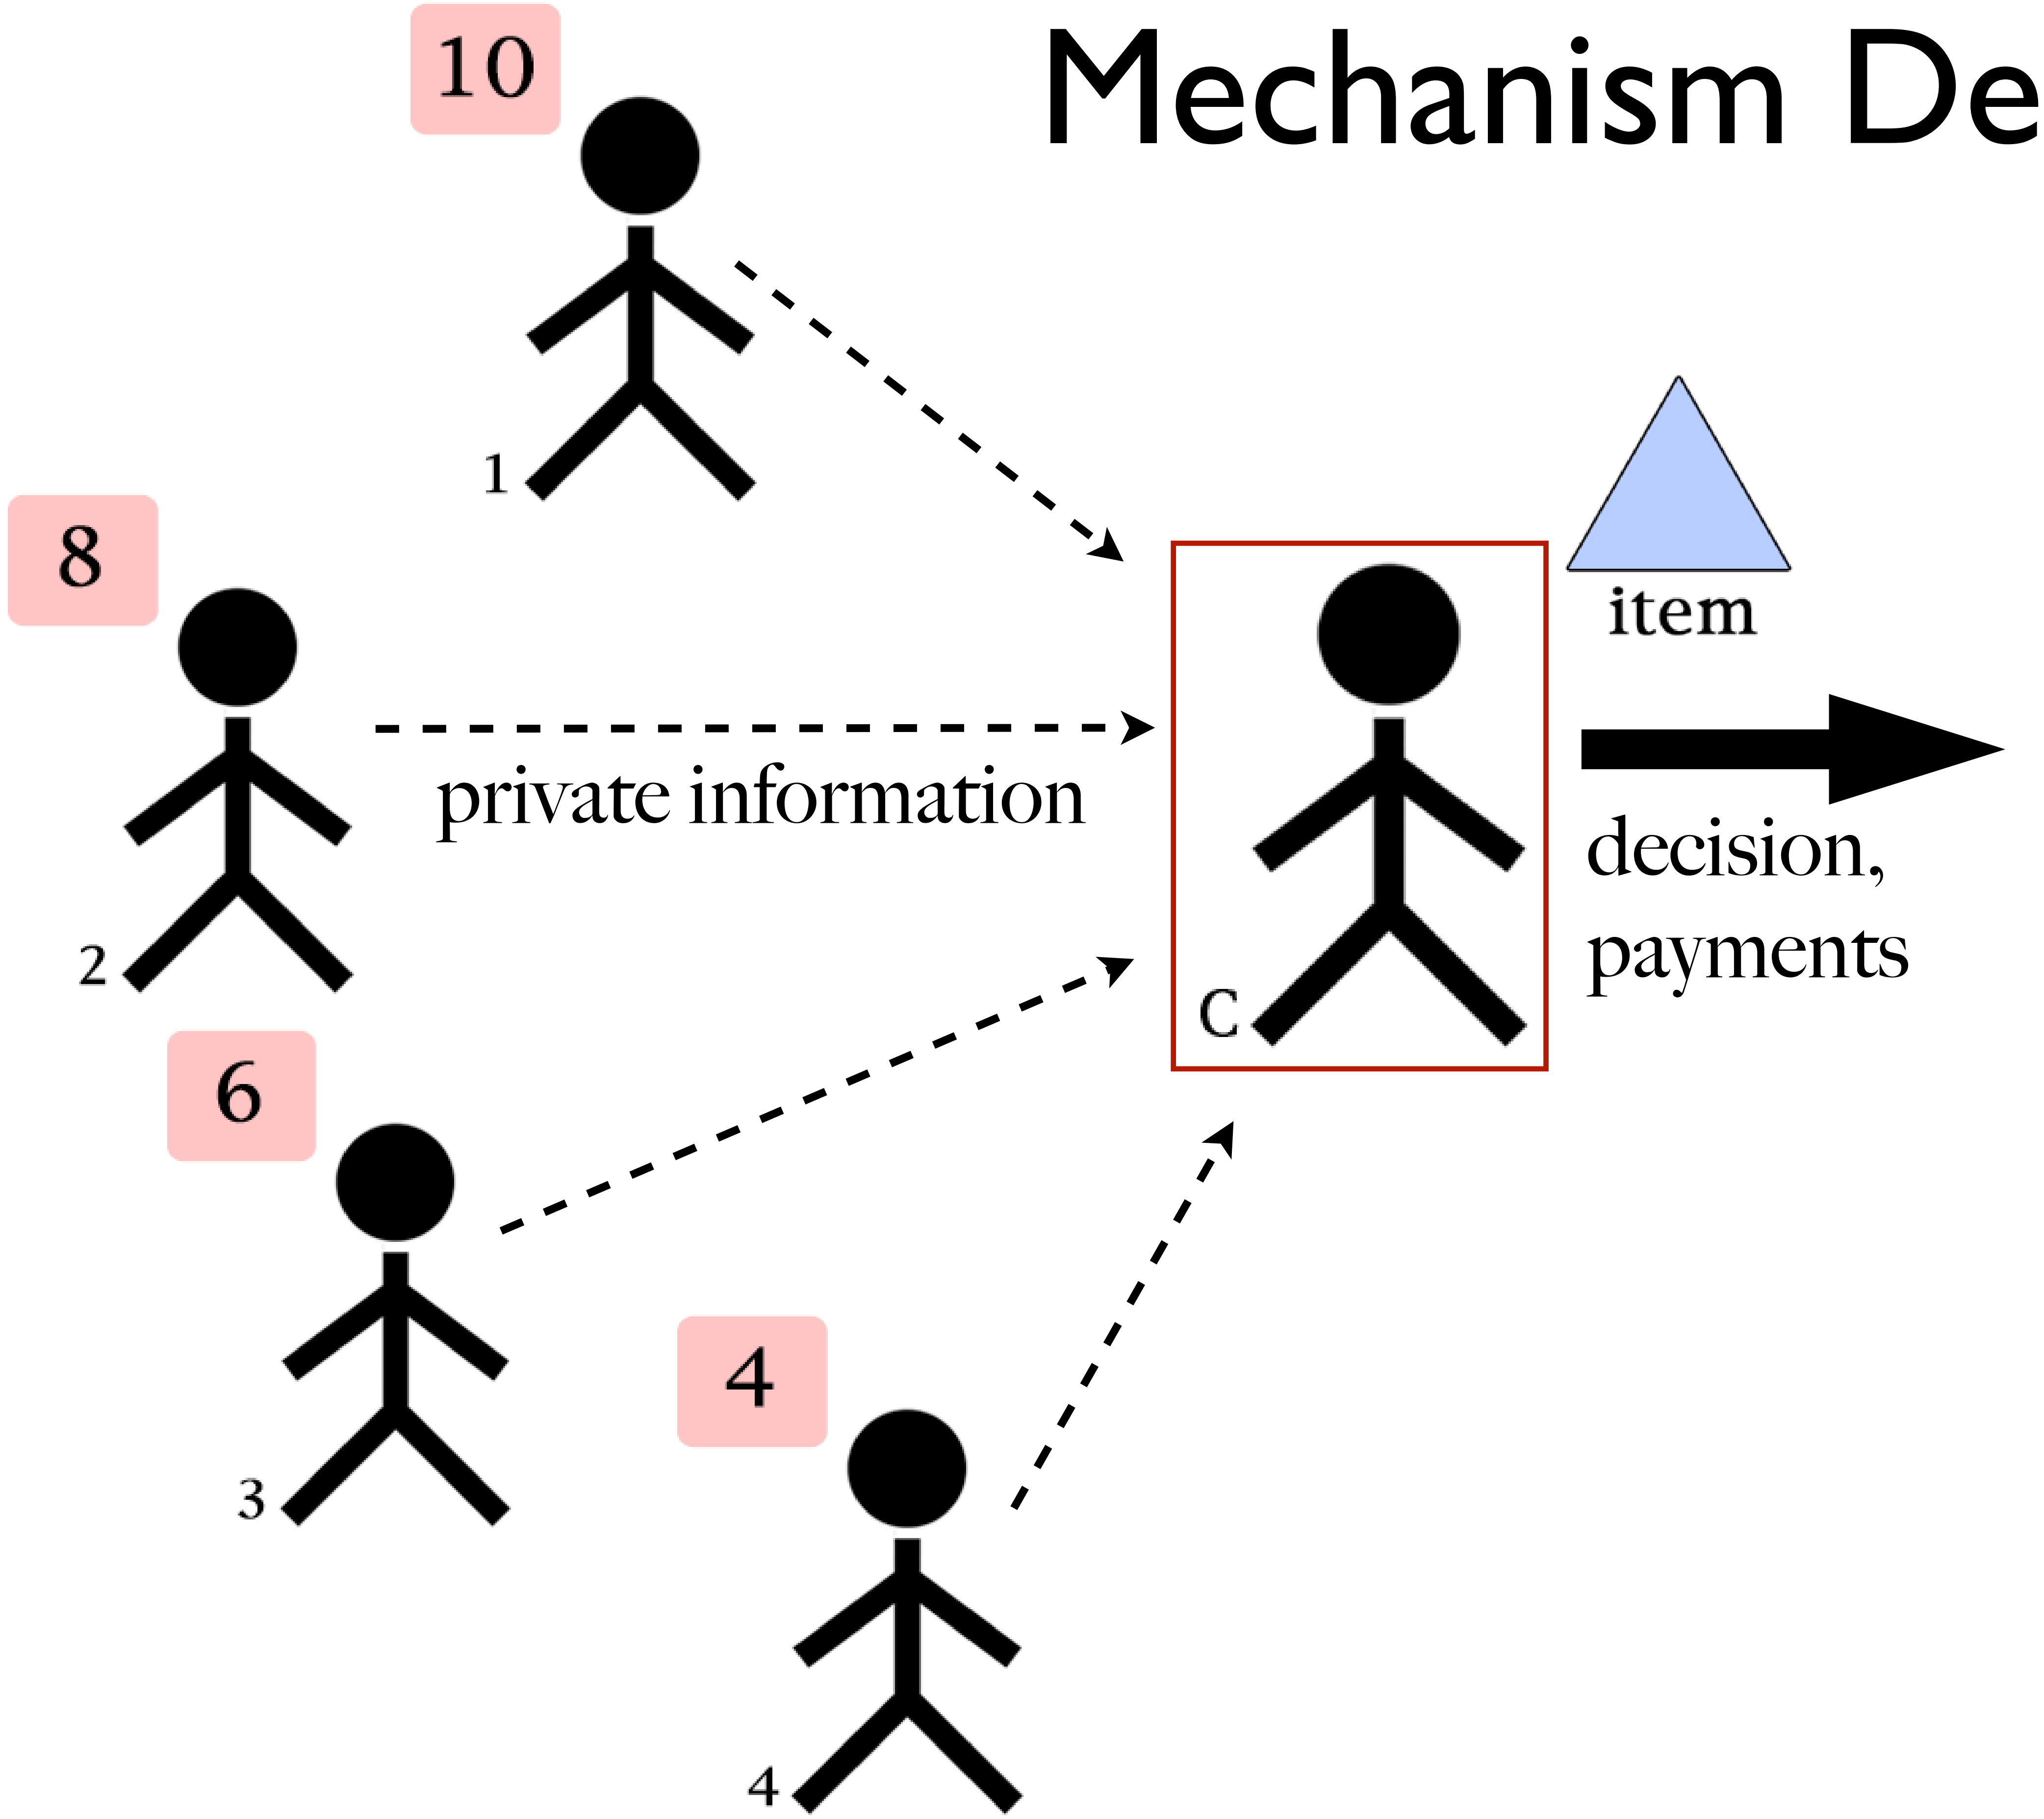

Mechanism design

- Specify decision rule (outcome selection), plus a monetary charge/payment imposed on each agent.
- Outcome/payments enforced by a *center*.
- Criteria for success:
  - **social-welfare maximizing** (a.k.a., *efficient*)
  - individual rational (no agent worse off)
  - budget properties

Solution concepts

- <u>Strategyproof</u>: reporting true type is always a utility-maximizing strategy, regardless of what other agents do.
- <u>Ex post incentive compatible</u>: reporting true type is utility-maximizing, whatever the types of other agents, assuming they're truthful. [same as strategyproof in private values setting]
- <u>Bayes-Nash incentive compatible</u>: reporting true type is utility-maximizing, in expectation given distribution over others' types, assuming other agents are truthful.

Efficient static mechanisms: the Groves class

- Choose outcome that is social welfare maximizing according to agent reports.
- Pay each agent the (combined) reported value of all other agents for the chosen outcome... minus some quantity independent of the agent's report.

[Vickrey, 1961; Clarke, 1971; Groves 1973]

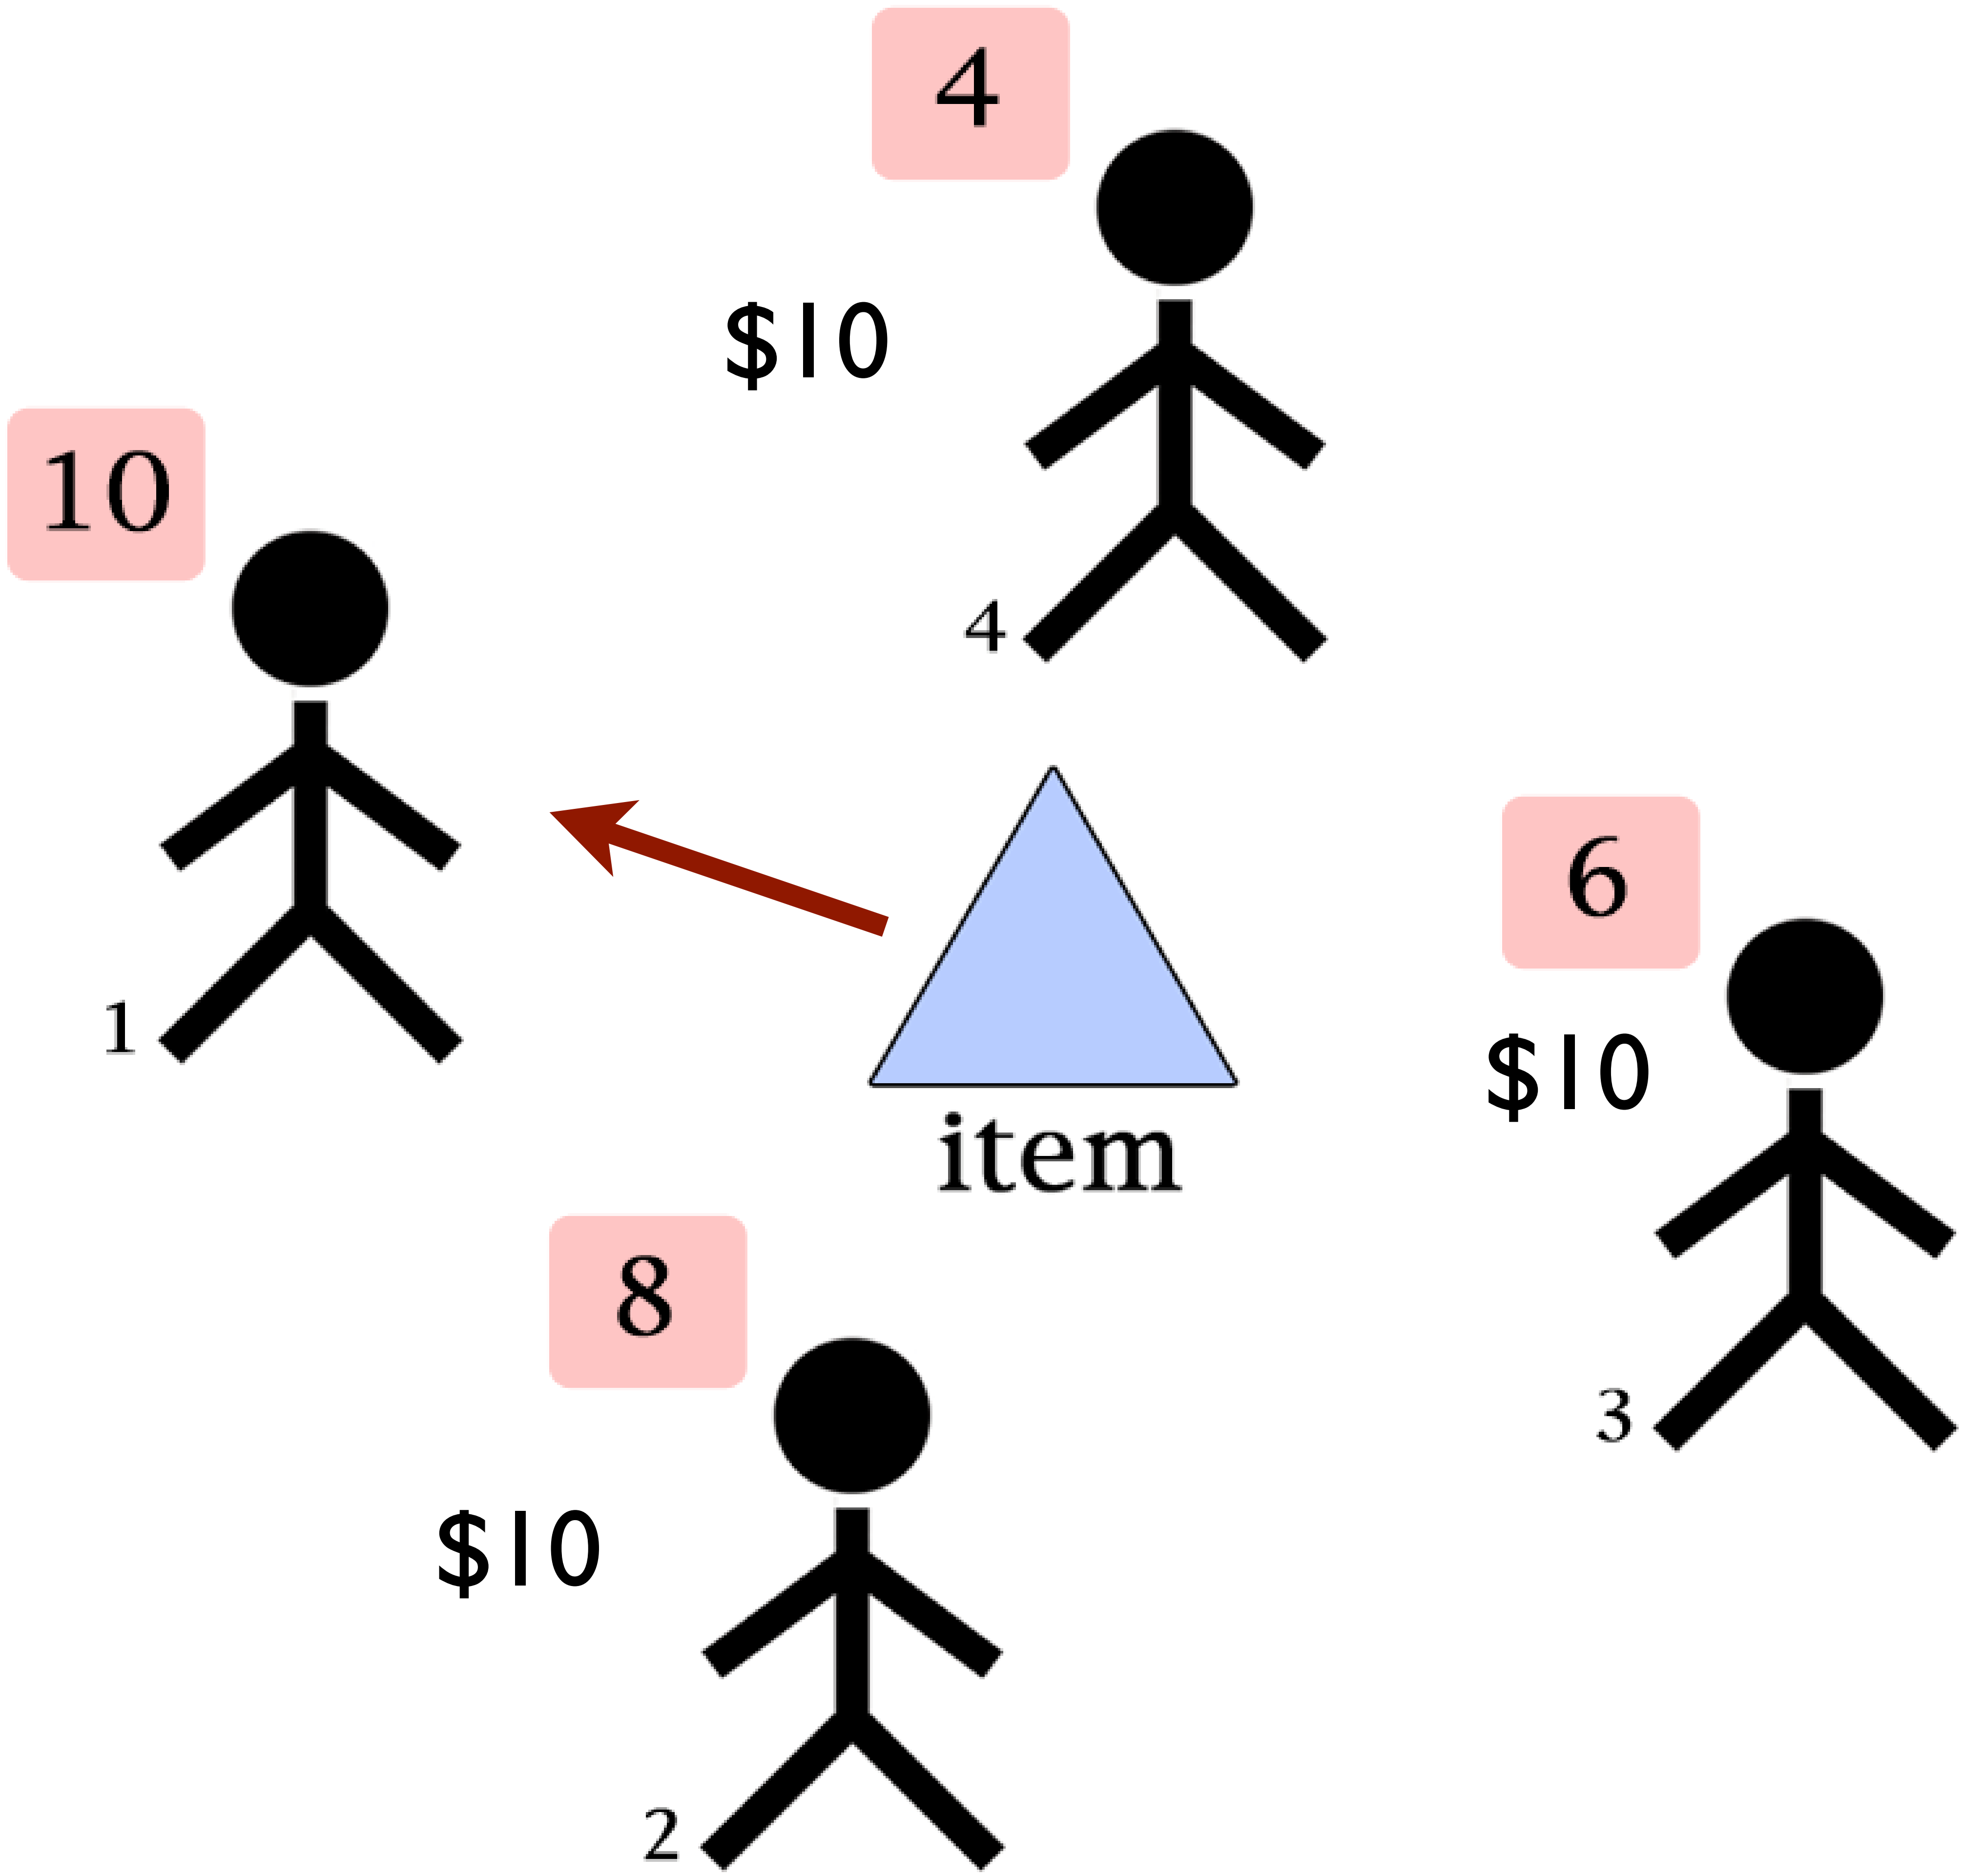
4
$10
4
10
1
6
$10
3
item
8
$10
2

Groves (and nothing else) works

- The set of Groves mechanisms exactly corresponds to those that are efficient in dominant strategies.*

[Green & Laffont, 77], strengthened by [Holmstrom, 79]

➡ Our freedom is limited to defining the agent-independent "charge" term

*For sufficiently rich domains ("for all practical purposes").

Efficient mechanism design boiled down

1. **Align incentives** – make each agent's payoff equal to social welfare.
2. **Recover funds** – have each agent make a payment independent of his behavior.

Efficient mechanism design boiled down

1. **Align incentives** – make each agent's payoff equal to social welfare.
2. **Recover funds** – have each agent make a payment independent of his behavior.

A little more subtle in dynamic setting...

The VCG mechanism

- A Groves mechanism.
- Defines "charge" term for each agent i equal to the value other agents *could have* obtained if i's interests were ignored.
  - Each agent's utility equals contribution to social welfare.
  - Ex post individual rational (if agents have non-negative values for all outcomes)
  - No-deficit... in fact often yields high *revenue*.

The expected externality (AGV) mechanism

[Arrow, 79; d'Aspremont & Gerard-Varet, 79]

- Each agent's payment is expected social welfare others will get given his report, minus some uninfluencable quantity.
- Efficient in Bayes-Nash equilibrium.
- Ex ante individual rational.
- **Strongly budget-balanced**.

Redistribution mechanisms

- AGV maintains *all* value with agents, but is only weakly efficient and IR, unlike VCG.
- Idea: try to return revenue under VCG back to agents – thus improving social welfare – without weakening equilibrium or running deficit.
- First reference: [Bailey, 97] for certain allocation settings.

Redistribution mechanism [Cavallo, 06]

Idea: leverage domain information to obtain "revenue-guarantees".

- For each agent i, compute minimum revenue that i could cause to result, given reports of other agents (G_i).
- Run VCG.
- Give each i payment of G_i/n.

Applicable to any setting (e.g., *combinatorial* allocation). In single-item allocation, coincides with [Bailey, 97] mechanism.

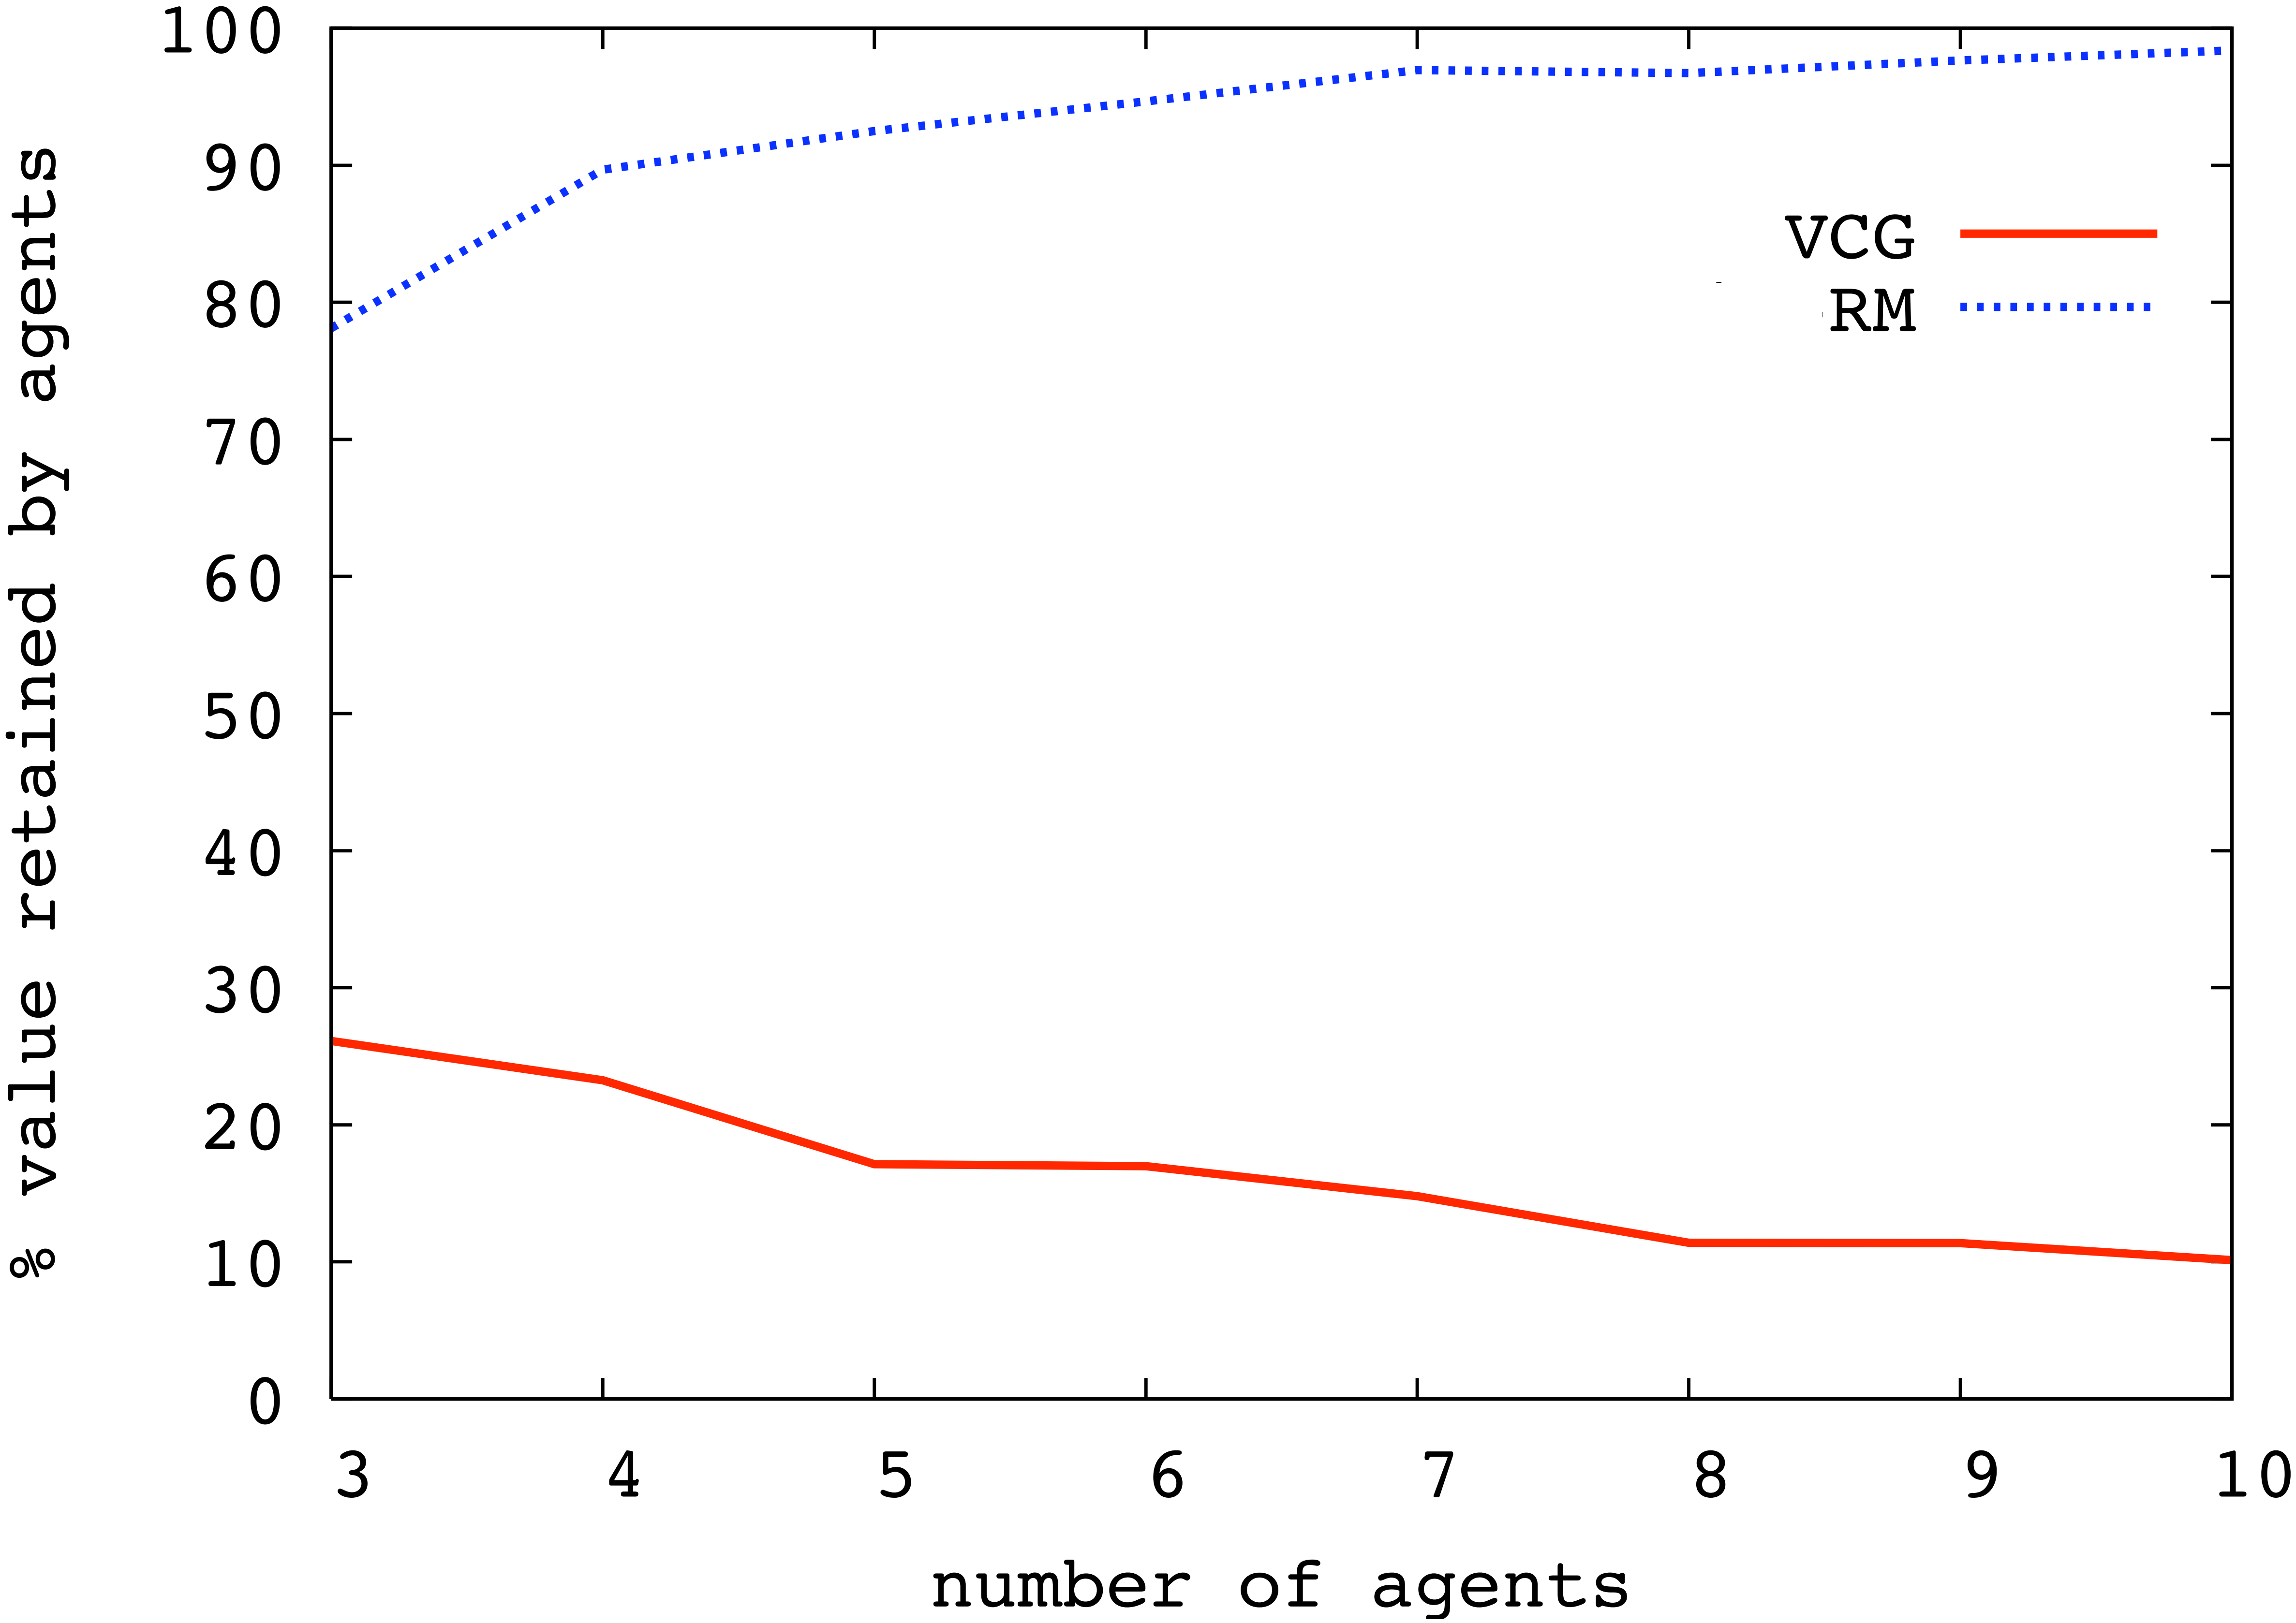

100
90
80
70
60
50
40
30
20
10
0
3
4
5
6
7
8
9
10
% value retained by agents
number of agents
VCG
RM

Lots of interesting recent work for case of multi-unit auctions

- [Guo & Conitzer, 07; Moulin, 2007] – worst-case optimality.
- [Guo & Conitzer, 08] – optimal-in-expectation mechanism.
- [Hartline & Roughgarden, 2008] – money burning when payments not possible.
- [de Clippel, Naroditskiy, Greenwald, here].

Much not discussed here, e.g.,

- Interdependent (“common”) values settings
- Inefficient mechanism design, concerned with, e.g., revenue maximization, maximizing the minimum utility, etc.

Basics of the dynamic setting

Aspects of the problem

- At each time period each agent holds some private information ("local state").
- At each time period, the center selects an action to execute, which generates value (of varying degree) for agents and yields new local states.
- The (predicted) effects of taking any given action depend on state.
- Agents perceive utility of value x obtained k steps in future to be $\gamma^k x$, for some $0 < \gamma \leq 1$.

➡ Key variable: local state.

Local state

- Encapsulates all information required to determine:
  - a conditional distribution over value the agent would obtain for every possible action
  - a conditional distribution over future local states for every possible action

Assumption

- Given the action executed by the center, value obtained and subsequent local state for each agent are independent of other agents' local states.

➡ Dynamic version of private values.

Markov decision processes (MDPs)

- State space
- Action space
- Reward function
  - When a given action is taken in a given state, what value results?
- Non-deterministic transition function
  - When a given action is taken in a given state, what new state results?

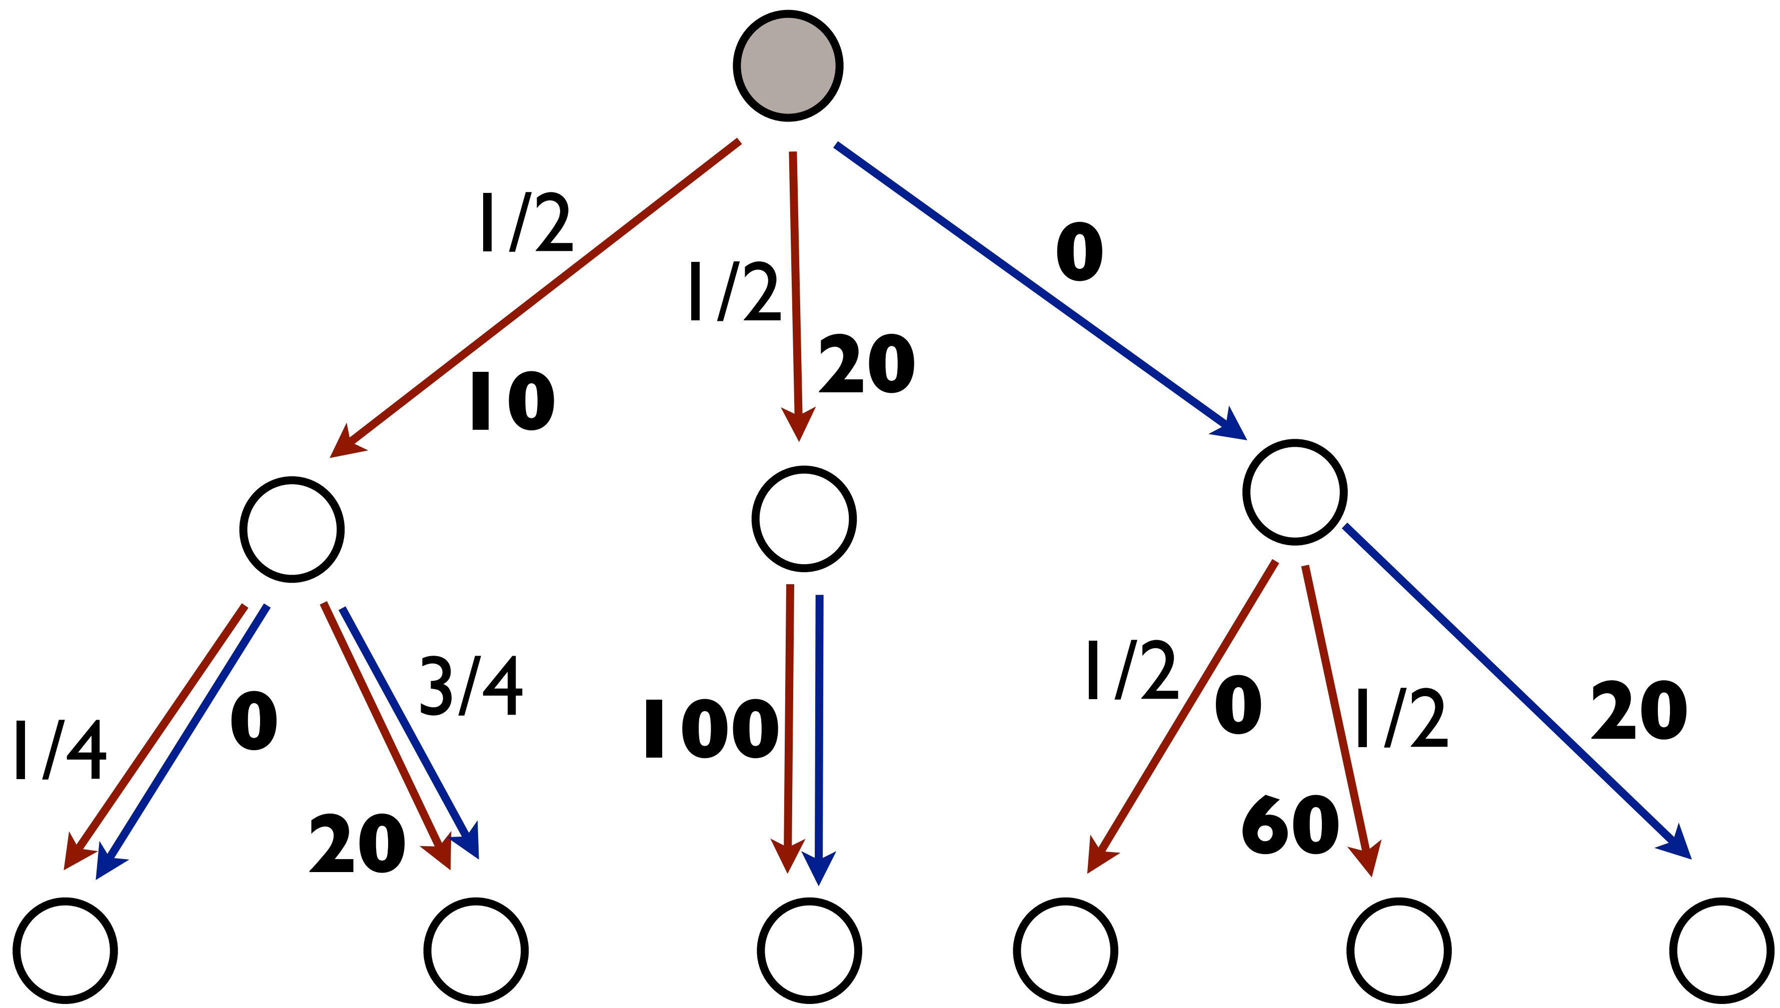

- Two possible actions (red and blue).
- Two time periods.

So what’s a dynamic type?

- It’s an MDP.
  - Transition dynamics between local states.
  - Value function for state-action pairs.
  - Indicator of “current” state.

- In static setting, type is "complete" and reportable. In dynamic setting, type is gradually revealed *to the agent* by nature over time.
- It's not the multiple time steps alone, it's the *uncertainty*.
- If types are MDPs with no stochastic state transitions, we're in a static MD setting – just decide policy at time 0.

A simple (and important) special case: MABs

- **Multi-armed bandit** problems: a special case of the general sequential decision-making framework.
- Captures, e.g., single-item repeated allocation scenarios.

A simple (and important) special case: MABs

- Each agent's dynamics can be represented by a Markov chain: no multiplicity of actions.
- A single action associated with each agent. When an agent's action is chosen, his state changes; otherwise, it doesn't.

Captures, e.g., repeated allocation of a resource

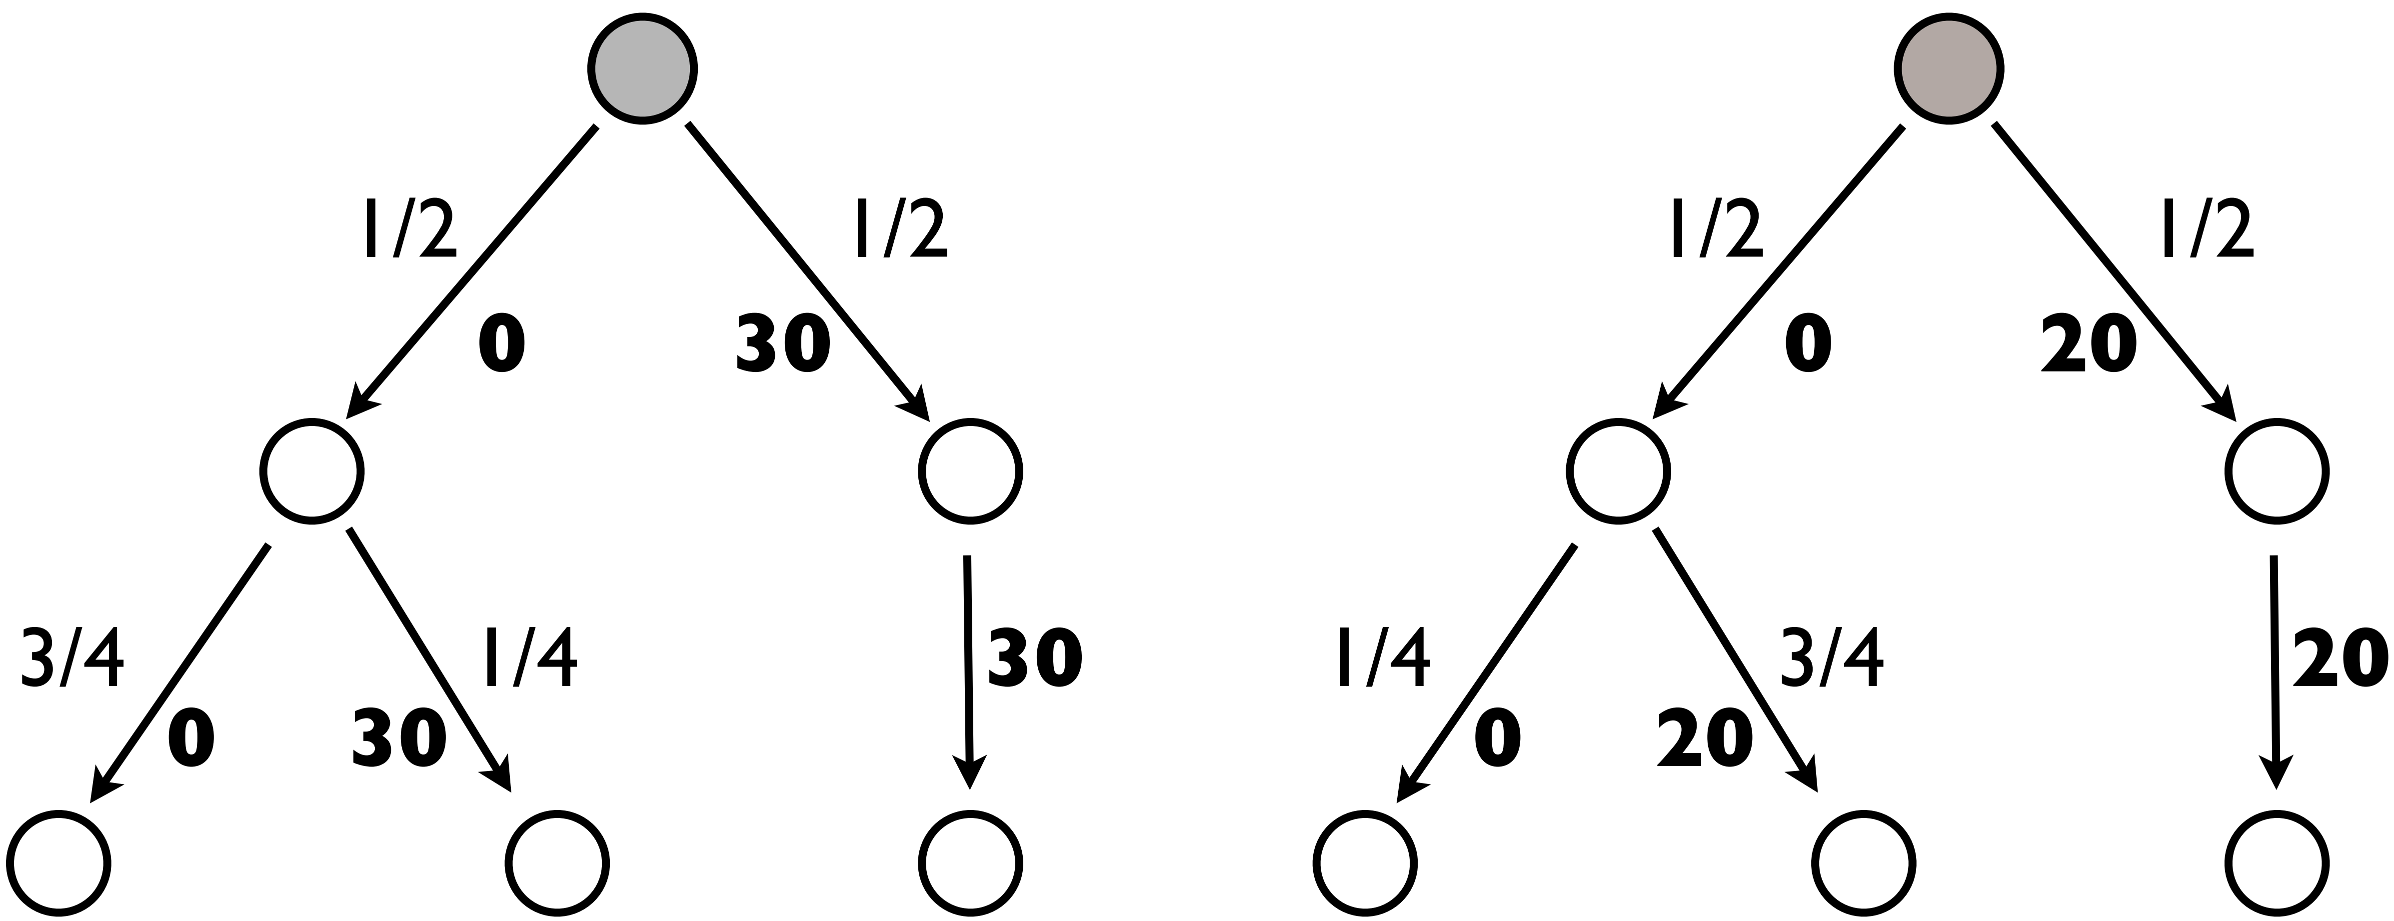

Captures, e.g., repeated allocation of a resource

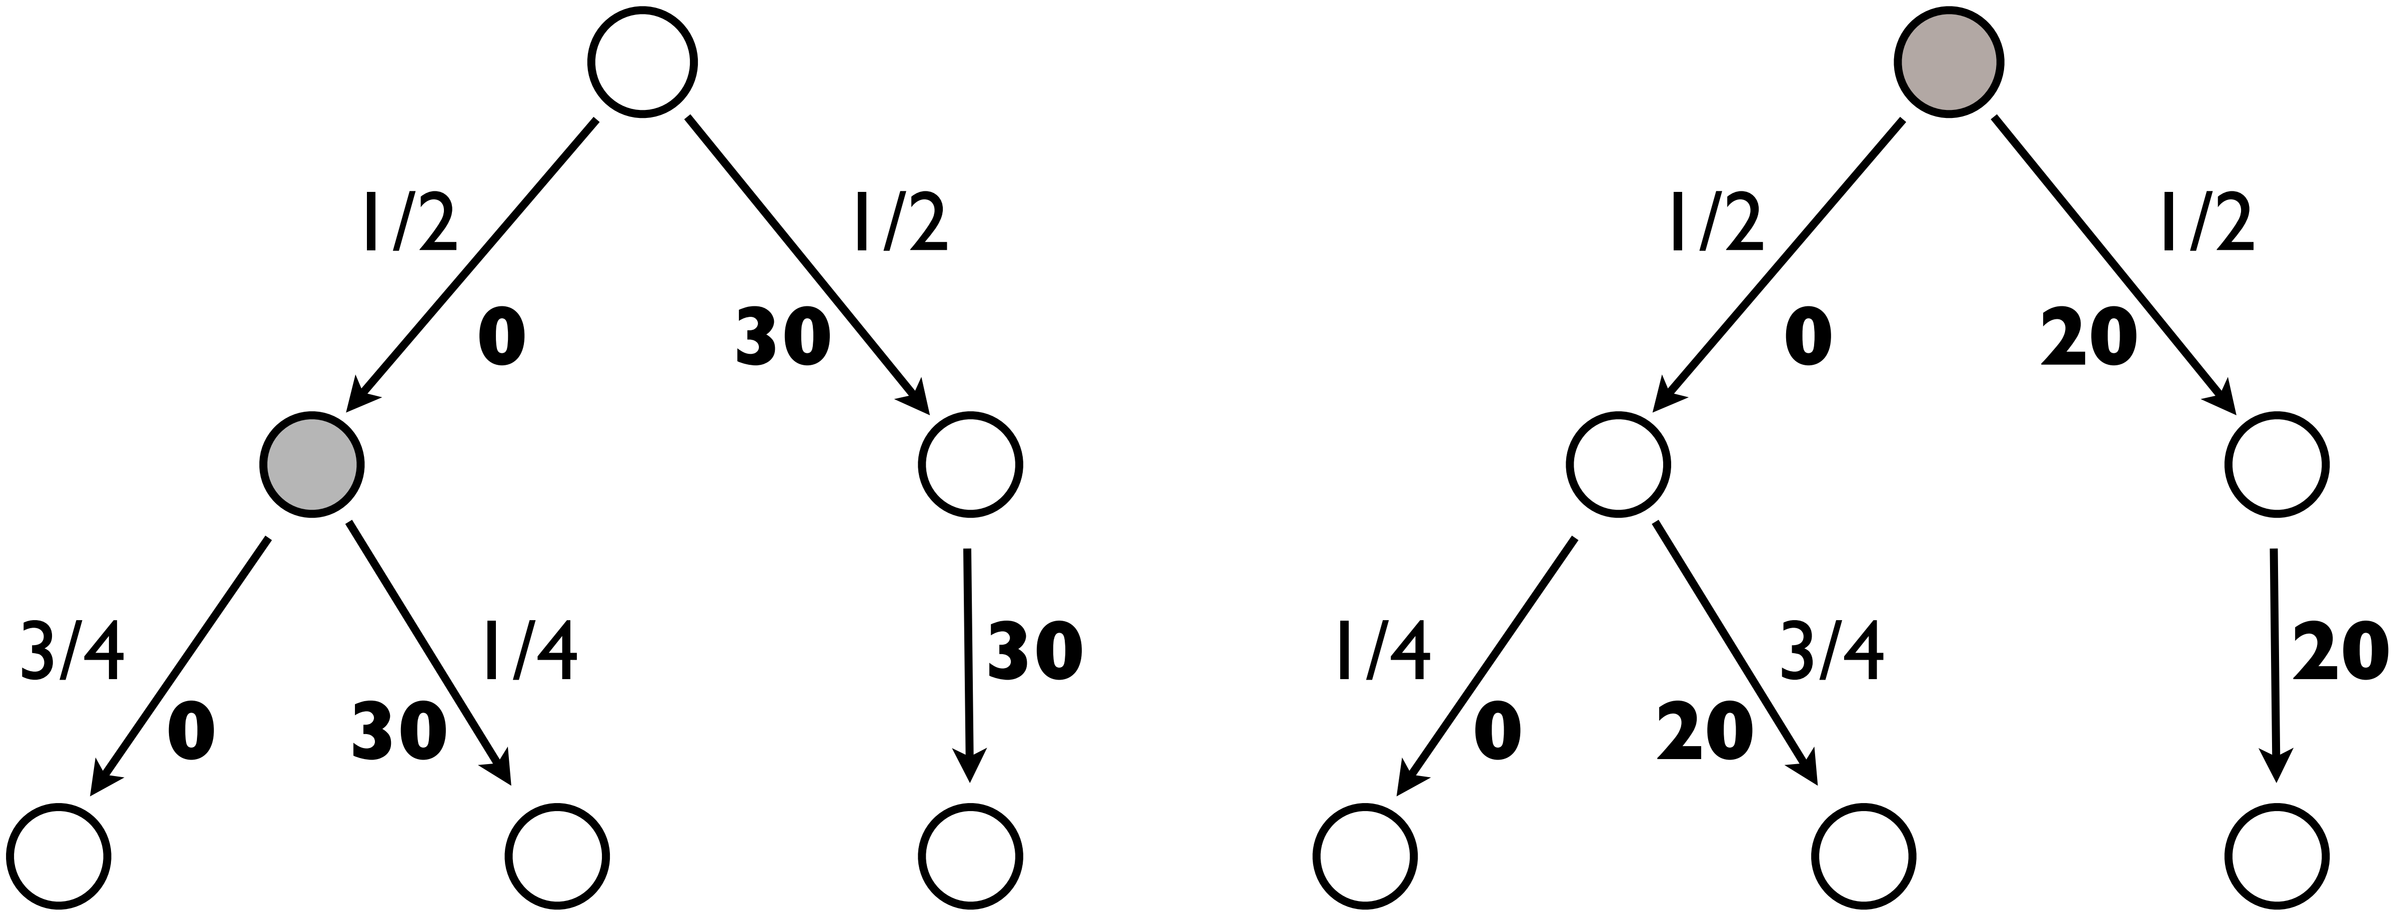

Allocate to agent 1, who finds no value.

DMD setup

- There is a set of actions.
- Each agent has a type represented by an MDP.
- In each period agents report types and the center takes an action.

- A dynamic mechanism specifies two things:
  - a **decision policy:** a function that maps a joint type to an action.
  - a **transfer function**: that maps a joint type to a payment for each agent.

Basics of the dynamic setting: equilibrium concepts

Within-period ex post Nash equilibrium

If all other agents play the equilibrium strategy in the future, no agent can benefit from deviating – regardless of what the joint state is and regardless of what came before.

Within-period ex post incentive compatibility

If all other agents report types truthfully in the future, no agent can benefit from misreporting type – regardless of what the joint type is and regardless of what came before.

No incentive to deviate even if agents know everything one *can* know – without being able to see the future.

This is the gold standard

- In a dynamic setting, agents needs to make predictions about the future in determining how to maximize utility – and this requires positing some behavior for other agents.
- Weaker than dominant strategy.
- But if others' future types were irrelevant to the agent's utility, incentives couldn't possibly be aligned.

Bayes-Nash equilibrium

Given distribution over other agents' types, no agent can expect to gain from deviating if others don't.

- Within-period ex post also involves expectation, but expectation is over uncertain type transitions, *not* current types.

Mechanism desiderata

- Efficiency: social-welfare maximizing decisions achieved in equilibrium.
- Individual rationality: no agent expects to lose from participating.
  - Within-period ex post: at *every* time-step, for *every* joint type.
  - Ex ante: from beginning of the mechanism, for whatever the joint type is then.
- Budget-balance / no-deficit.

By the way...

- A dynamic analog of the revelation principle holds [Myerson, 1986].
- So we can think only about direct revelation mechanisms, without loss of generality.

Some solutions so far

A basic efficient dynamic mechanism

- Dynamic team mechanism
  [Athey & Segal, 07]
  - Follows efficient policy given agent reports.
  - In each period, pays each agent the expected *immediate* value obtained by other agents given reported types ("Groves payment").

Dynamic team mechanism
example

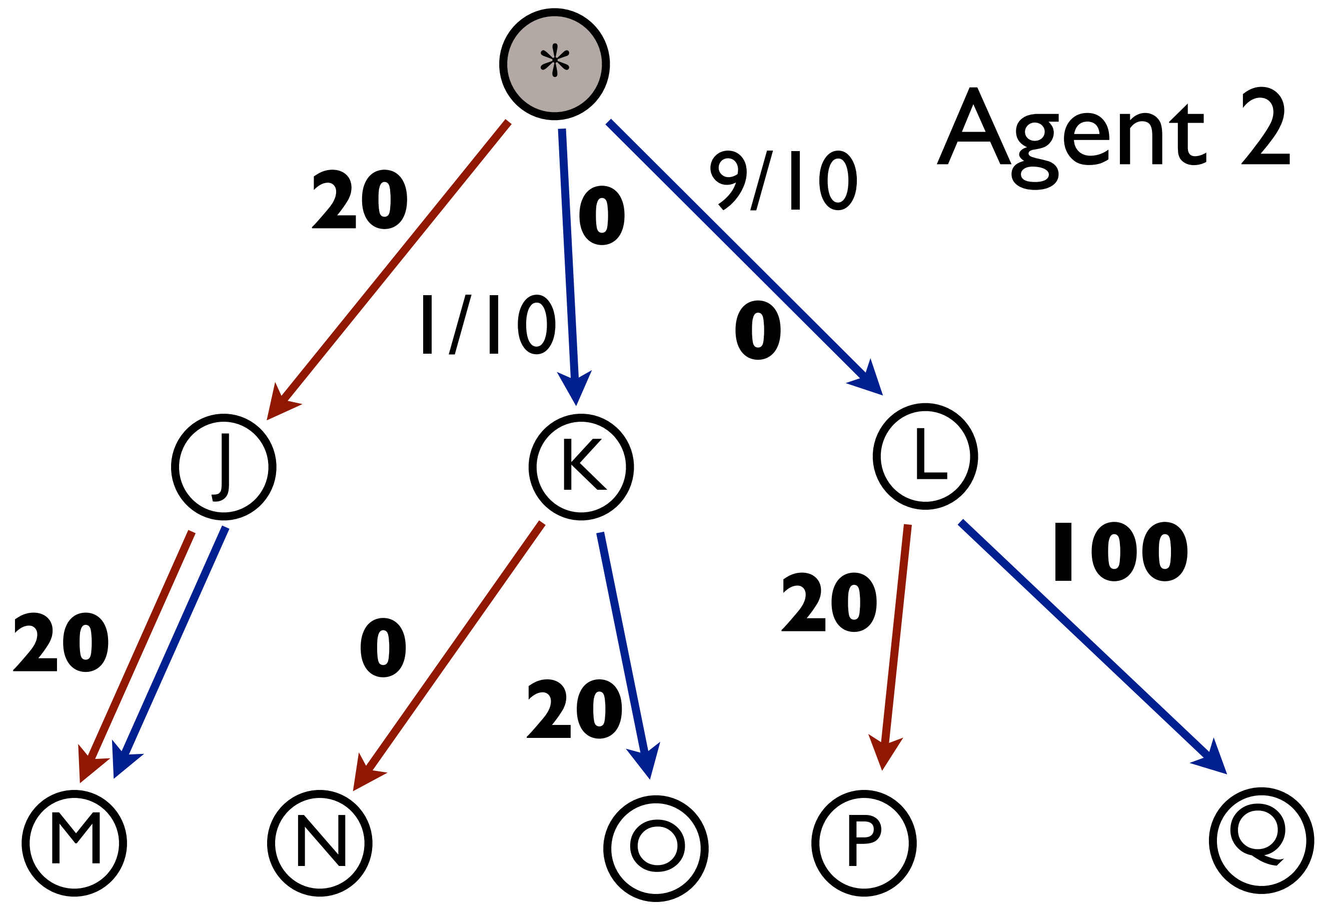

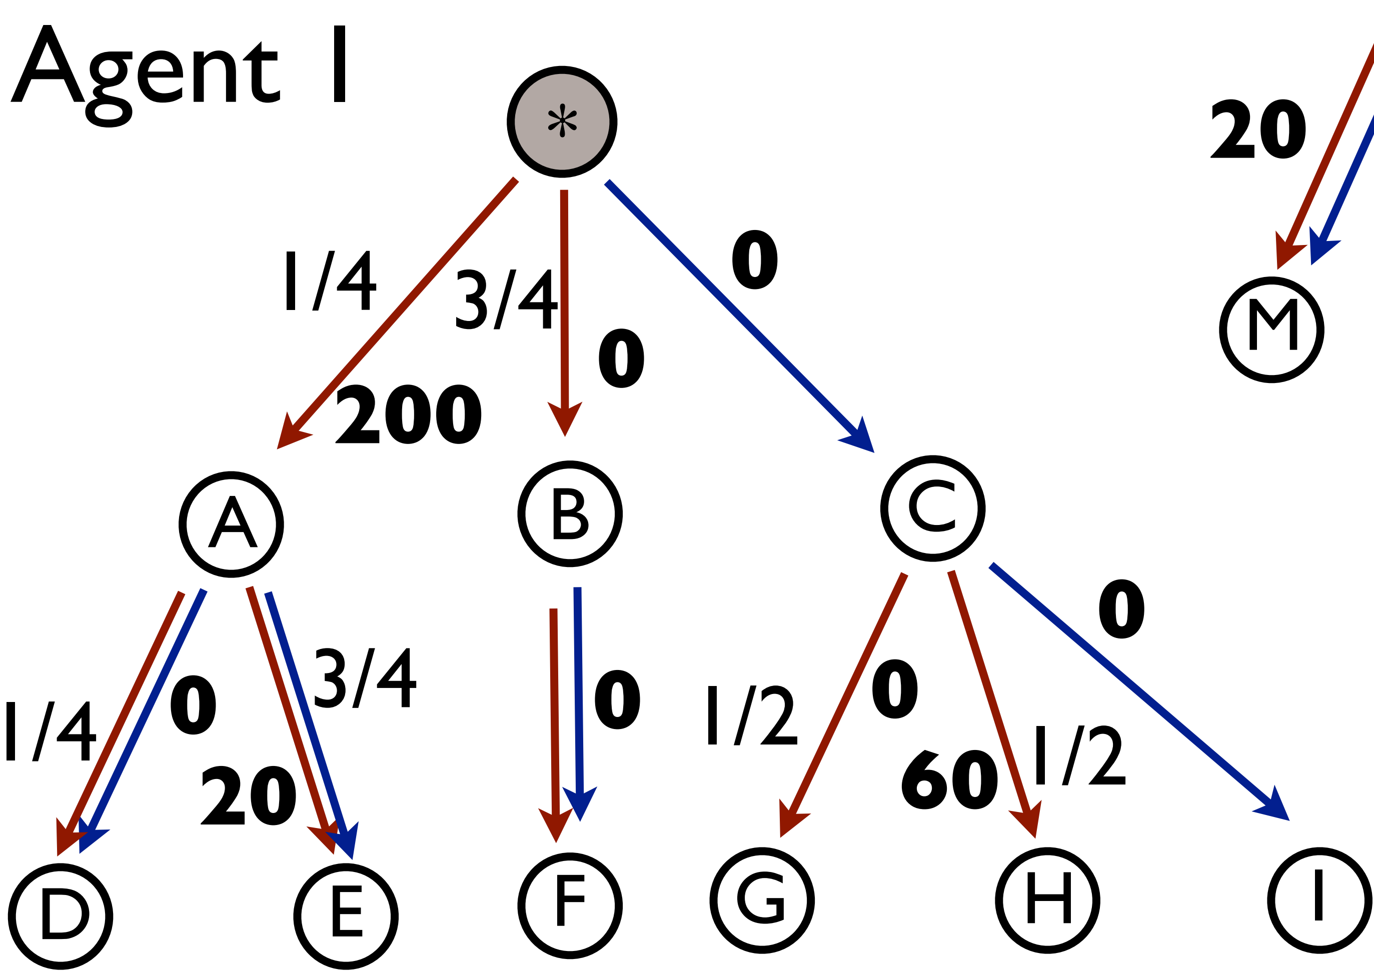

Dynamic team mechanism example

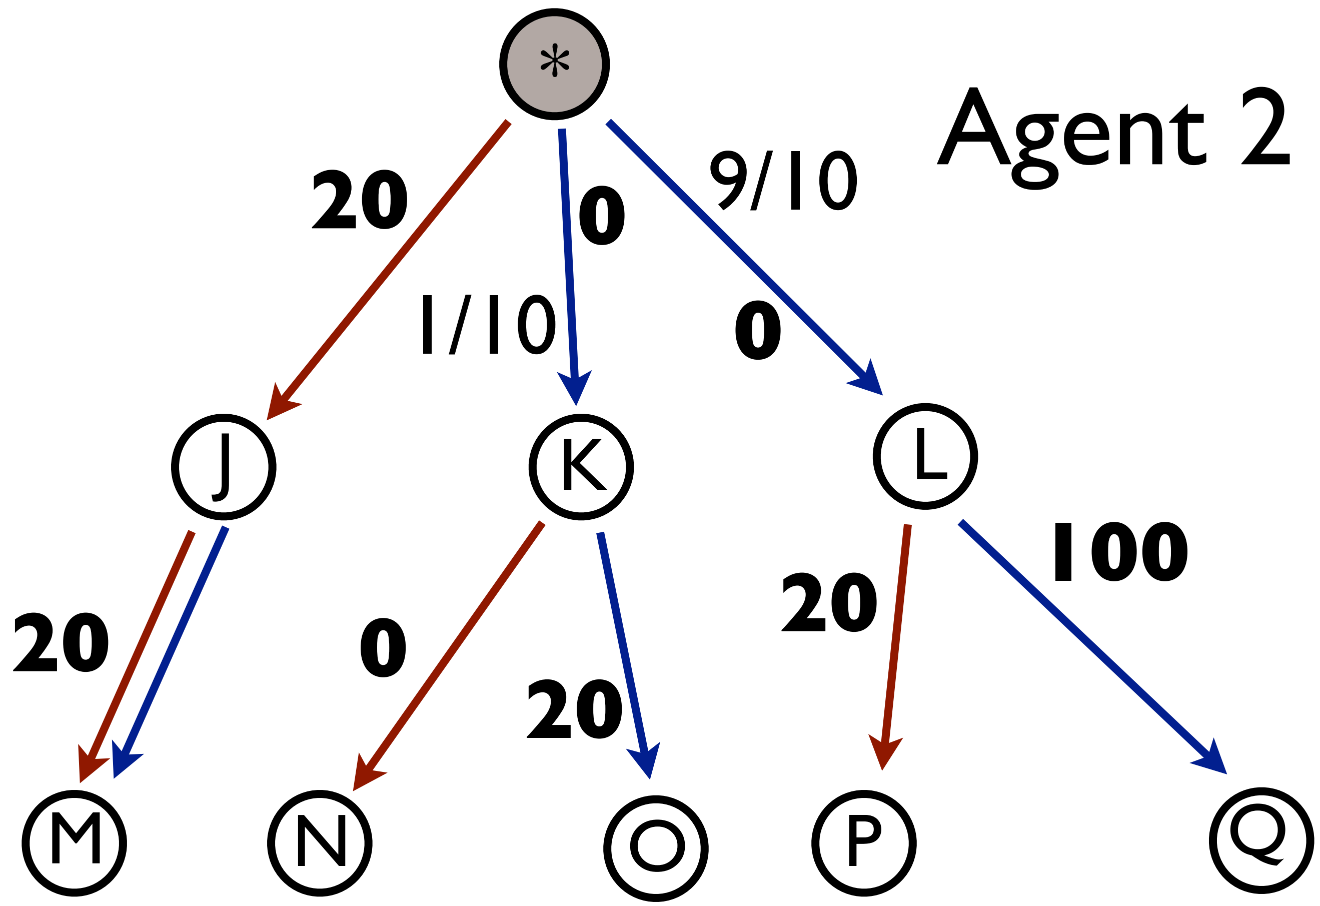

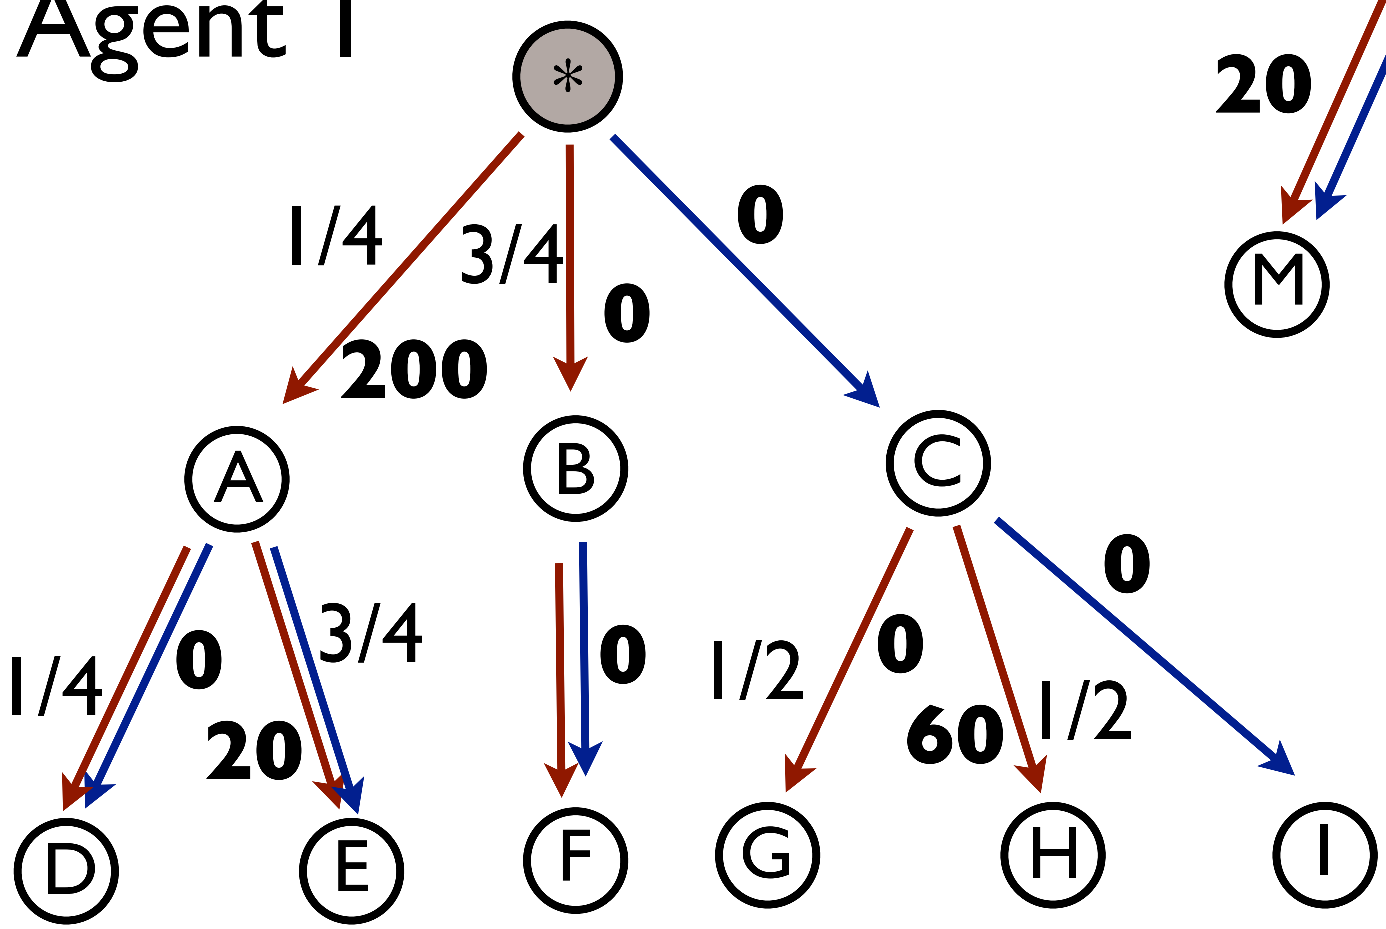

optimal policy (γ close to 1):

- * → blue
- AJ → red or blue
- BJ → red or blue
- CK → red
- CL → blue

Dynamic team mechanism example

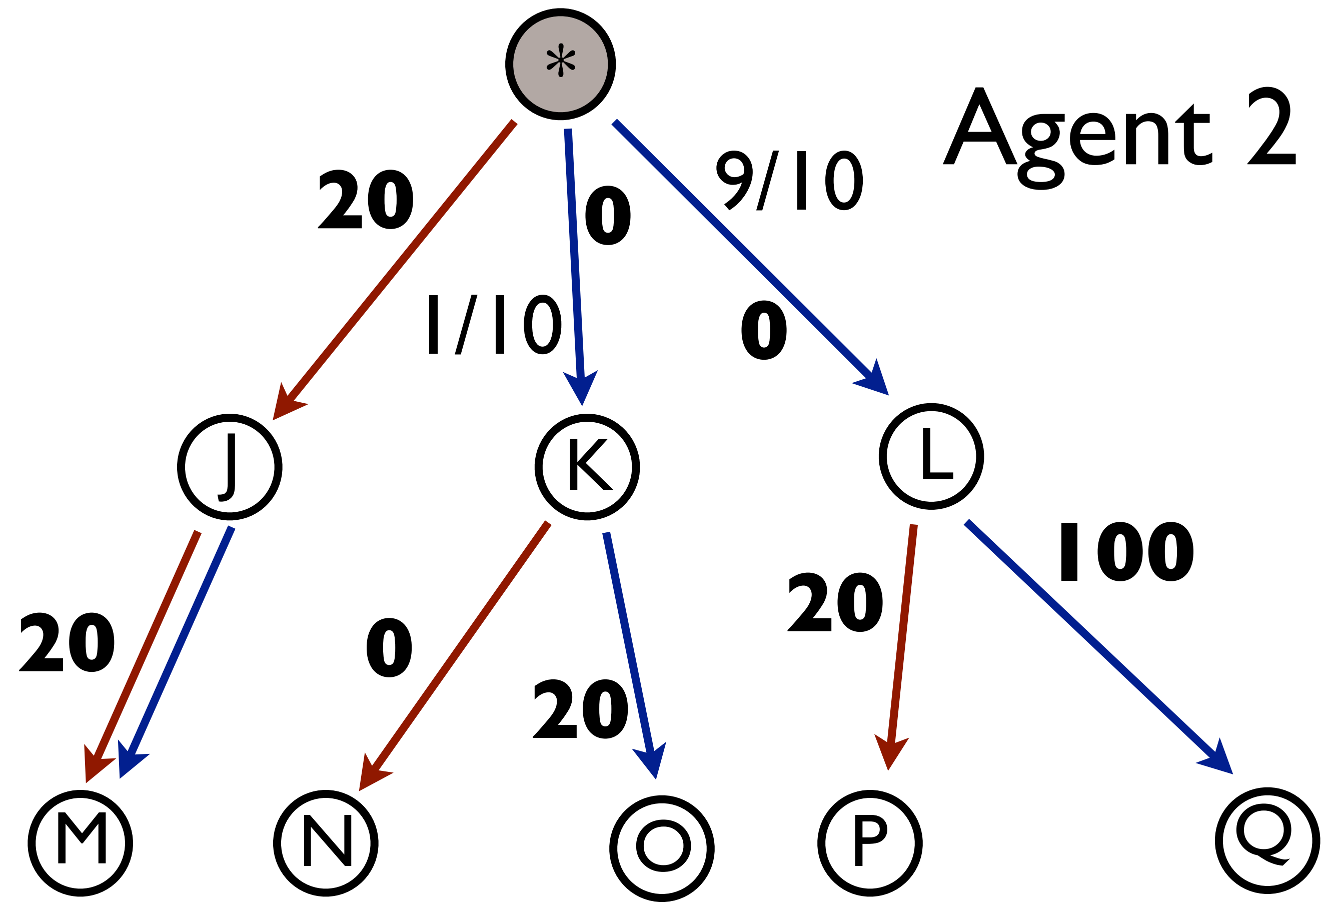

Agent 1

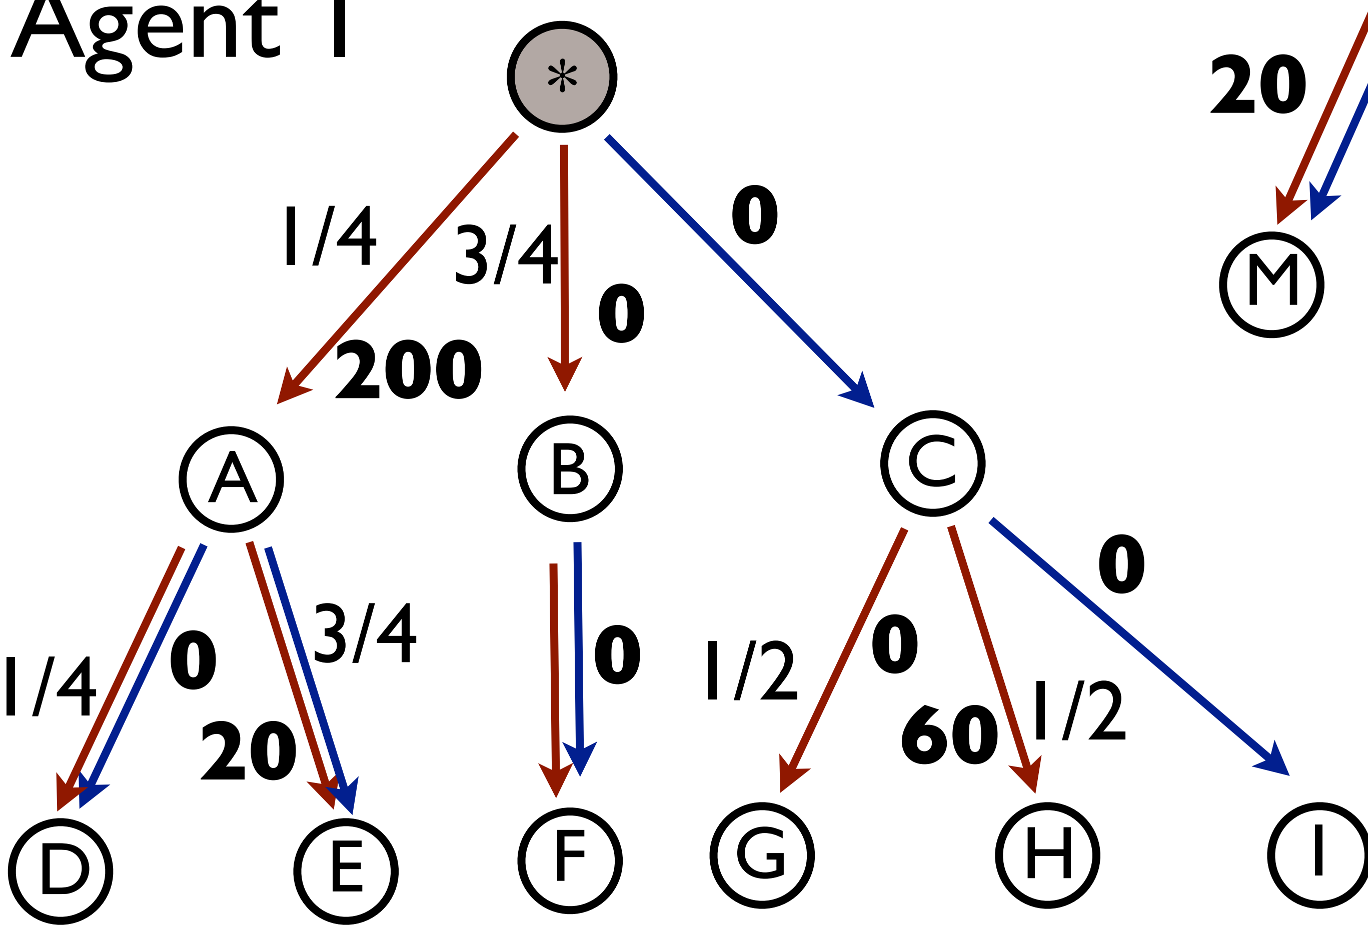

- $T_1(*) = 0$, $T_2(*) = 0$

Dynamic team mechanism example

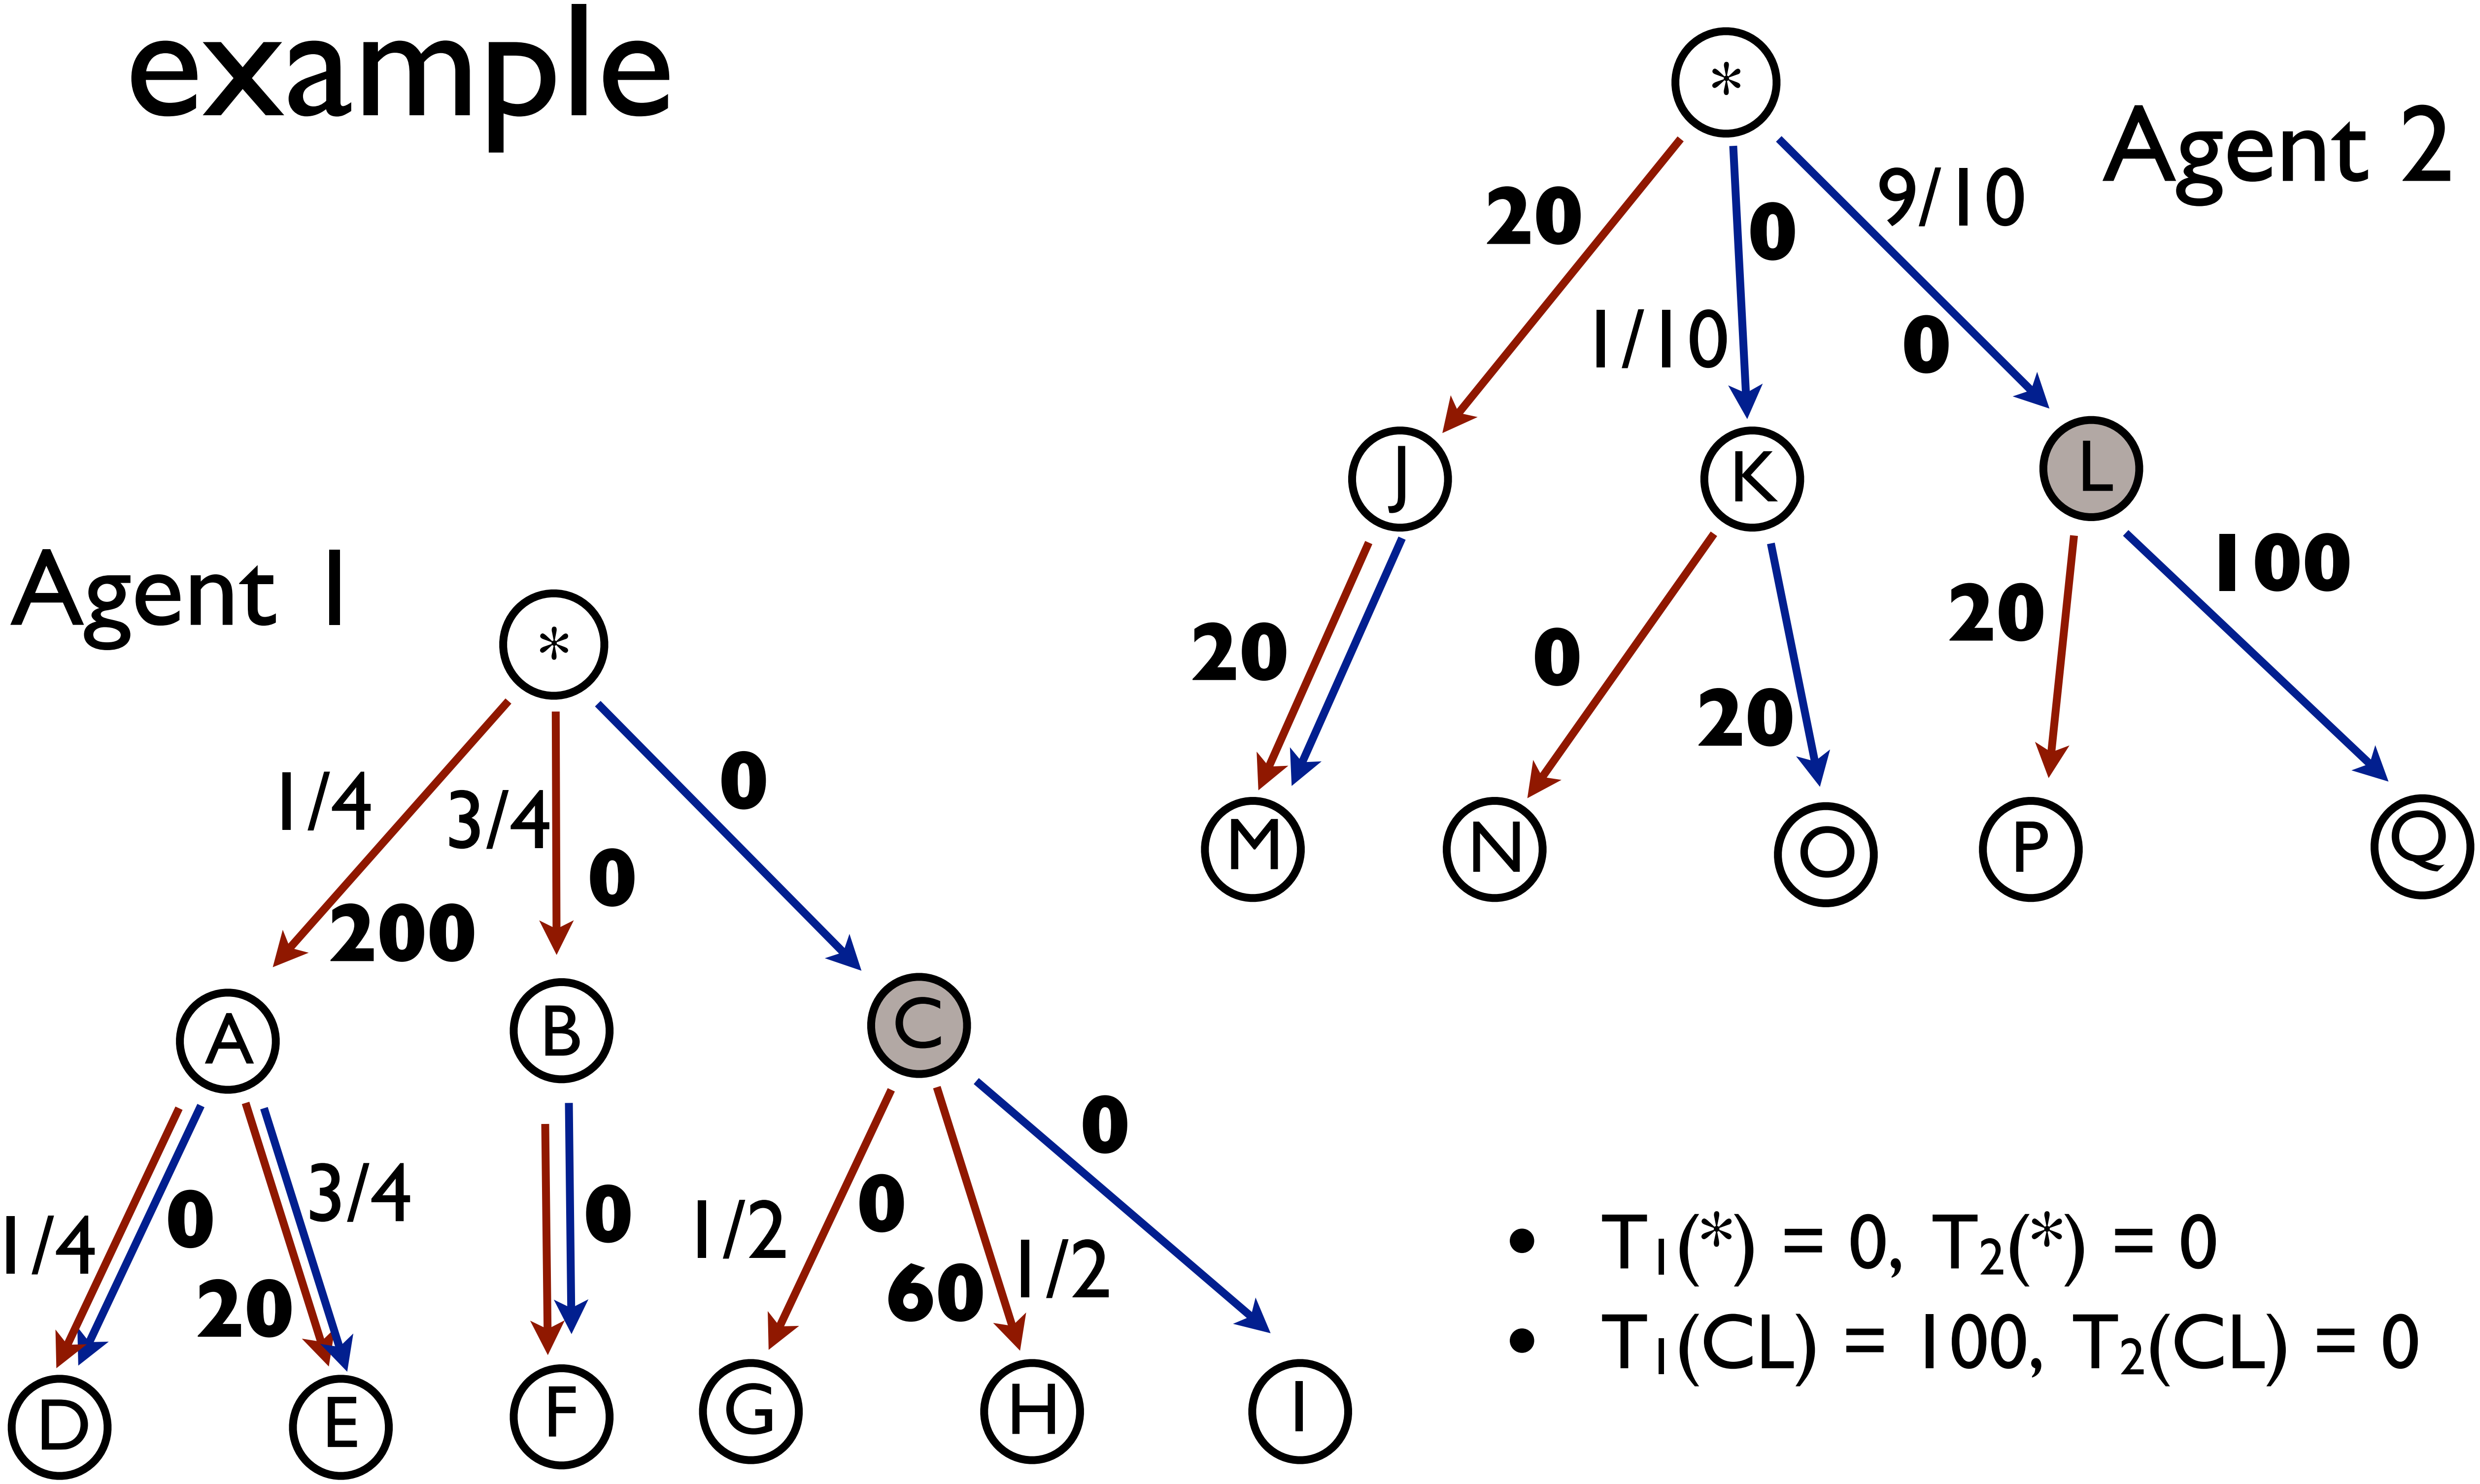

- $T_1(*) = 0$, $T_2(*) = 0$
- $T_1(CL) = 100$, $T_2(CL) = 0$

Dynamic team mechanism

<u>Theorem</u>: The dynamic team mechanism is truthful and efficient in within-period ex post Nash equilibrium.

[Athey & Segal, 07]

Dynamic-Groves mechanism class

- Follows efficient policy given agent reports; defines payments such that:
  - Each agent's expected sum of payments when he follows strategy σ equals the expected value other agents obtain when he follows σ, minus some quantity independent of σ.

Dynamic-Groves mechanism class

Theorem: Every dynamic-Groves mechanism is truthful and efficient in within-period ex post Nash equilibrium.

[Cavallo, Parkes, & Singh, 07]

Proof: Each agent obtains social utility (aligns incentives) minus some constant (doesn't distort).

Dynamic-Groves: all efficient mechanisms

Theorem: For unrestricted types, the dynamic-Groves class exactly corresponds to the history-independent dynamic mechanisms that are truthful and efficient in within-period ex post Nash equilibrium. [Cavallo, 08]

→ For within-period ex post efficient (and history-independent) dynamic mechanism design, dynamic-Groves is the only game in town.

Dynamic-Groves: all efficient mechanisms

Theorem: For unrestricted types, the dynamic-Groves class exactly corresponds to the history-independent dynamic mechanisms that are truthful and efficient in within-period ex post Nash equilibrium. [Cavallo, 08]

- Generalizes [Green & Laffont, 77] (Groves class unique for static settings).
- Proof idea: If non-Groves, there is always some type for which incentives are sufficiently distorted from efficiency.

Budget & participation

- Given characterization theorem, if we demand efficiency in strongest sense, we know what the possibilities are.
- Now pick mechanisms in class with desirable budget/participation properties.
  - basic “team mechanism” won’t fly – extreme budget imbalance
  - need to recover payments...

Recovering payments:
ex ante charge (EAC)

Charge agents some quantity computed “ex ante” of anything they report.

Recovering payments:
ex ante charge (EAC)

- At every time-step:
  - Choose efficient decision given reported types.
  - Make Groves payments.
  - Charge each agent a quantity based only on the reported types of other agents *in the first time-step*: $(1-\gamma)$ times *total* value other agents would obtain, in expectation from beginning of mechanism, if policy optimal for them was chosen.

$$T_i(\theta^t) = r_{-i}(\theta^t_{-i}, \pi^*(\theta^t)) - (1-\gamma)V_{-i}(\theta^0_{-i})$$

[Cavallo, Parkes, & Singh, 06]

Dynamic-EAC mechanism example

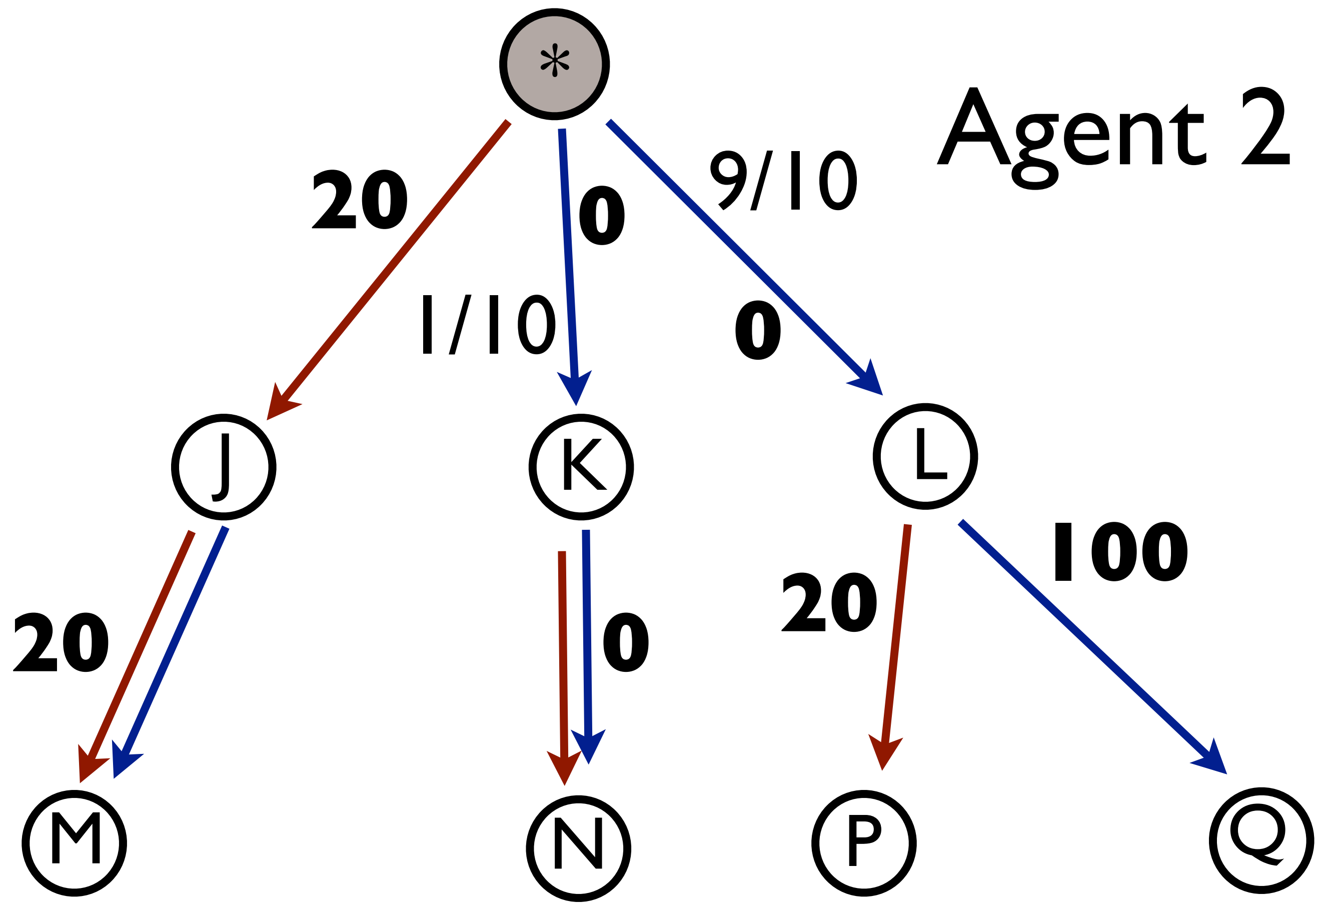

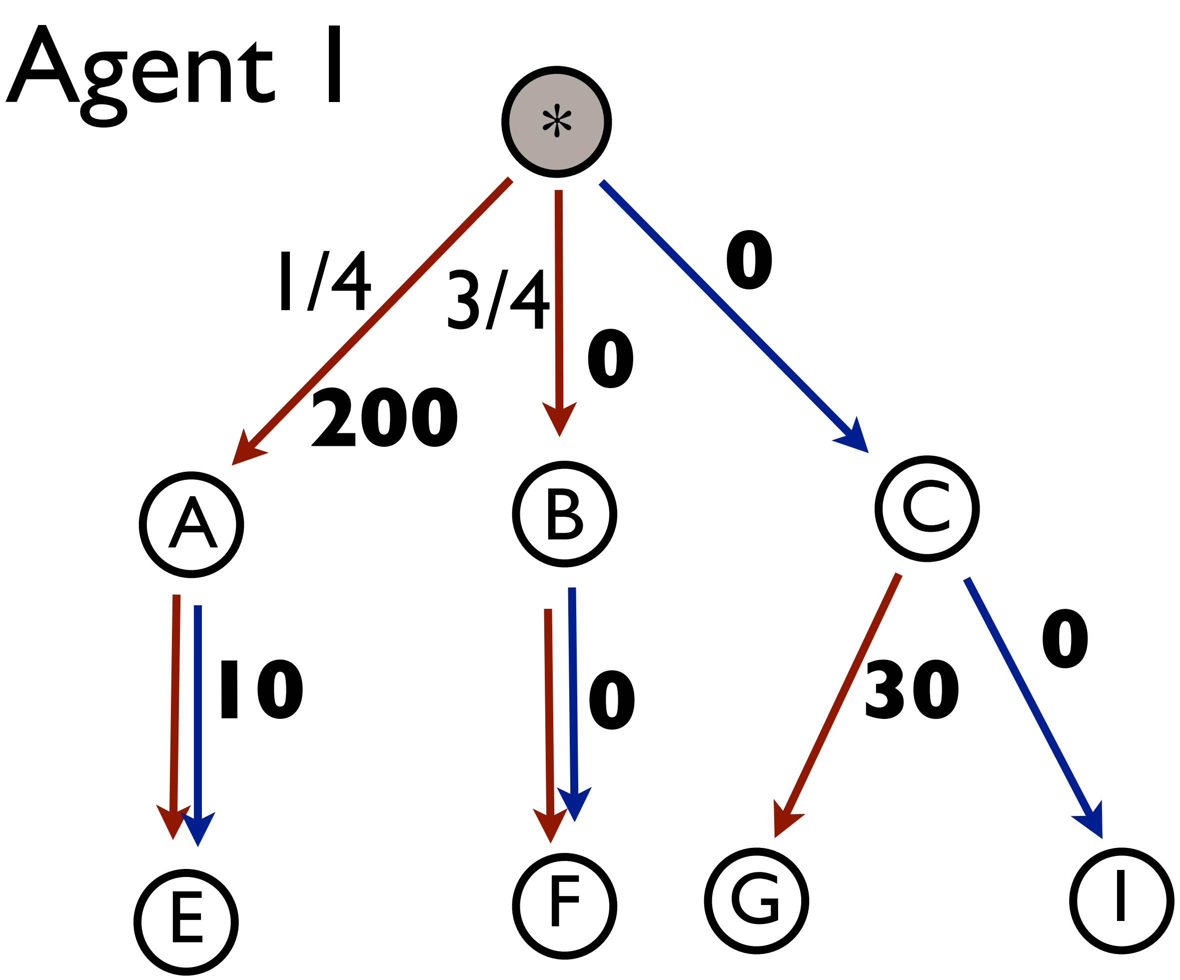

- $\gamma = 0.9$
- $V_1(*_{-2}) = 50 + \gamma 10 = 59$
- $V_2(*_{-1}) = \gamma 90 = 81$
- $T_1(*) = -8.1$, $T_2(*) = -5.9$
- $T_1(CL) = 100 - 8.1$, $T_2(CL) = -5.9$
- ...

Recovering payments:
ex ante charge (EAC)

Theorem: The dynamic-EAC mechanism is truthful and efficient in within-period ex post Nash equilibrium, ex ante individual rational, and ex ante no-deficit.

[Cavallo, Parkes, & Singh, 06]

Weak IR and budget-balance properties

- With dynamic-EAC scheme agents will "sign up" at beginning of mechanism, but may wish to back out...
- Same for center.
- Can we strengthen?

Dynamic-VCG
[Bergemann & Valimaki, 08]

- At each time-step, pay each agent i the expected value other agents would obtain if i were ignored after one step, minus the value they'd obtain if i were always ignored.

➡ Each agent has to pay the amount he inhibits other agents from obtaining value (now and in the future) by his current report.

Dynamic-VCG [Bergemann & Valimaki, 08]

- At each time-step, pay each agent i the expected value other agents would obtain if i were ignored after one step, minus the value they'd obtain if i were always ignored.

$$T_i(\theta^t) = r_{-i}(\theta^t_{-i}, \pi^*(\theta^t)) + \gamma \mathbb{E}[V_{-i}(\tau(\theta^t_{-i}, \pi^*(\theta^t)))] - V_{-i}(\theta^t_{-i})$$

Dynamic-VCG mechanism

example

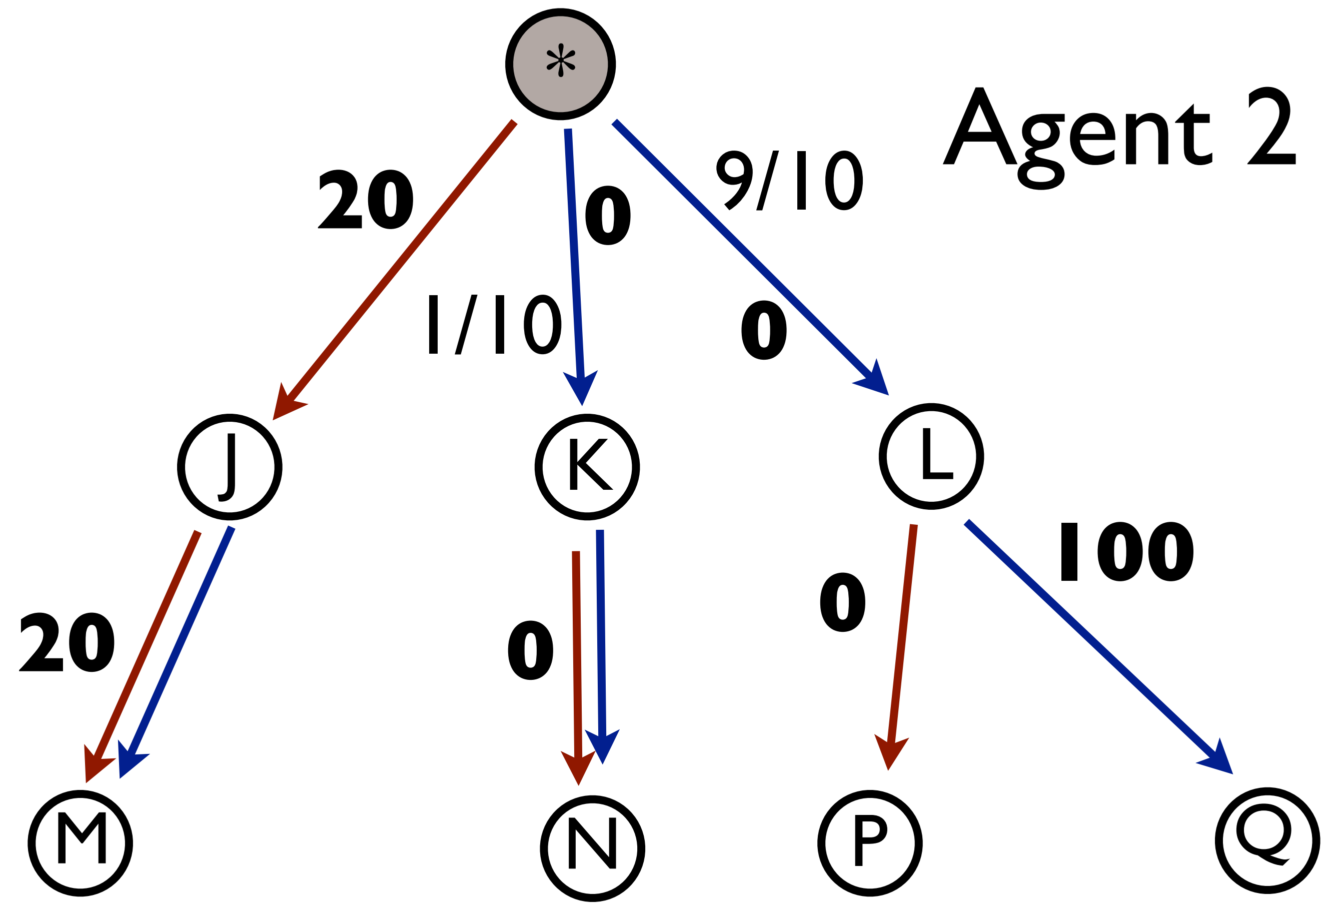

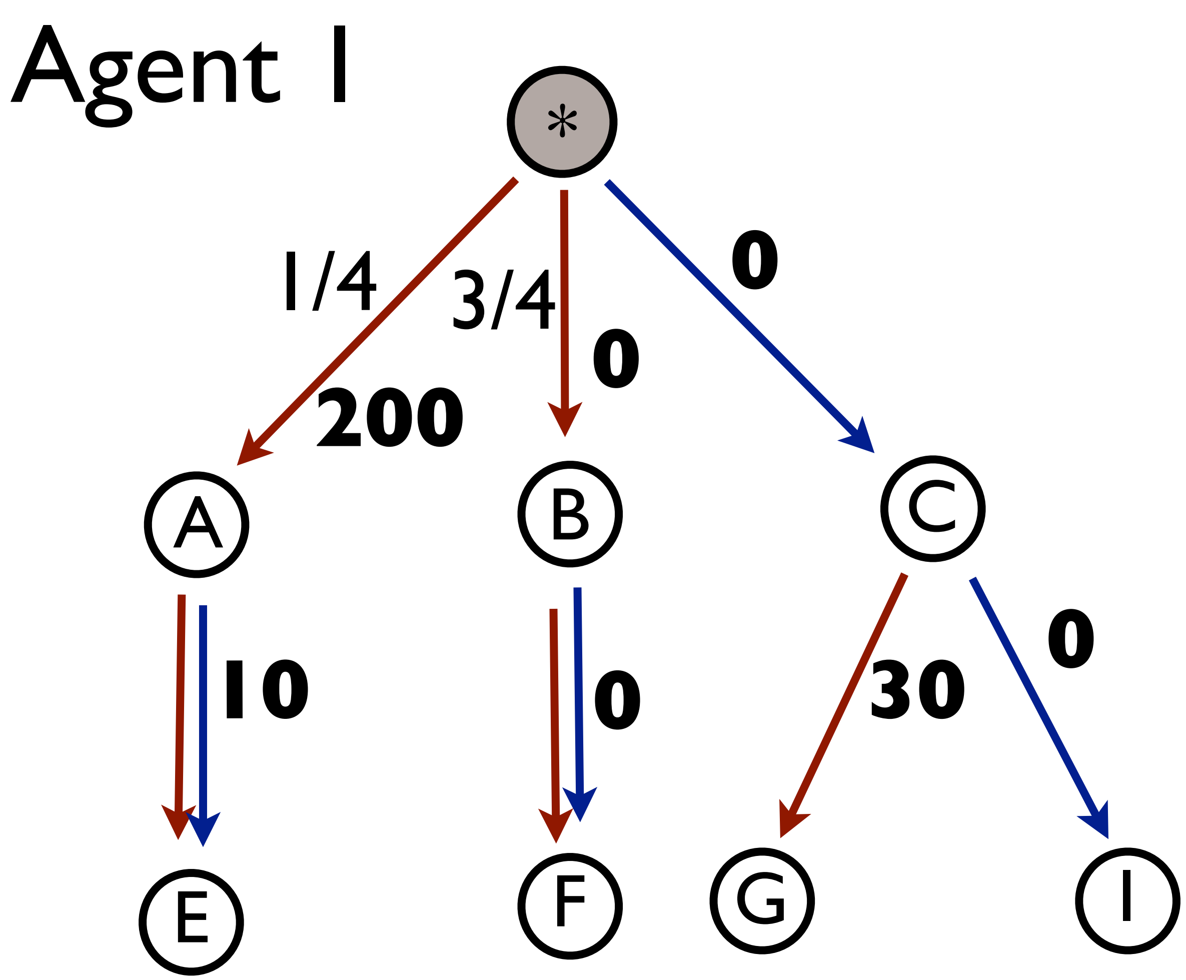

- $\gamma = 0.9$
- $T_1(*) = \gamma 90 - \gamma 90$
- $T_2(*) = \gamma 30 - (50 + \gamma 10)$
- $T_1(CL) = 100 - 100$
- $T_2(CL) = 0 - 30$
- ...

Dynamic-VCG [Bergemann & Valimaki, 08]

$$T_i(\theta^t) = r_{-i}(\theta^t_{-i}, \pi^*(\theta^t)) + \gamma\mathbb{E}[V_{-i}(\tau(\theta^t_{-i}, \pi^*(\theta^t)))] - V_{-i}(\theta^t_{-i})$$

- No payment to any agent in any period is positive.

$$r_{-i}(\theta^t_{-i}, \pi^*(\theta^t)) + \gamma\mathbb{E}[V_{-i}(\tau(\theta^t_{-i}, \pi^*(\theta^t)))] \leq V_{-i}(\theta^t_{-i})$$

- Expected future payoff to every agent i, from any joint state, at any time t, is:

$$V(\theta^t) - V_{-i}(\theta^t_{-i}) \geq 0$$

(NB: assumes no negative values)

Dynamic-VCG
[Bergemann & Valimaki, 08]

Theorem: The dynamic-VCG mechanism is truthful and efficient in within-period ex post Nash equilibrium, within-period ex post individual rational, and ex post no-deficit.

Dynamic-VCG: good social-welfare?

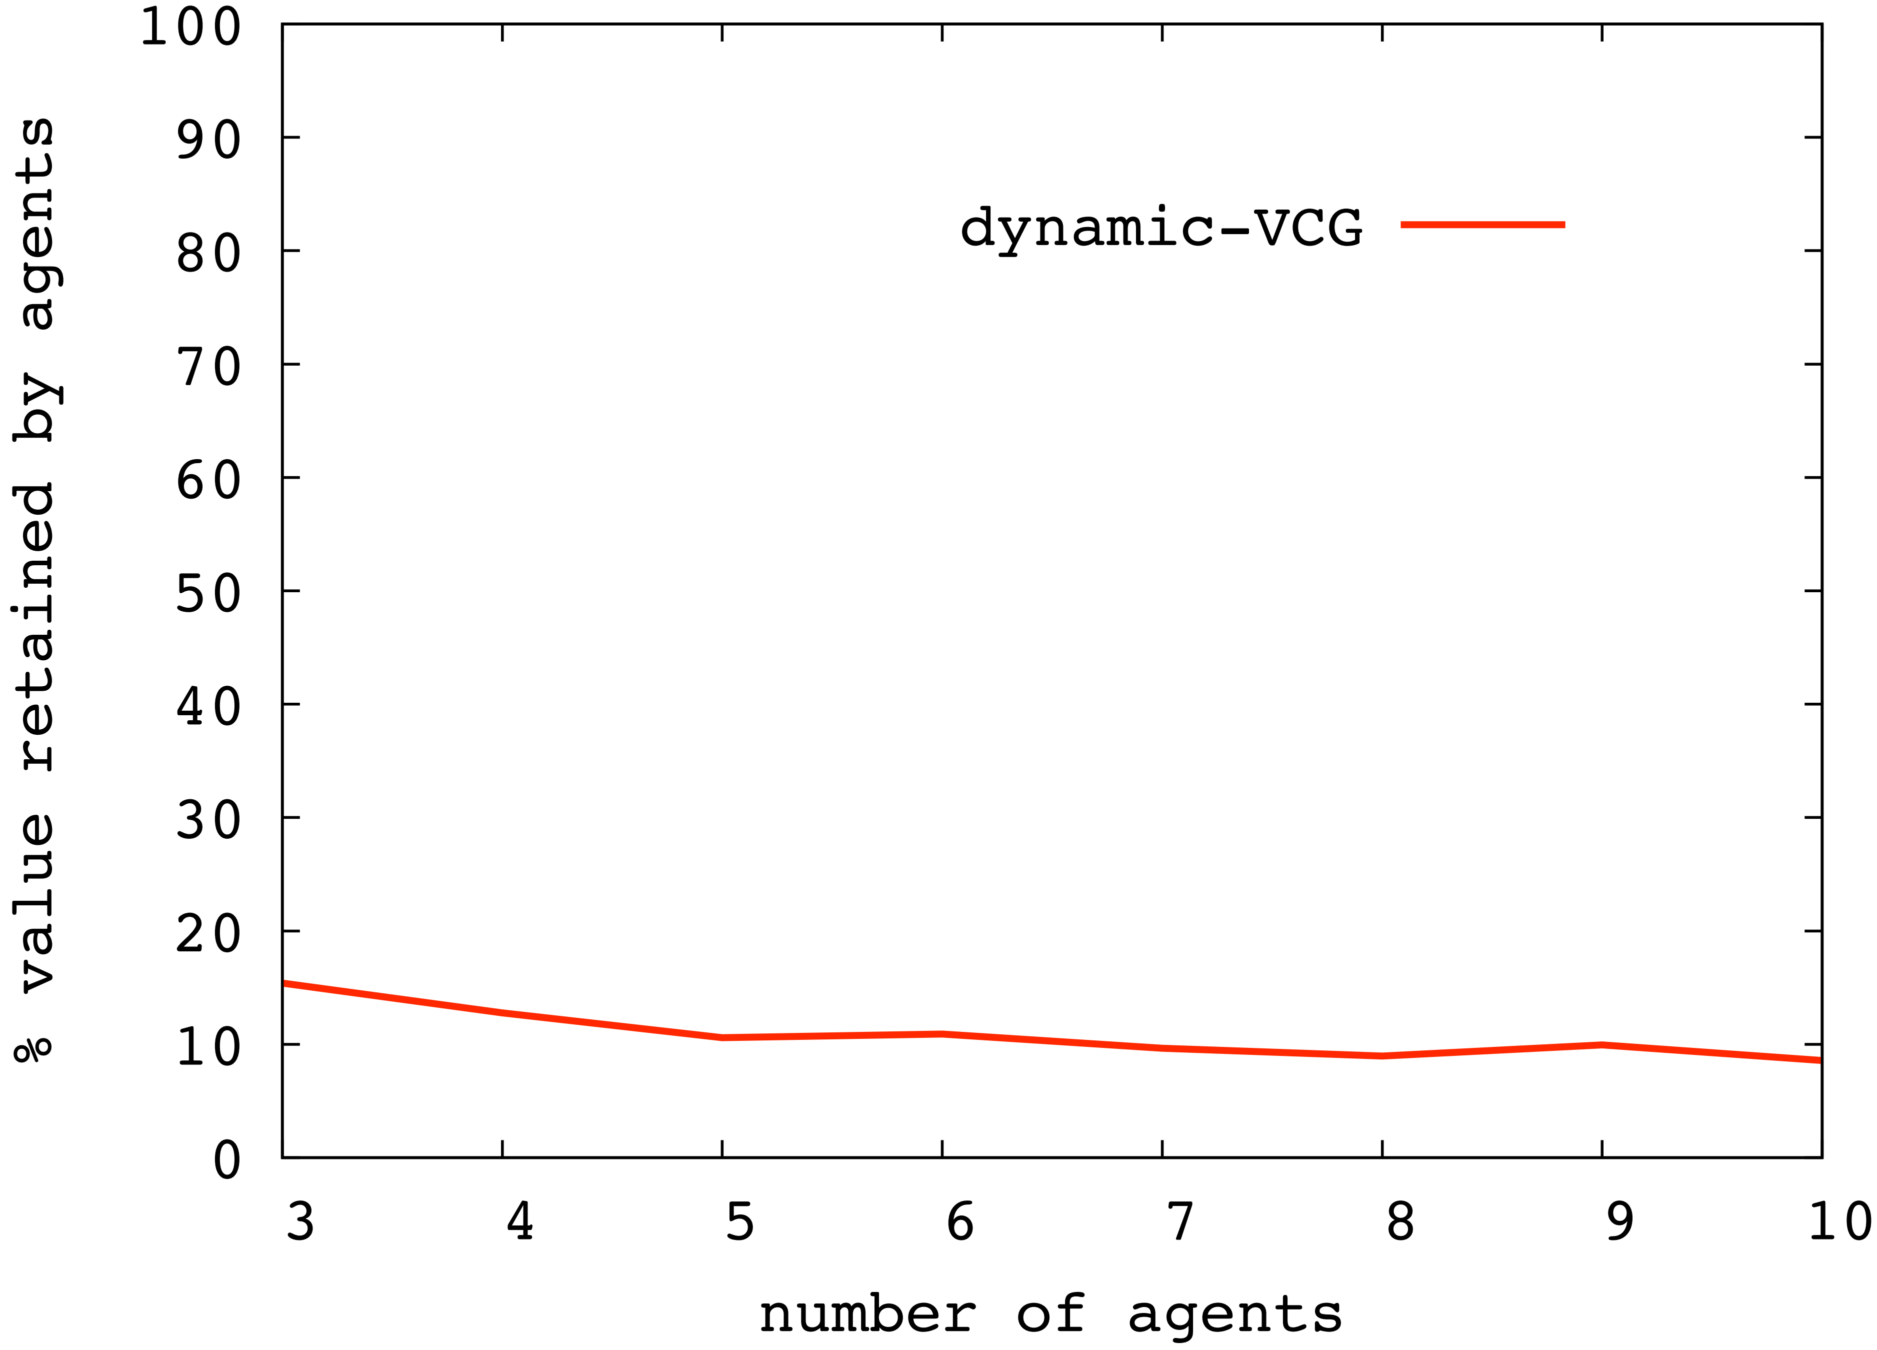

In a single-item allocation setting, with values normally distributed.

Dynamic-VCG: good social-welfare?

Theorem: Among all history-independent mechanisms that are efficient in within-period ex post Nash equilibrium and within-period ex post individual rational, dynamic-VCG yields the *most* expected revenue, for *every* joint type.

[Cavallo, 08]

Dynamic-VCG: good social-welfare?

- Since dynamic-VCG can be so bad for the agents, what do we do?
- Think back to the static setting... better budget balance was achieved by redistribution mechanisms; strong budget-balance by moving to Bayes-Nash equilibrium.

A dynamic redistribution mechanism?

- Redistribution much more complicated in the dynamic setting. Now redistribution payment computed in later time periods can potentially be influenced via an agent's reports in earlier periods... in subtle ways.
  - Focus on worlds representable as multi-armed bandits.

Dynamic-VCG for MABs reduces to:

- Determine optimal agent `i` to activate.
  - `i` pays `(1-γ)` times the expected value other agents would get if `i` were always ignored.
  - Other agents pay nothing.

Dynamic-VCG for MABs:

- Winner pays (1-γ) times the expected value other agents would get if he were always ignored.
- Other agents pay nothing.

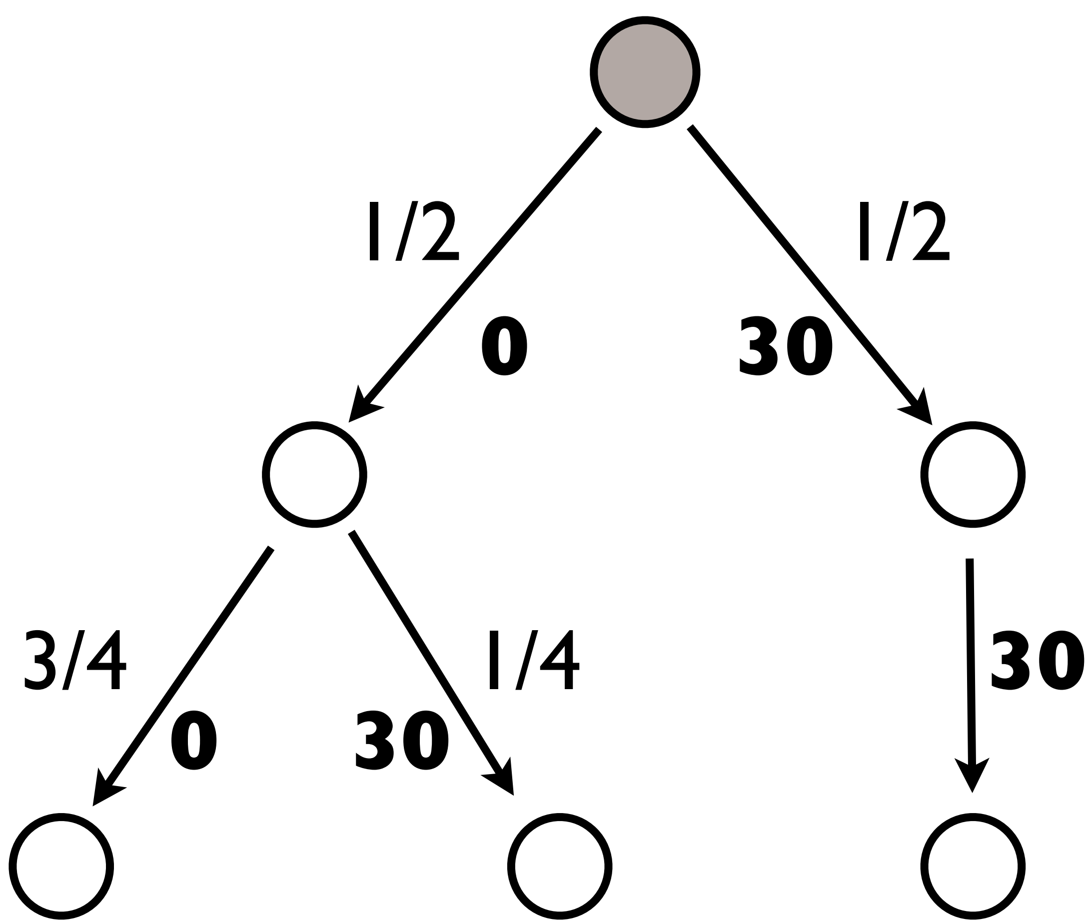

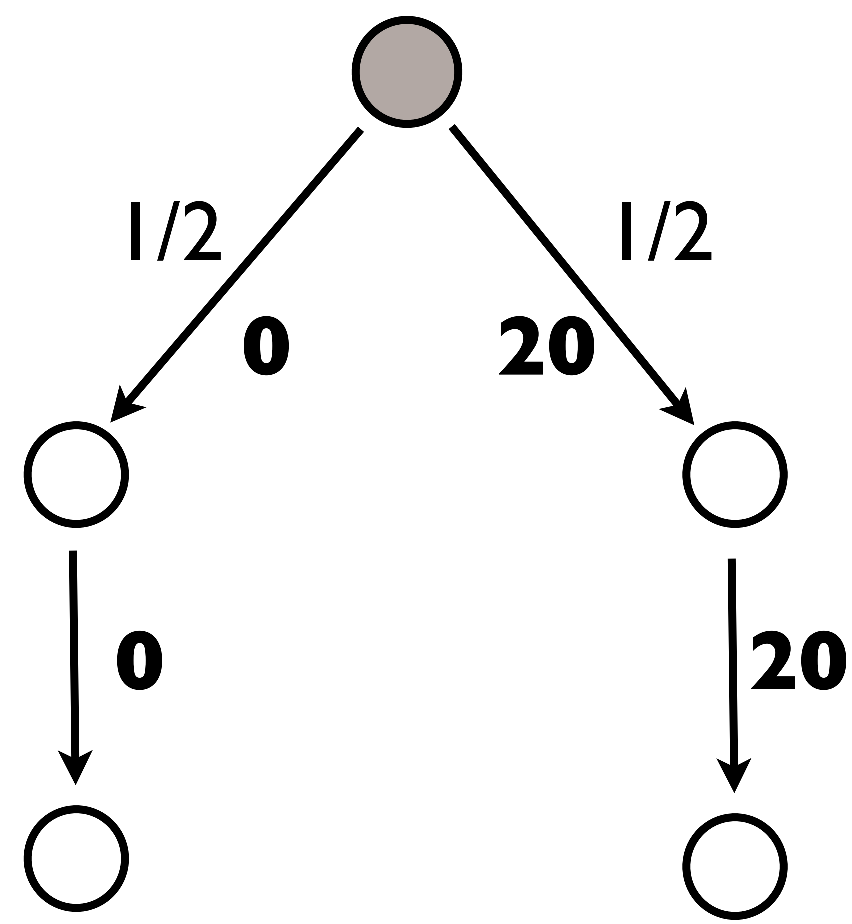

Dynamic-VCG for MABs:

- Winner pays (1-γ) times the expected value other agents would get if he were always ignored.
- Other agents pay nothing.

$T_1 = -(1-\gamma)\ (10 + \gamma 10)$

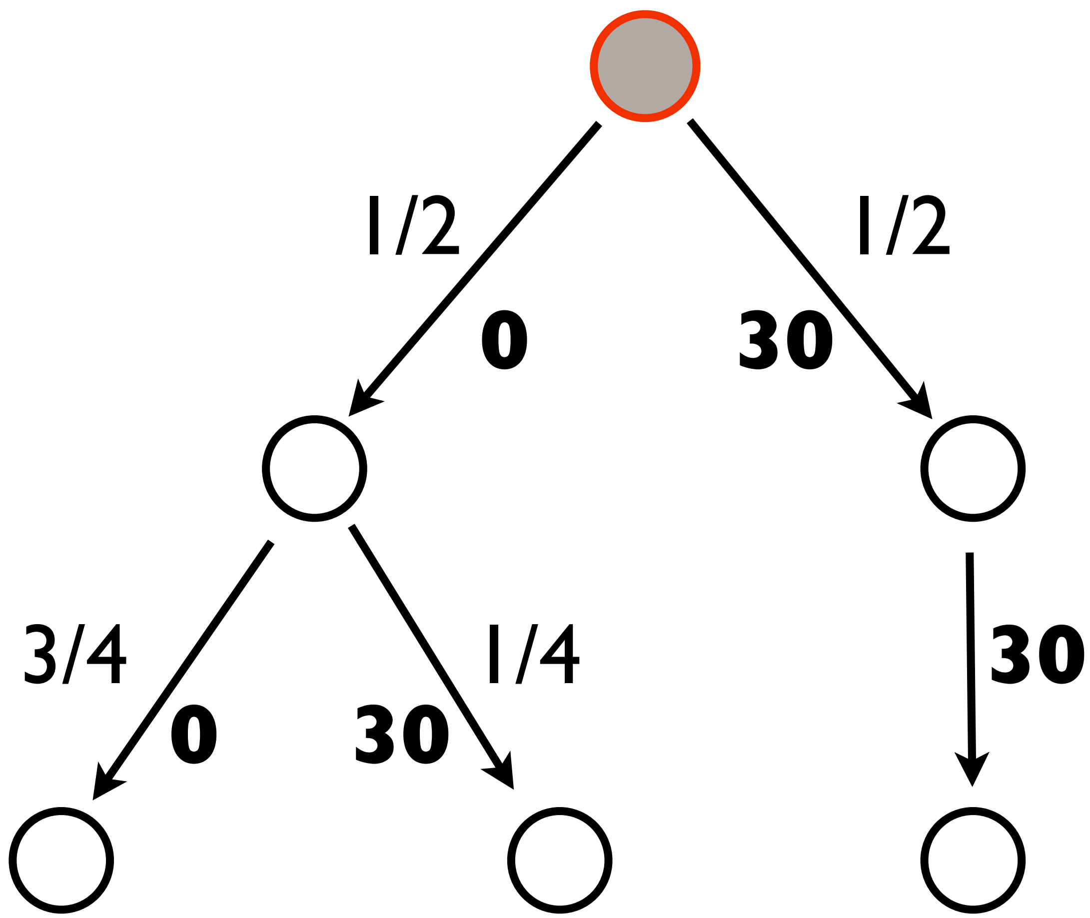

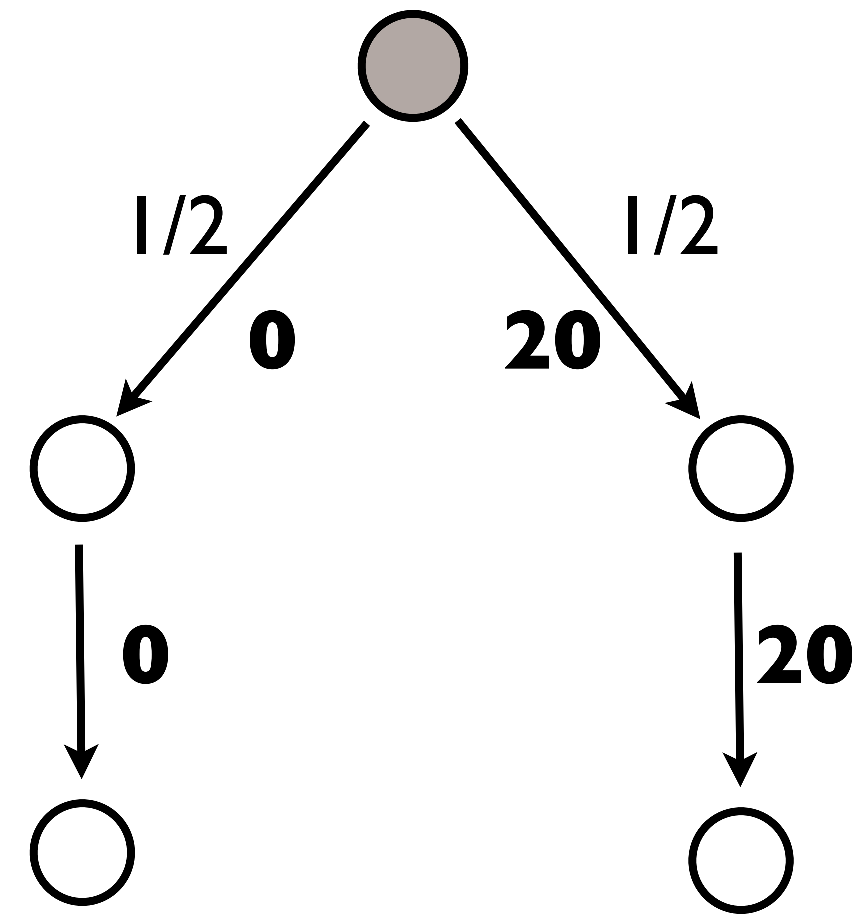

Dynamic-VCG for MABs:

- Winner pays (1-γ) times the expected value other agents would get if he were always ignored.
- Other agents pay nothing.

$T_1 = -(1-\gamma)\ (10 + \gamma 10)$

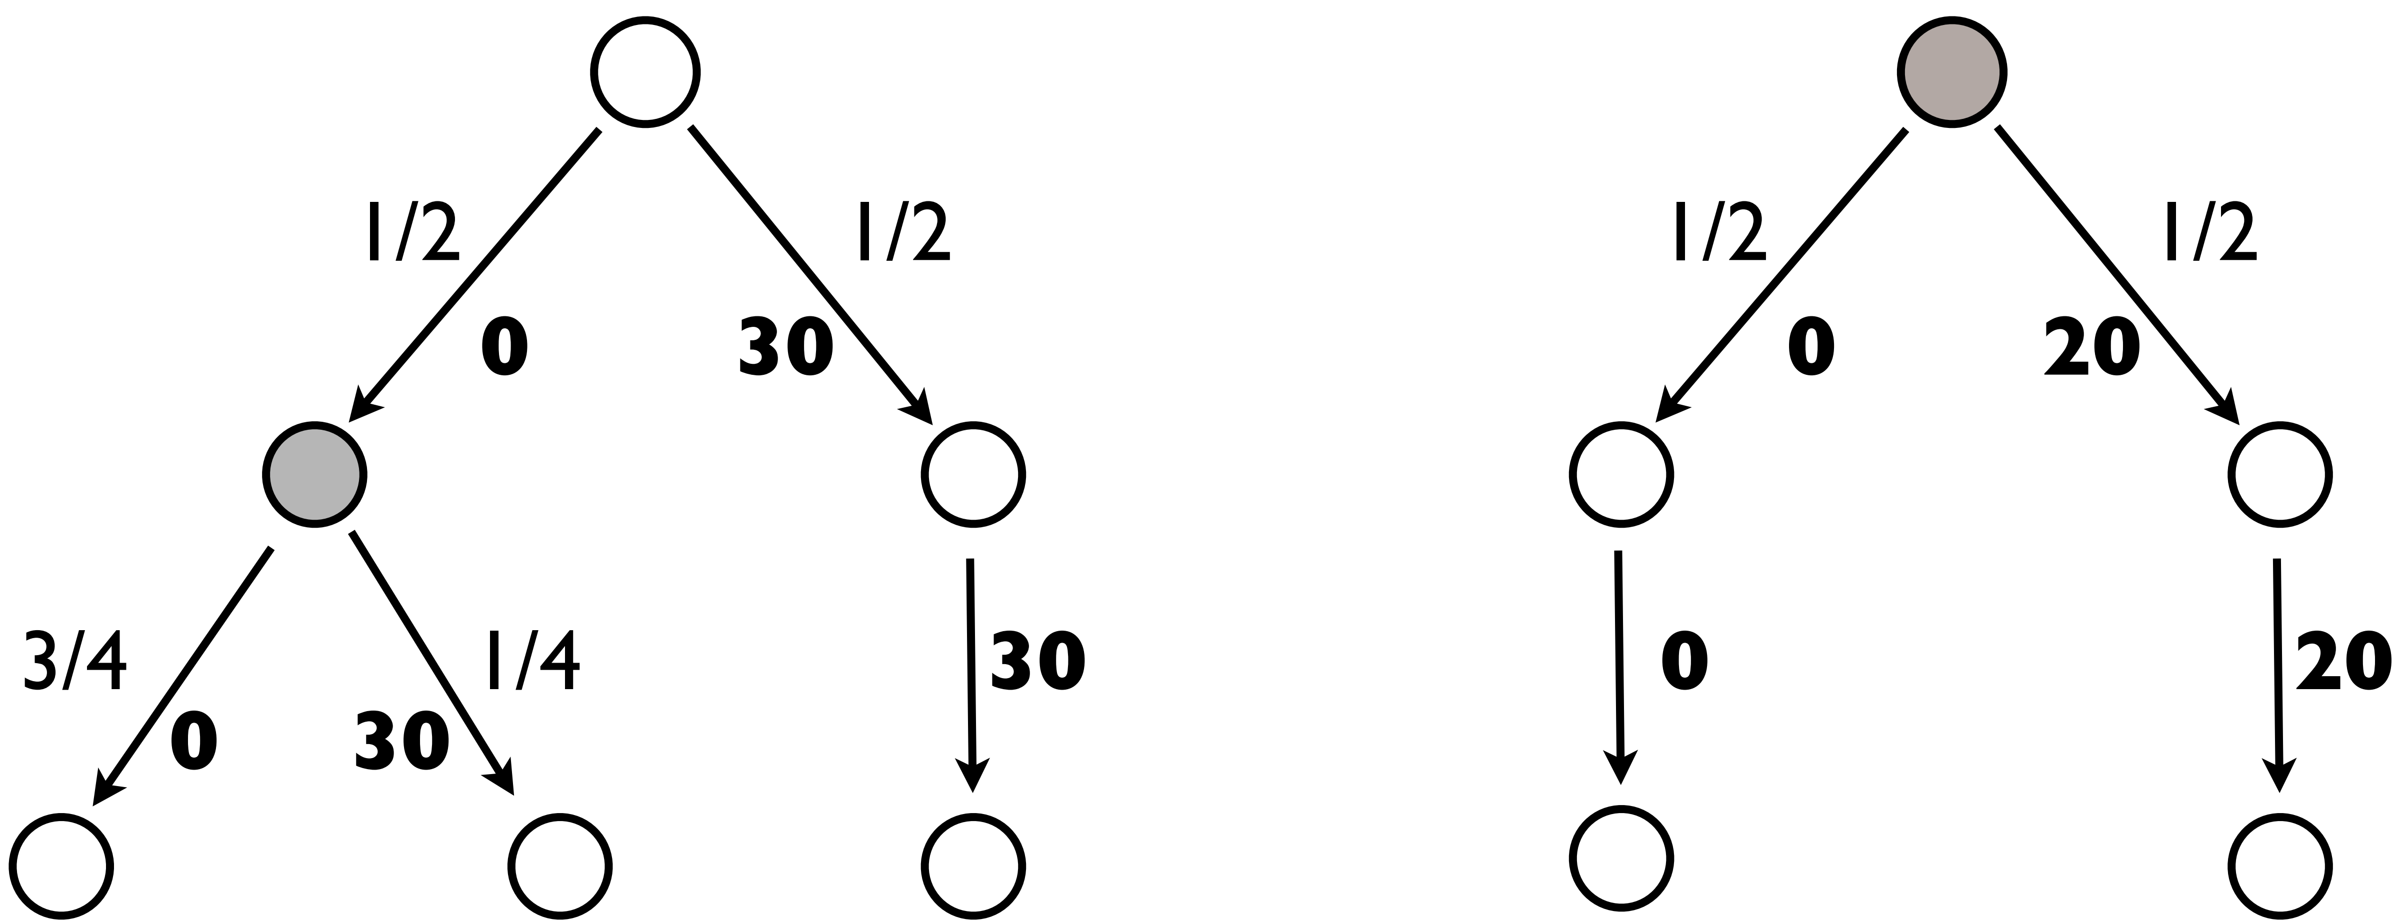

Dynamic-VCG for MABs:

- Winner pays $(1-\gamma)$ times the expected value other agents would get if he were always ignored.
- Other agents pay nothing.

$T_1 = -(1-\gamma)\ (10 + \gamma 10)$

$T_2 = -(1-\gamma)\ 7.5$

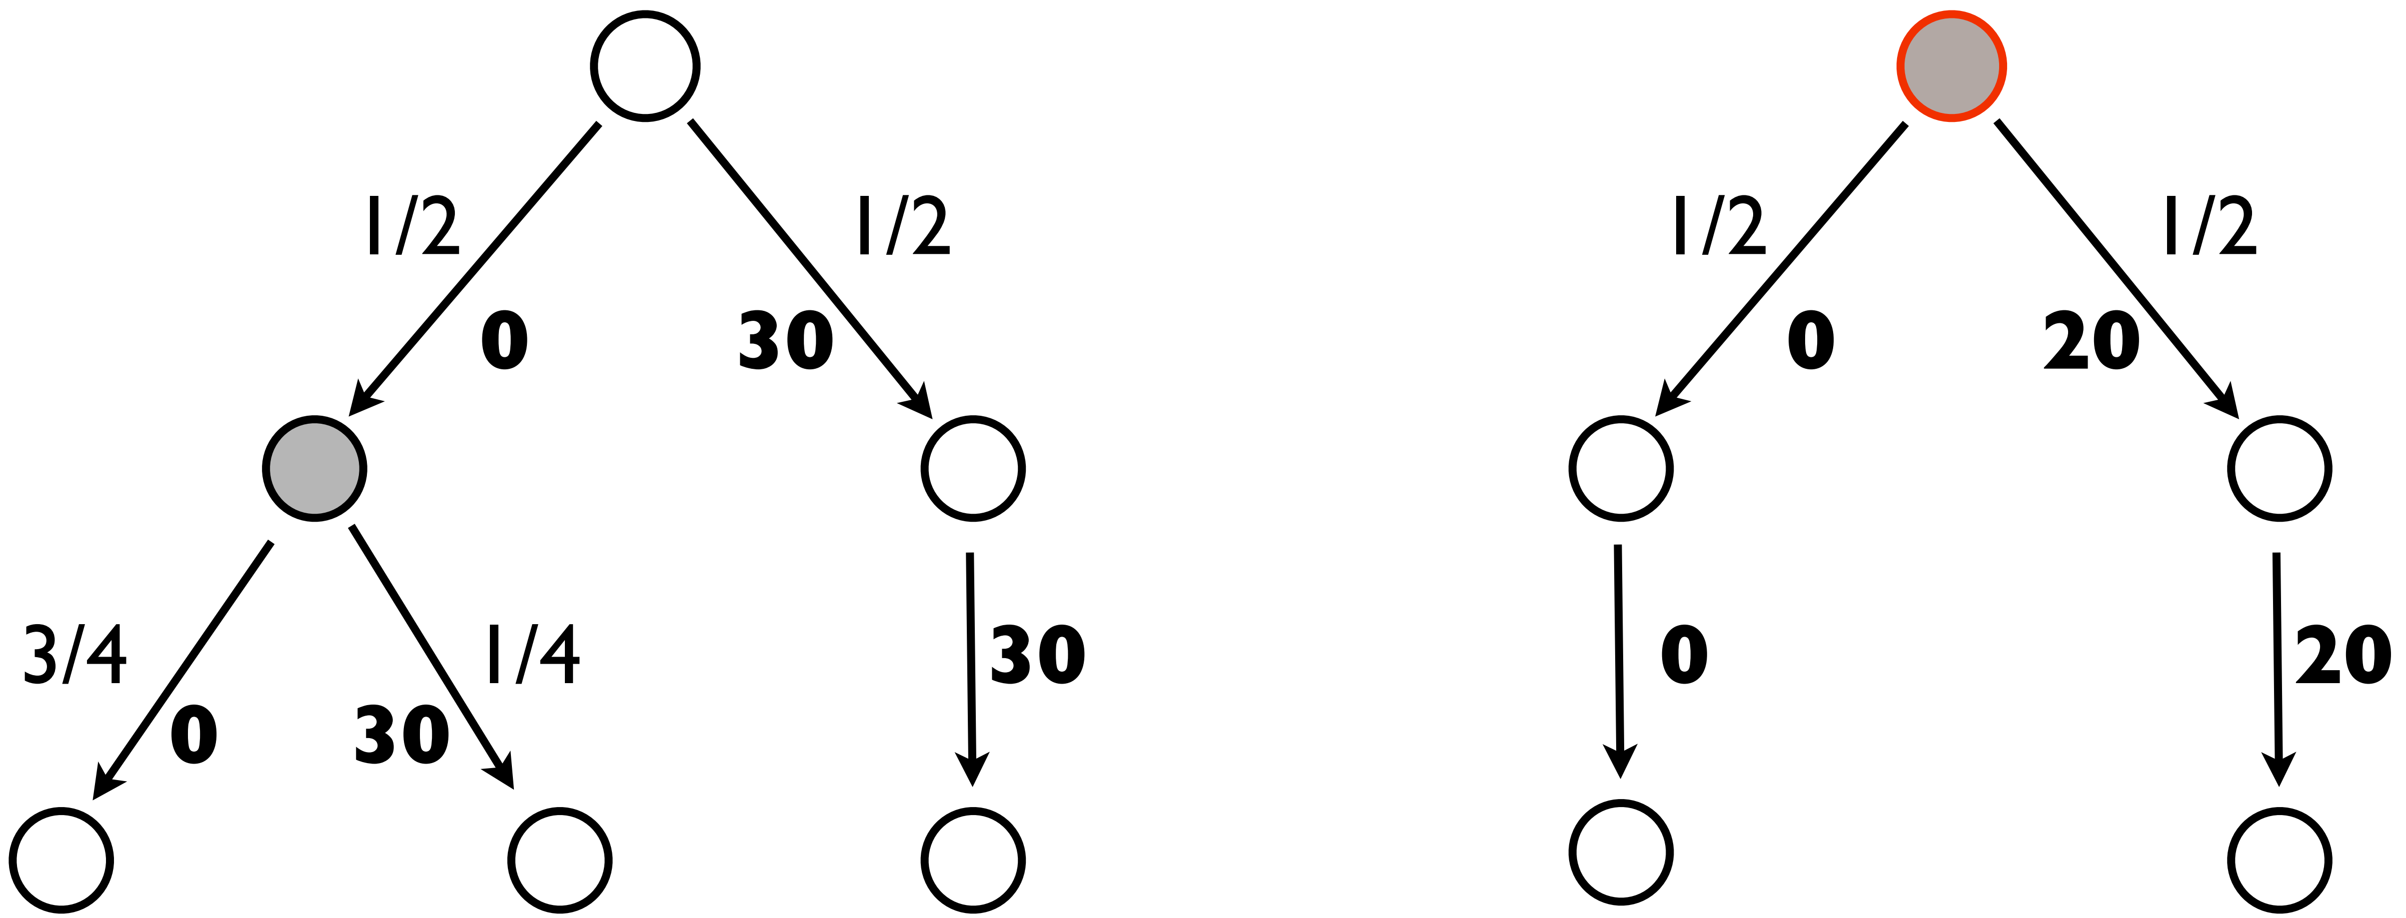

Dynamic-VCG for MABs:

- Winner pays $(1-\gamma)$ times the expected value other agents would get if he were always ignored.
- Other agents pay nothing.

$T_1 = -(1-\gamma)\ (10 + \gamma 10)$

$T_2 = -(1-\gamma)\ 7.5$

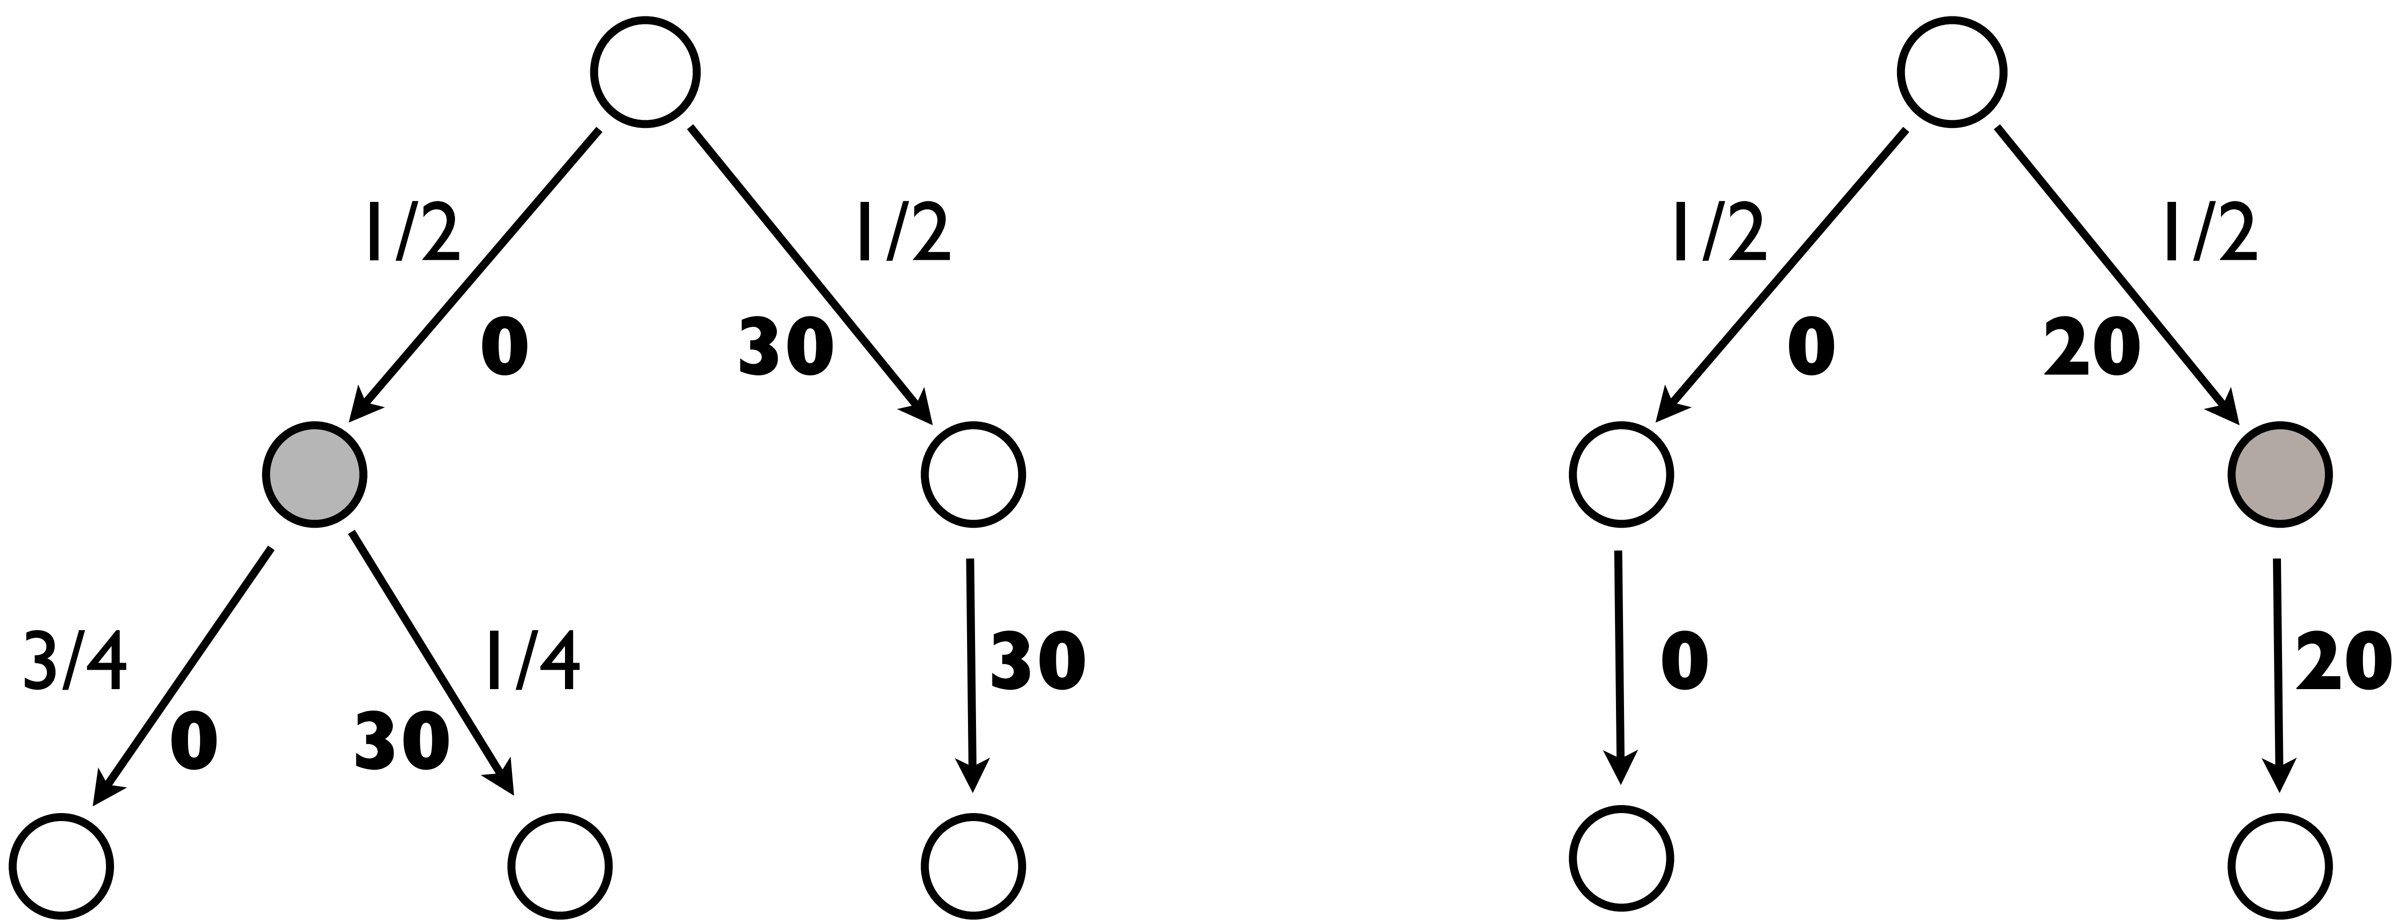

Dynamic-VCG for MABs:

- Winner pays $(1-\gamma)$ times the expected value other agents would get if he were always ignored.
- Other agents pay nothing.

$T_1 = -(1-\gamma)\ (10 + \gamma 10)$

$T_2 = -(1-\gamma)\ 7.5$

$T_2 = -(1-\gamma)\ 7.5$

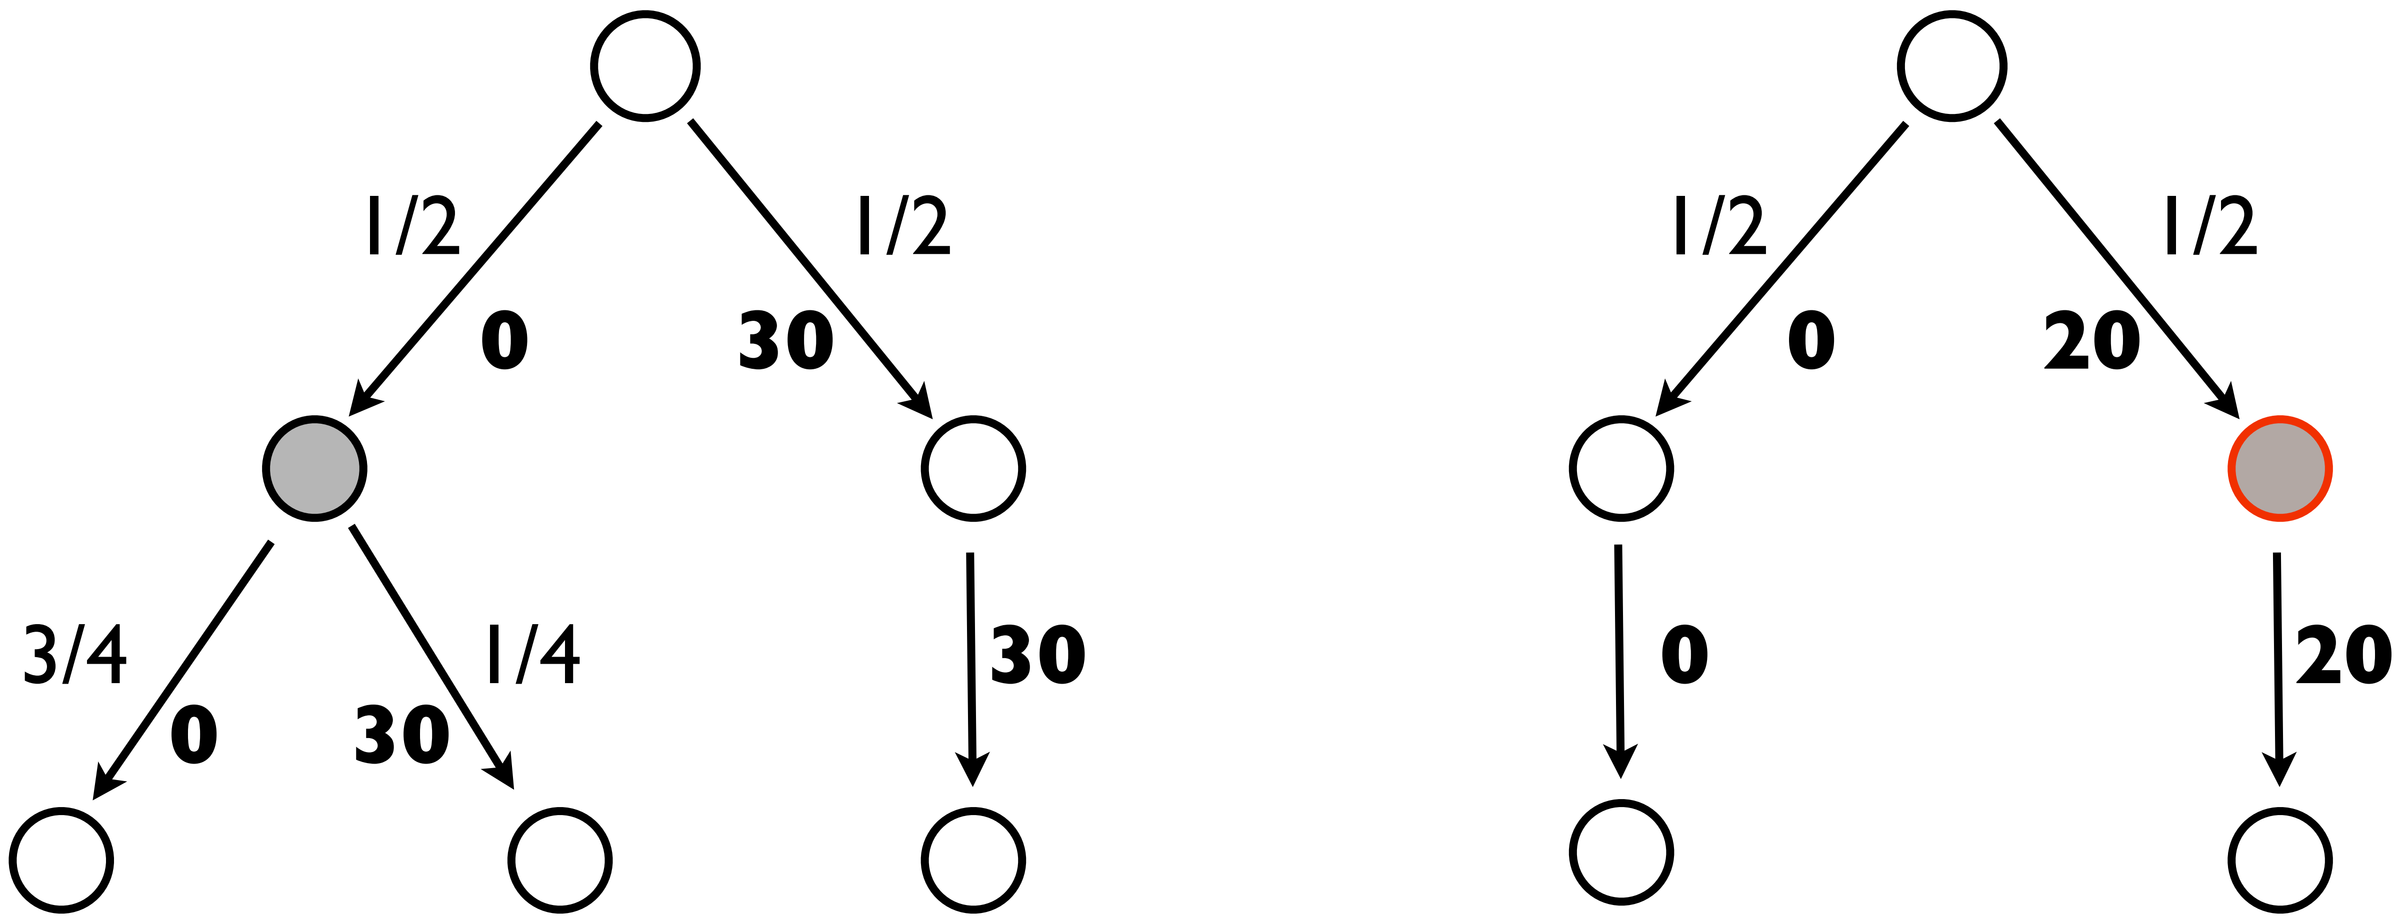

Dynamic-RM for MABs
[Cavallo, 08]

- Modify dynamic-VCG by adding the following payments *to* the agents each period:
  - For agent `i` receiving item: `(1-γ)/n` times the expected total discounted revenue that would result if `i` were ignored going forward.
  - For every other agent `j`: `1/n` times the expected *immediate* revenue that would have resulted this period if `j` were ignored.

Dynamic-RM for MABs
[Cavallo, 08]

<u>Lemma</u>: Whatever strategy an agent follows, his expected redistribution payments over time equal: a 1/n share of the expected total (over time) revenue that would result if the agent were not present.

(This is the hard part to prove. Once we have, it follows that dynamic-RM is a dynamic-Groves mechanism, and thus efficient.)

Dynamic-RM for MABs
[Cavallo, 08]

Theorem: Dynamic-RM is efficient in within-period ex post Nash equilibrium, within-period ex post IR, and *never runs a deficit*.

- And yields *significantly* more value for the agents than dynamic-VCG.
- Examples with three or more agents are tough to illustrate, so let's just look at aggregate results:

Value retained: normal distribution

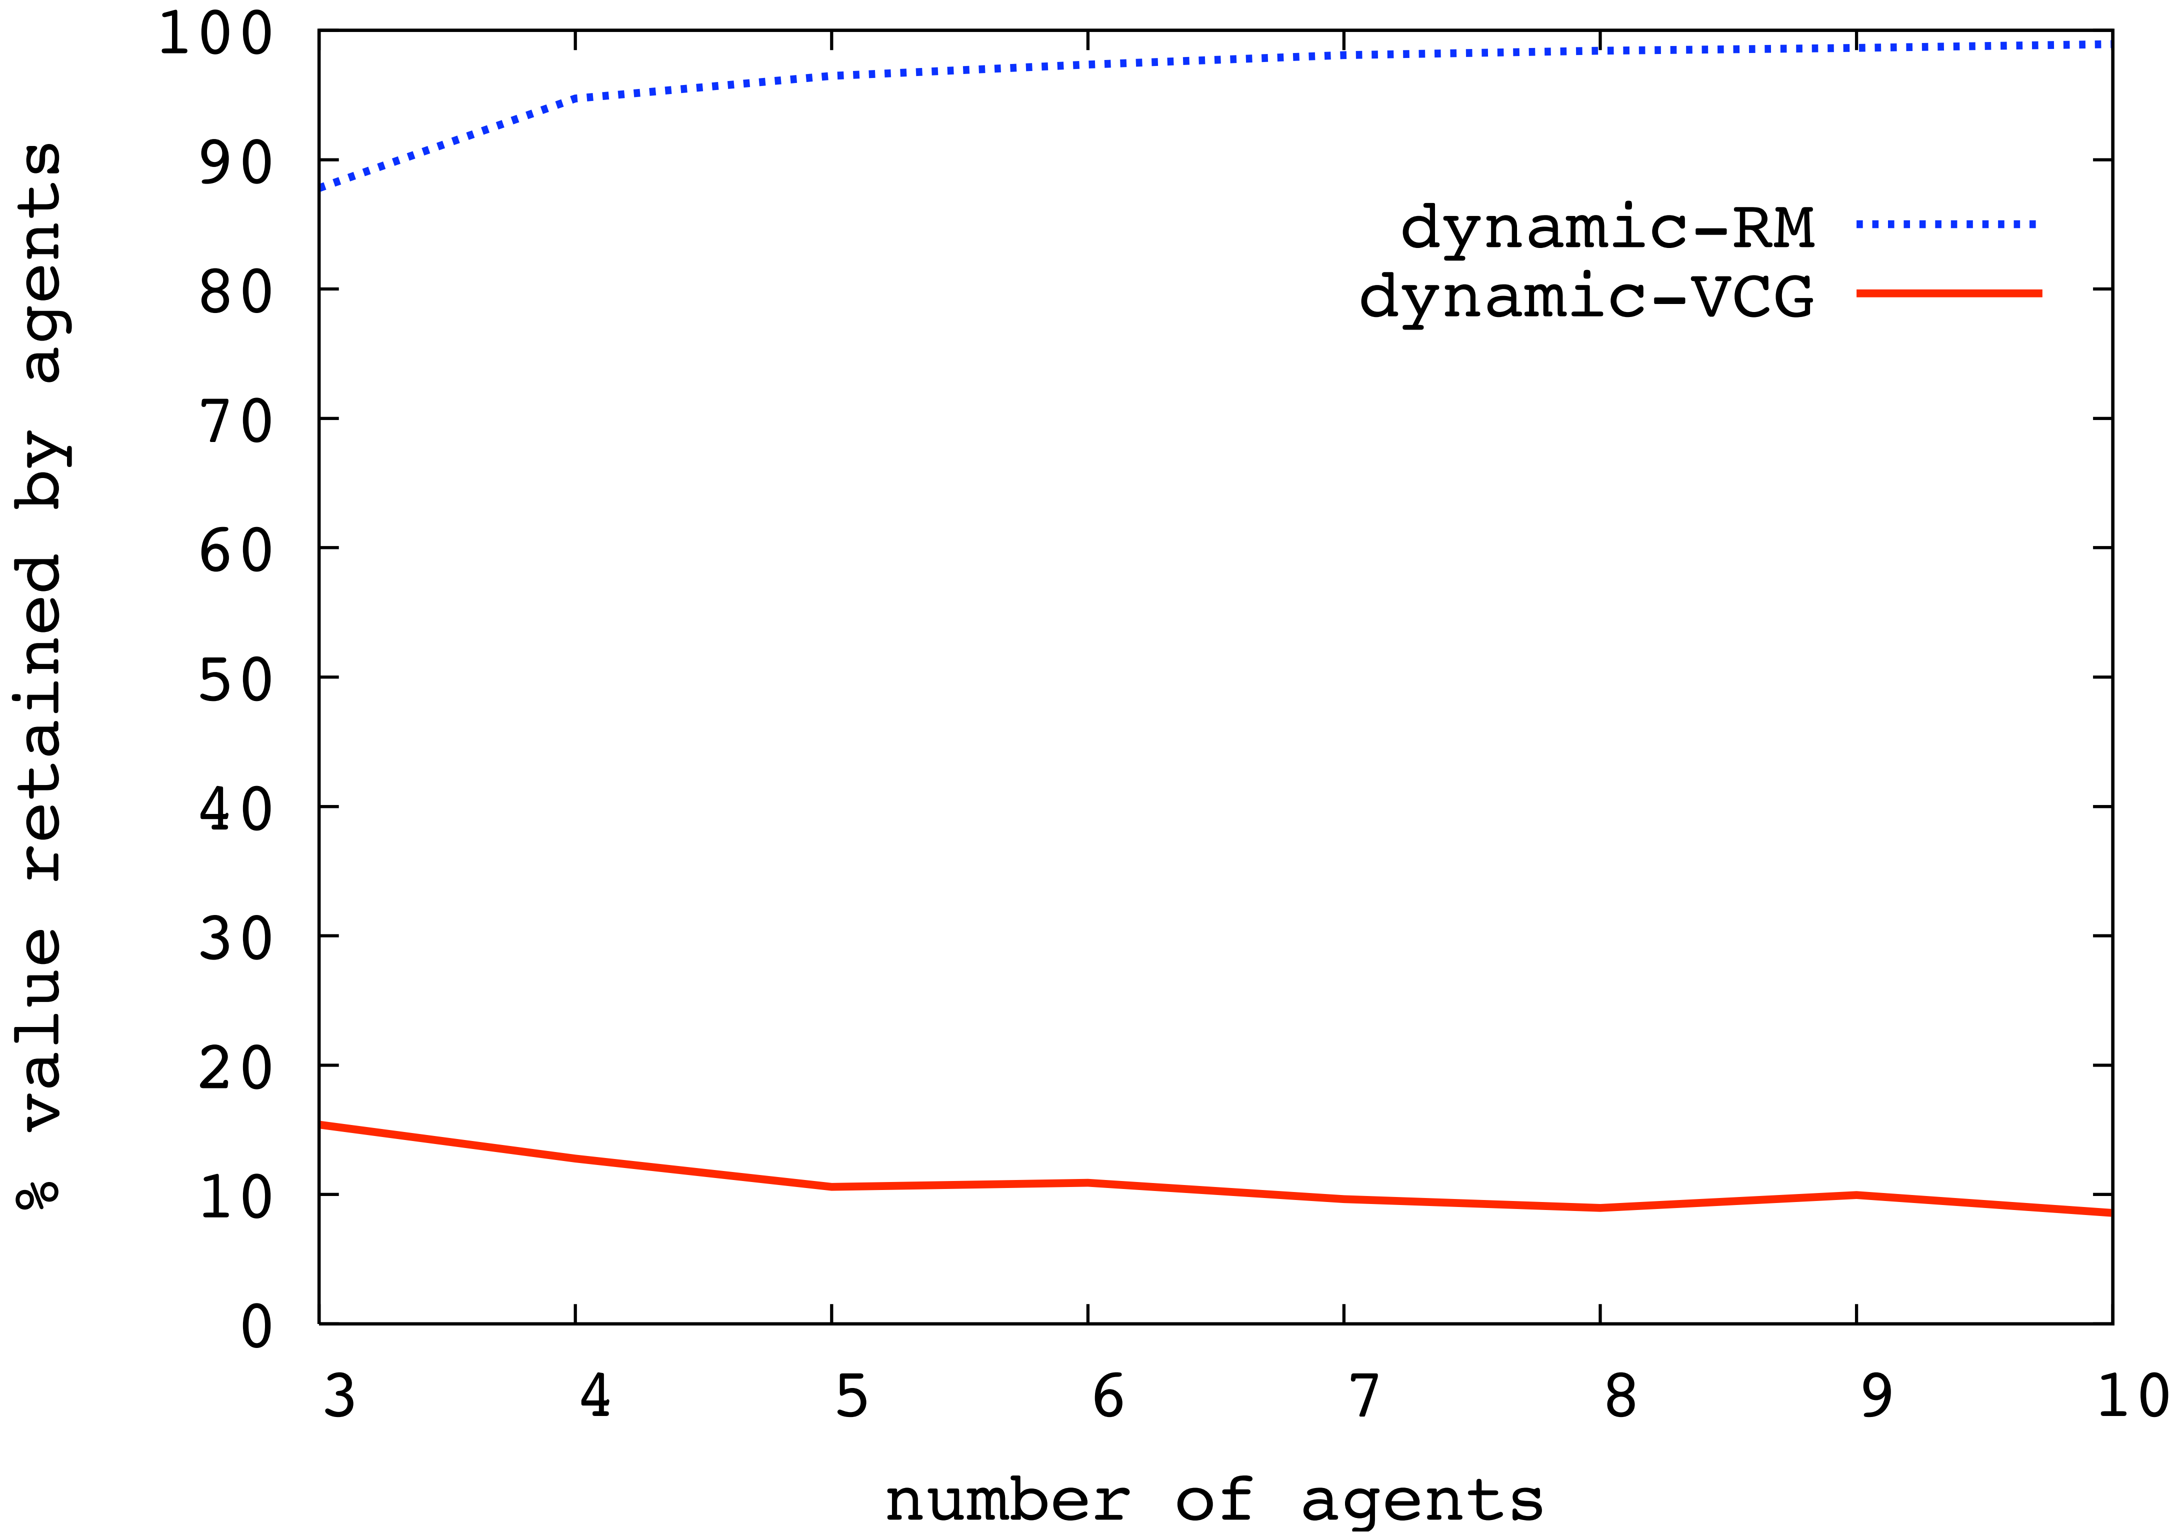

Value retained: uniform distribution

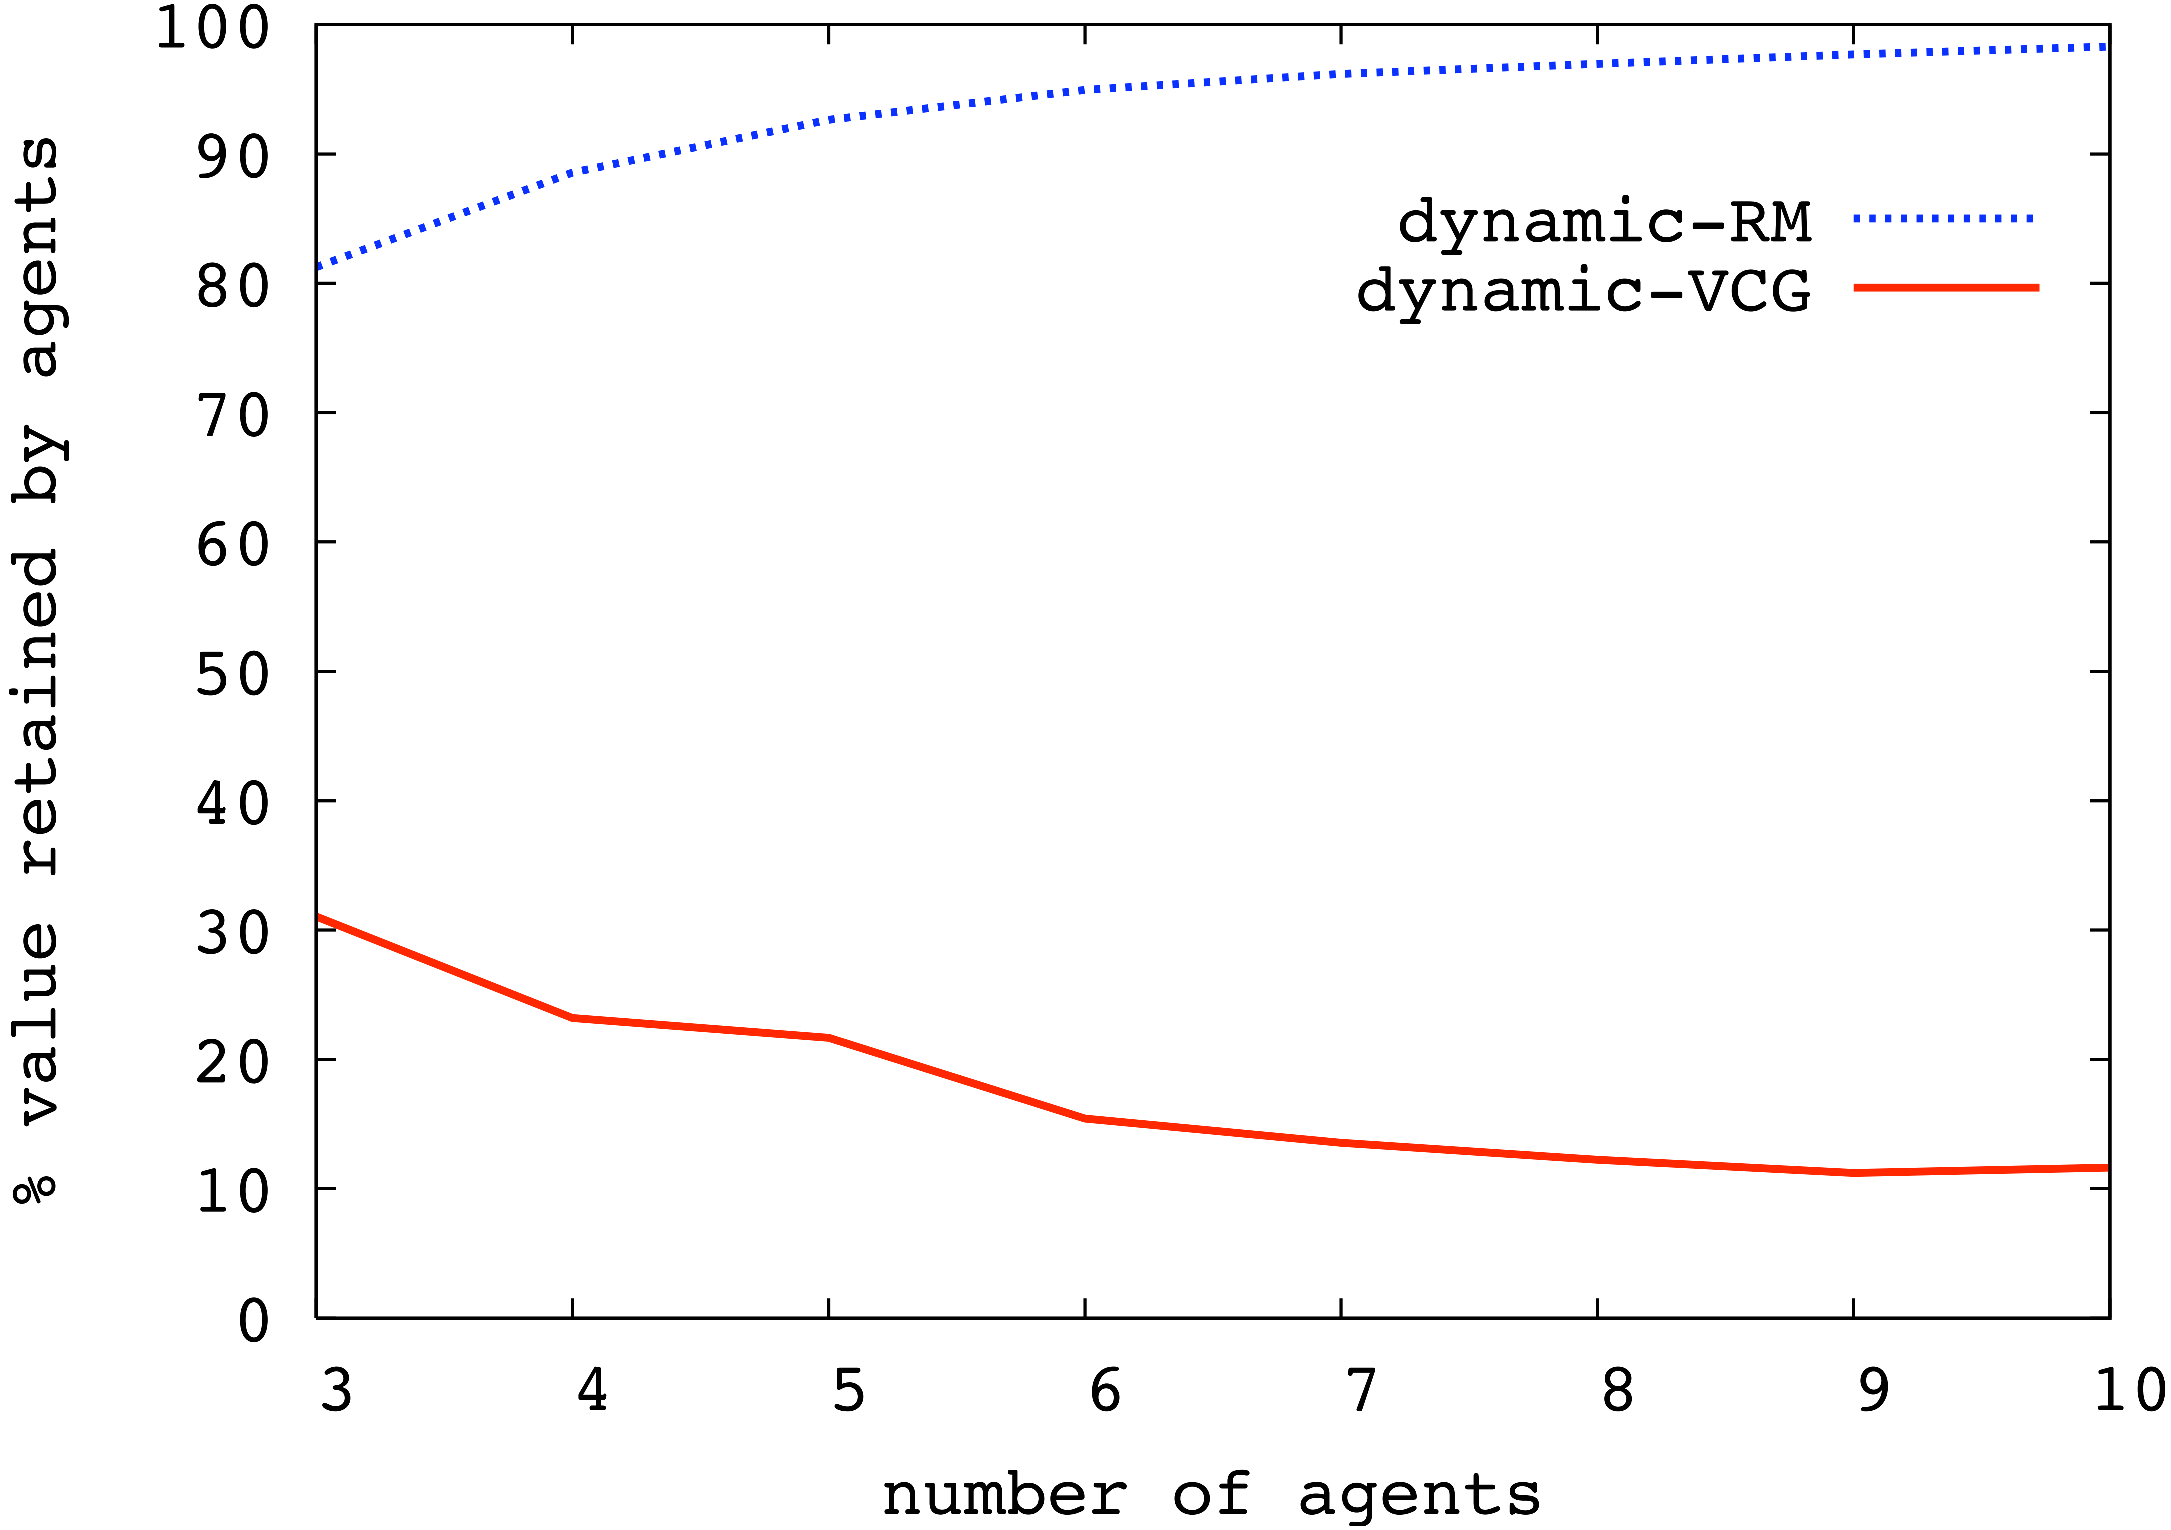

	efficiency	**IR**	**budget-balance**
team mechanism	w.p. ex post	w.p. ex post	huge deficit
dynamic-EAC	w.p. ex post	ex ante	ex ante no-deficit
dynamic-VCG	w.p. ex post	w.p. ex post	ex post no-deficit
dynamic-RM (only for MABs)	w.p. ex post	w.p. ex post	ex post no-deficit, much closer to perfect BB
balanced-mechanism	Bayes-Nash	ex ante	perfect

(Balanced team mechanism presented by Susan Athey)

Extensions

Dynamically changing populations of agents

- What's new: agents may – either temporarily or permanently – become "inaccessible", i.e., unable to communicate with the center or make/receive payments.
- Generalizes arrival/departure dynamics.

For instance:

- Imagine selling theater tickets to tourists who plan to see multiple shows over a period of days.
    - New tourists always arriving, others leaving (dynamic population).
    - A tourist may see a show, realize she likes the theater more/less (dynamic types).

Related area: *online* mechanism design

- Dynamic population (arrivals and departures), but *static* types – all private information an agent will ever obtain can be reported in arrival period.
  - [Friedman & Parkes, 03]
  - [Parkes & Singh, 03]
  - [Lavi & Nisan, 04]
  - [Porter, 04]

Online-VCG mechanism

[Parkes & Singh, 03]

- Collects a single payment from each agent in her "arrival period".
  - Within-period ex post efficient.
  - Ex post individual rational
  - Ex post no-deficit.

Online-VCG mechanism

[Parkes & Singh, 03]

- Collects a single payment from each agent in her "arrival period".
    - Within-period ex post efficient.
    - Ex post individual rational
    - Ex post no-deficit.

➡ But only for *static* types.

Dynamic populations, dynamic types
[Cavallo, Parkes, & Singh, 07]

- Unifies dynamic mechanism design and online mechanism design.
- The new challenges:
  - Optimal policy must consider accessibility/ inaccessibility dynamics
  - Agents may not be available for payment while still exerting influence on welfare of other agents.

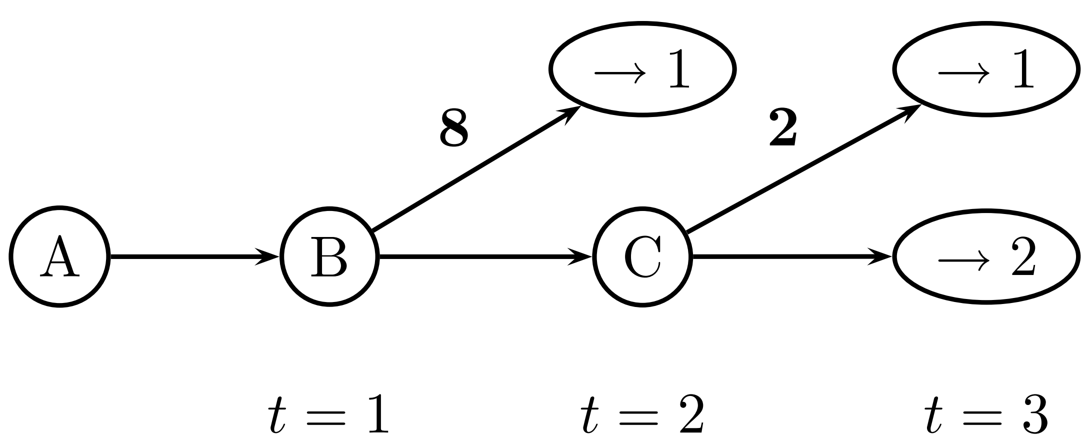

(a) Agent 1's Type.

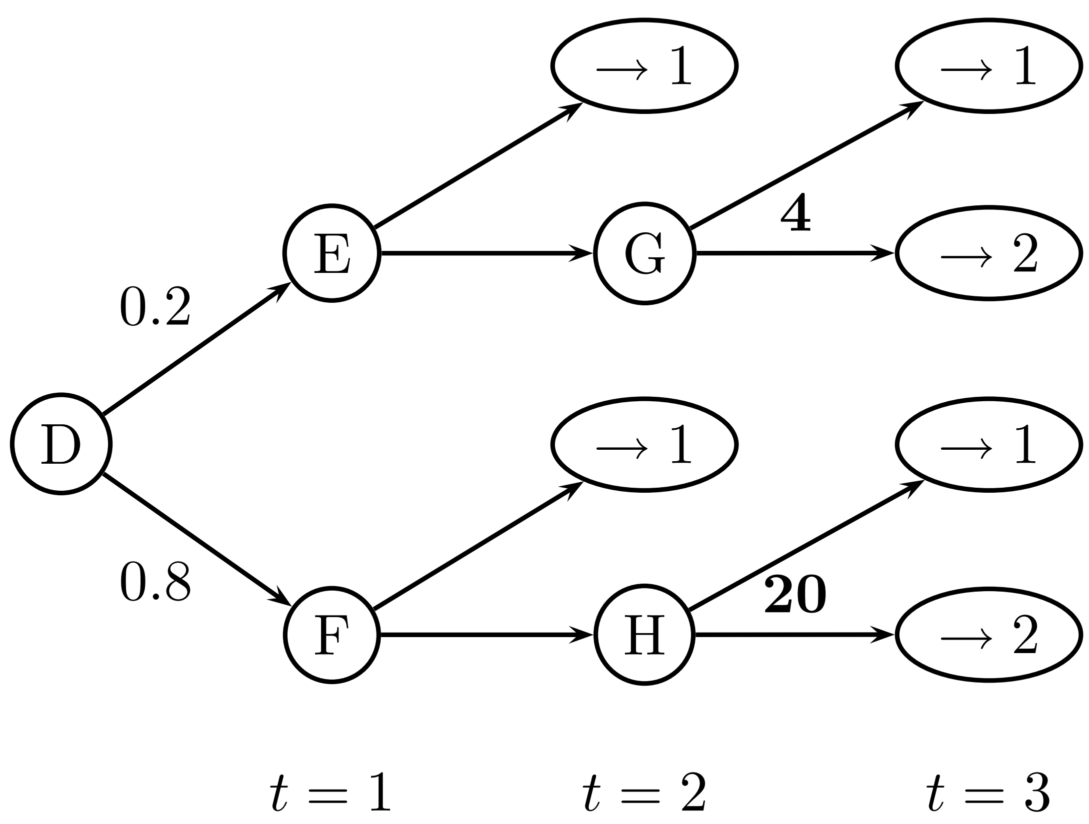

(b) Agent 2's Type.

- Imagine agent 1 accessible at t = 1, and agent 2 inaccessible at t = 1 but very likely to become accessible at t = 2.

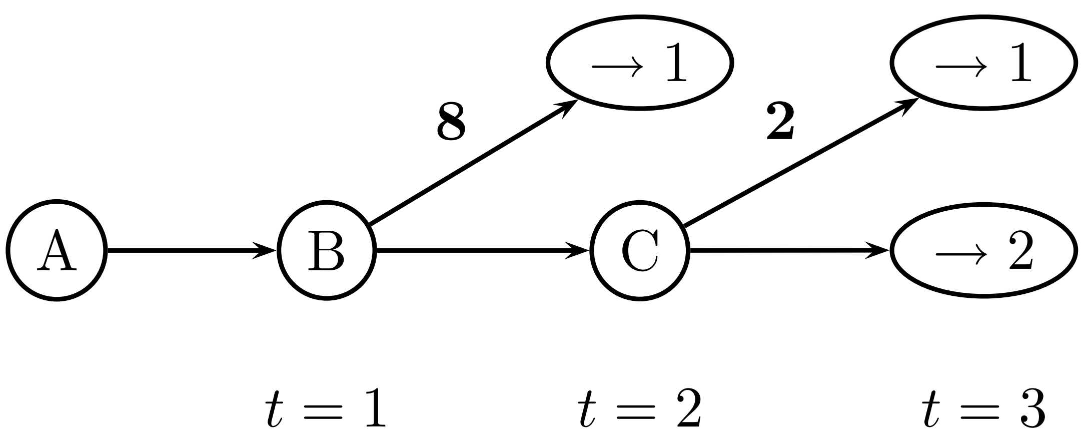

(a) Agent 1's Type.

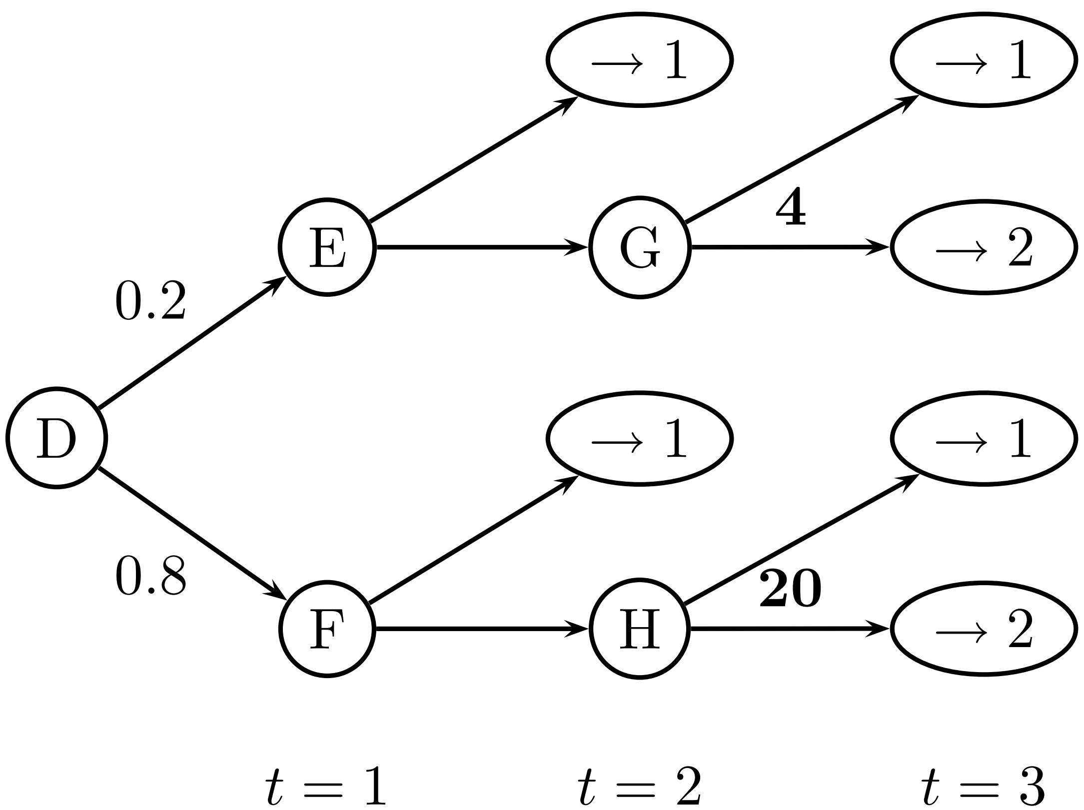

(b) Agent 2's Type.

- Imagine agent 1 accessible at t = 1, and agent 2 inaccessible at t = 1 but very likely to become accessible at t = 2.
- In "naive" dynamic-VCG mechanism, agent 1 better off "hiding" to improve social-welfare.

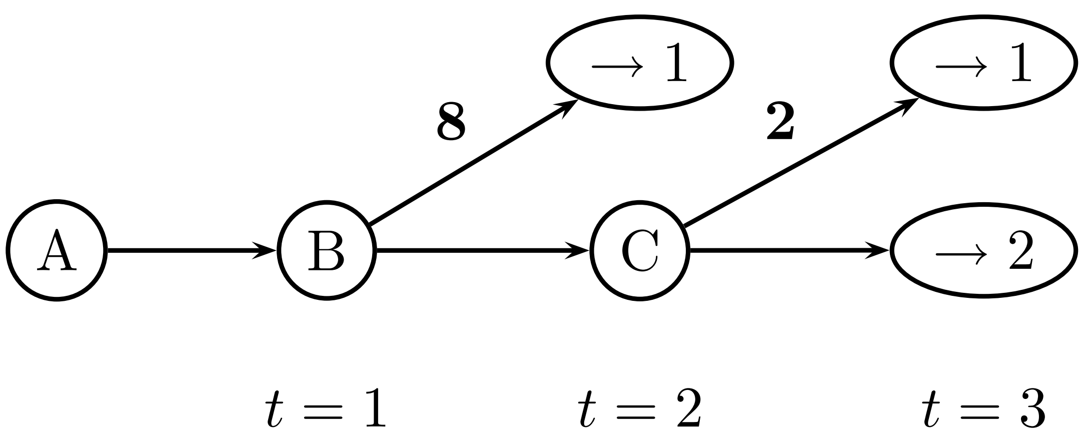

(a) Agent 1's Type.

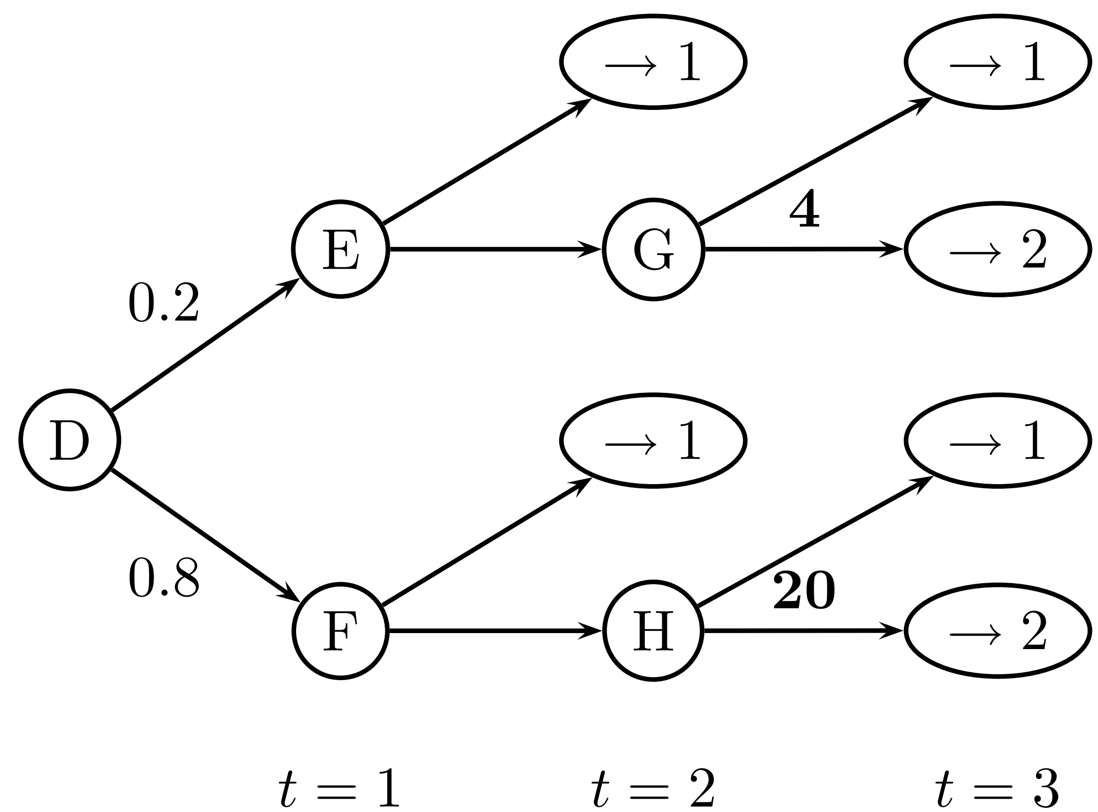

(b) Agent 2's Type.

- Imagine agent 1 accessible at t = 1, and agent 2 inaccessible at t = 1 but very likely to become accessible at t = 2.
- In "naive" dynamic-VCG mechanism, agent 1 better off "hiding" to improve social-welfare.
- In non-naive mechanism that makes dynamic-VCG payments only to accessible agents, agent 2 can benefit by hiding.

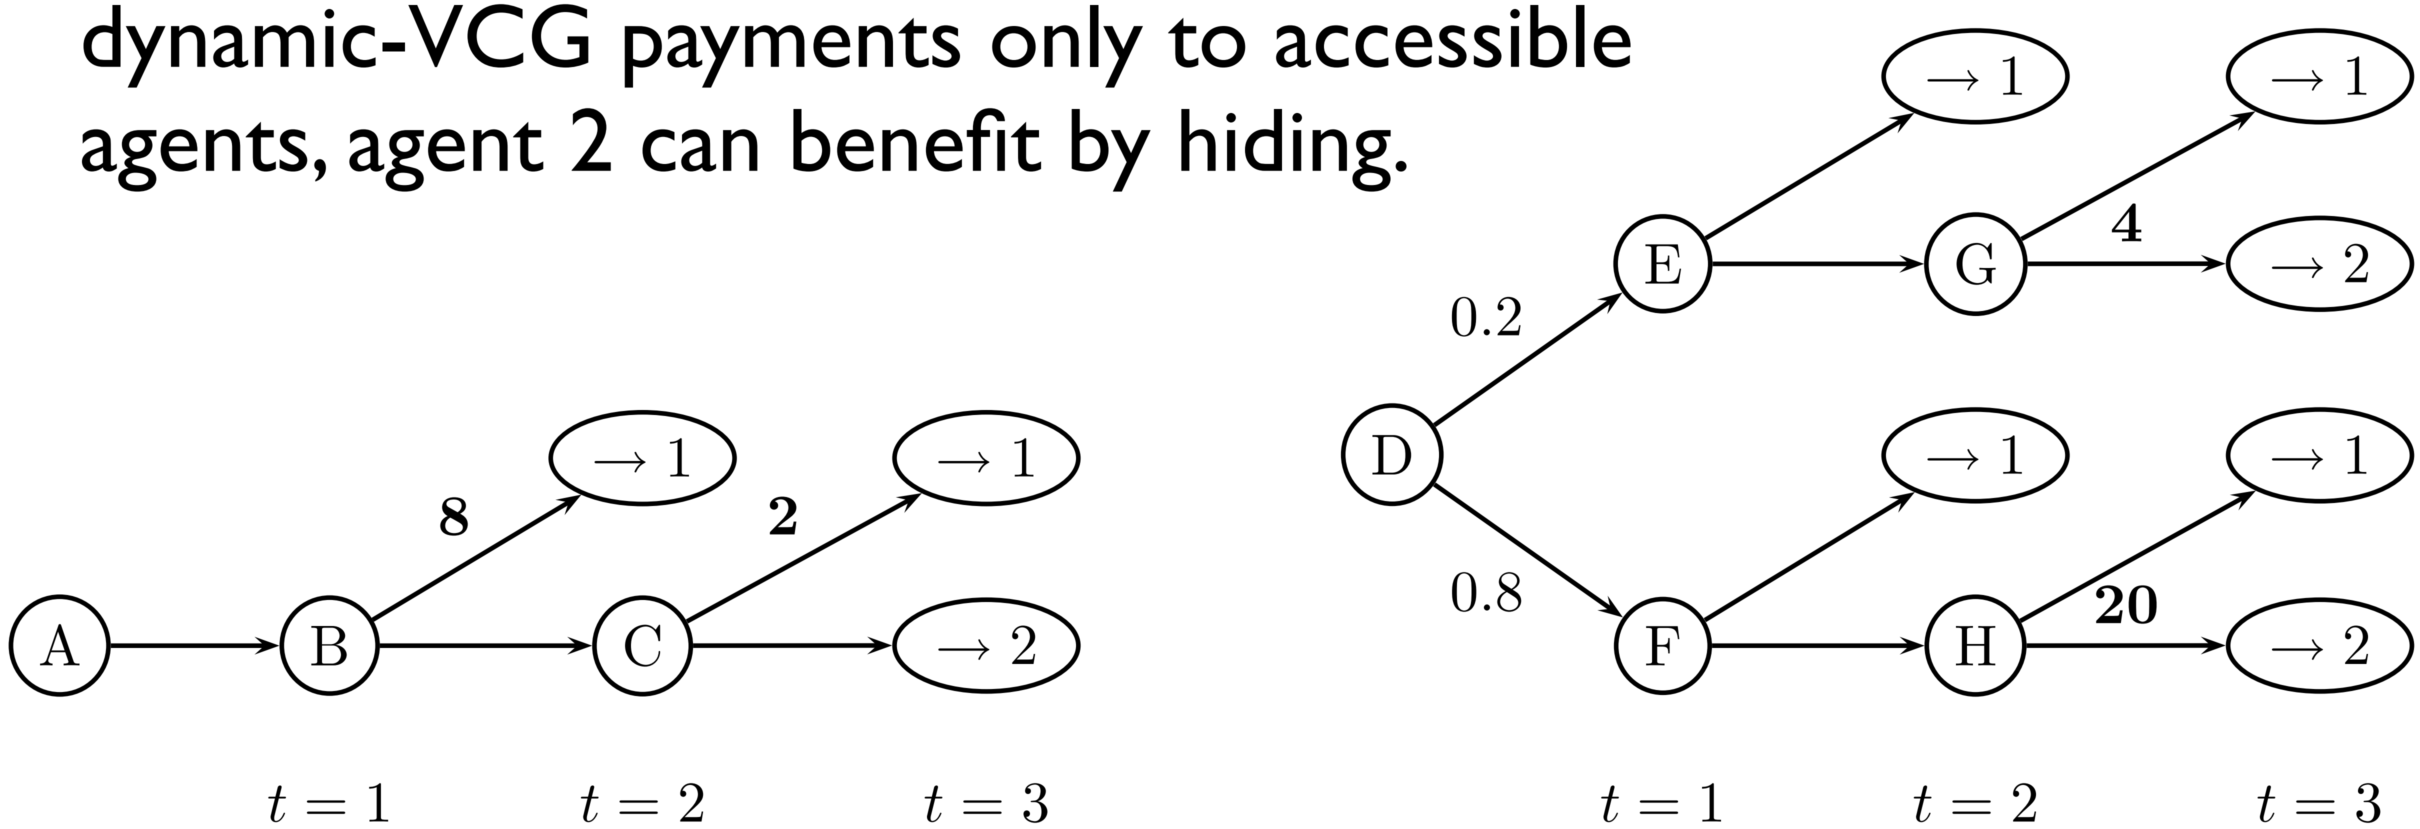

(a) Agent 1's Type.

(b) Agent 2's Type.

A fix

- For any inaccessible agent, keep log of payments dynamic-VCG would impose on agent; when the agent becomes accessible, execute "lump sum" payment, appropriately scaled for discounting.
- Requires that all agents eventually "come back".

Imagine both agents accessible in all periods. Should agent 2 feign inaccessibility until t = 2?

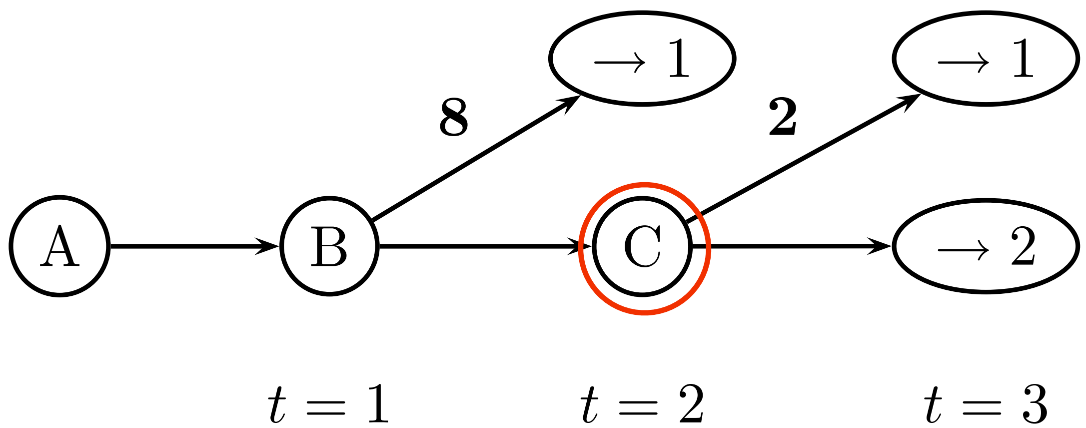

(a) Agent 1's Type.

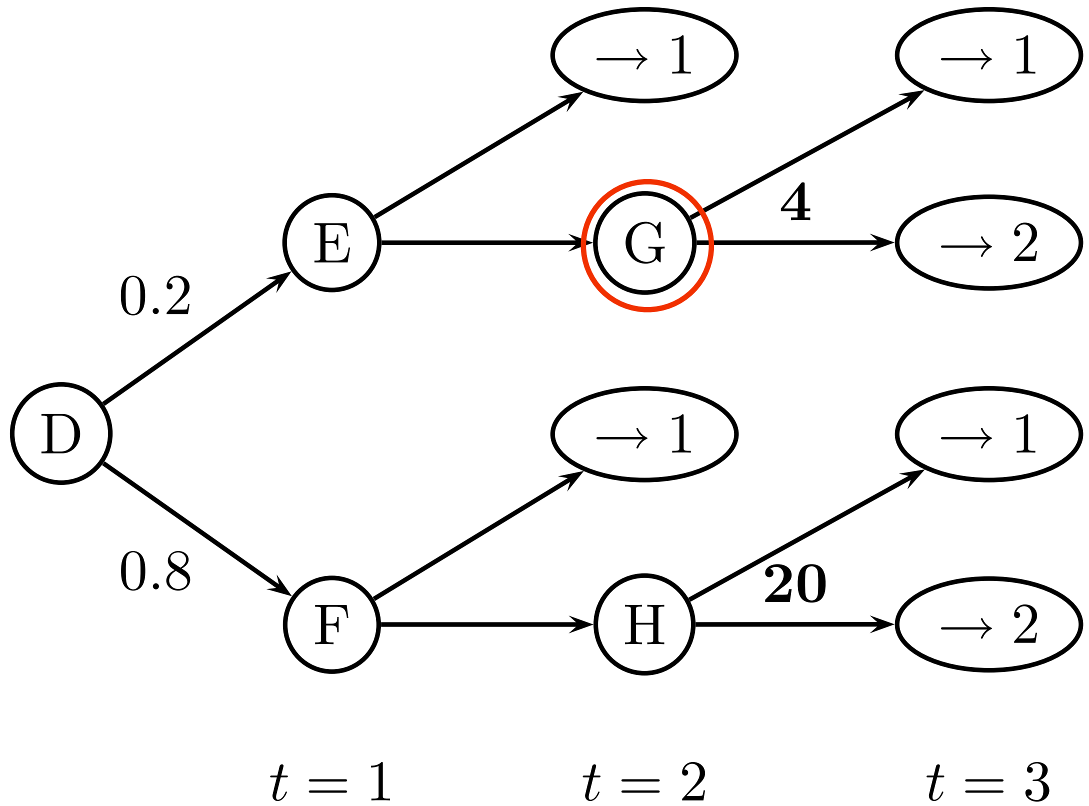

(b) Agent 2's Type.

Imagine both agents accessible in all periods. Should agent 2 feign inaccessibility until t = 2?

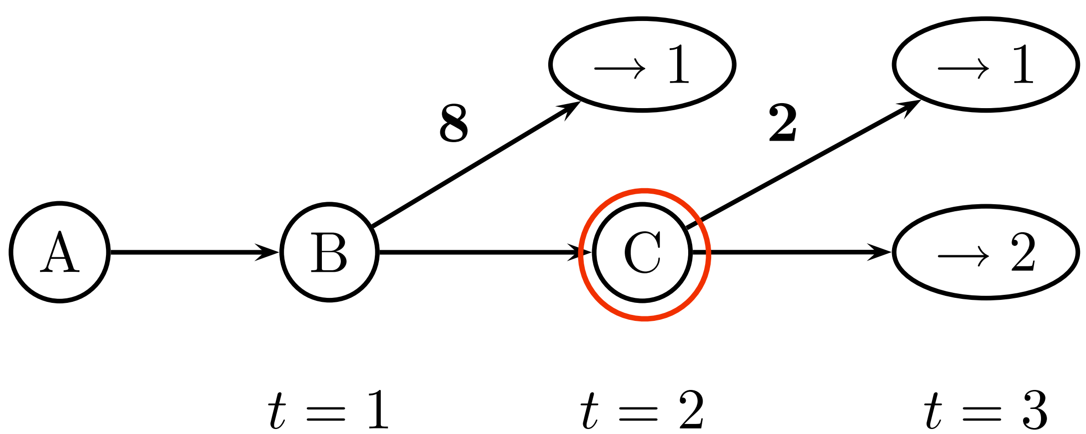

(a) Agent 1's Type.

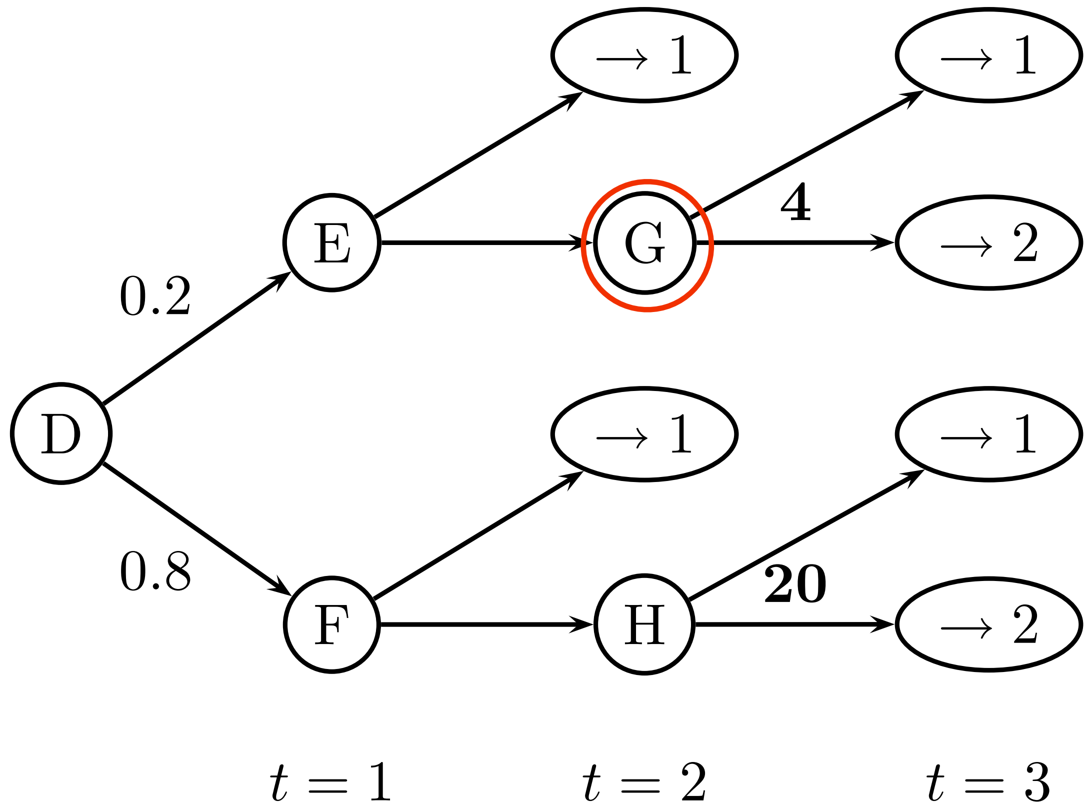

(b) Agent 2's Type.

T_2 = -6 - 2 = -8, same whether he hides at t = 1 or not.

Imagine both agents accessible in all periods. Should agent 2 feign inaccessibility until t = 2?

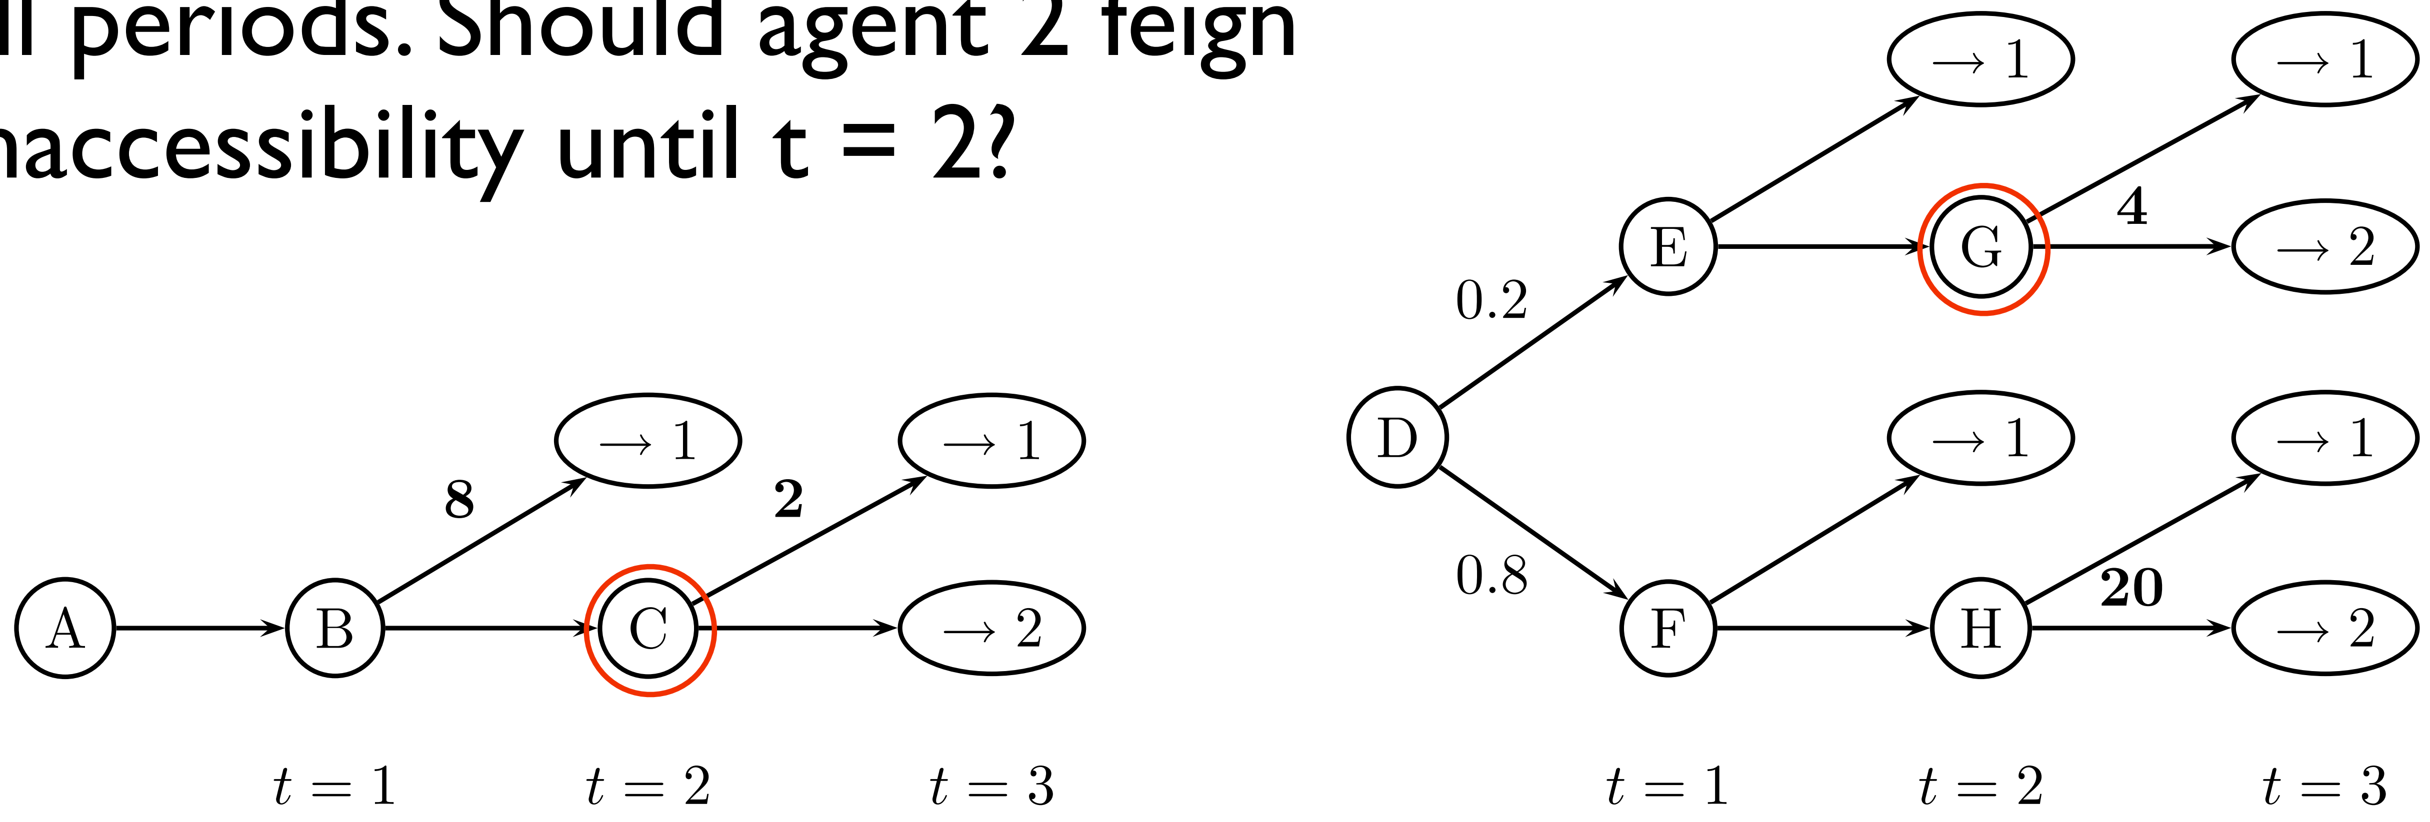

(a) Agent 1's Type. (b) Agent 2's Type.

T_2 = -6 - 2 = -8, same whether he hides at t = 1 or not.

difference in optimal value for agent 1, with and without agent 2 present at t=1

difference in optimal value for agent 1, with and without agent 2 present at t=2

What if agents *don't* always come back?

- In general, the scheme won't work.
- For an arrival/departure model, within-period ex post efficiency is recovered if agent arrivals are independent conditioned on actions chosen.

Randomly arriving agents, revenue maximization
[Gershkov and Moldovanu, 09]

- Goal is not efficiency, but rather revenue maximization.
- Agents arrive randomly over time.
- Set of resources to be allocated before a deadline.

Computation

The big bad secret

- Computing optimal policies is, in general, *very* hard... but often necessary.
- What can we do?
  - Approximations, yielding approximate equilibria (even this is hard)?
  - Identify tractable special cases.
  - Thankfully, MABs are such a case.

Computing optimal policies in MABs [Gittins & Jones, 74]

- At each period, compute Gittins index for each agent's Markov chain.
- "Activate" (e.g., allocate resource to) agent with highest index.

→ Complexity: Gittins indices are independent, so *linear in number of agents*.

Beyond simple repeated allocation

Coordination of value information acquisition preceding one-time allocation of a single item (“metadeliberation auctions”).

Metadeliberation auction
[Cavallo & Parkes, 08]

- A resource is to be allocated. Agents have initial valuations for the resource. Valuations can potentially be increased by costly “deliberation” (e.g., researching new ways of using the resource).
- How to coordinate deliberation/allocation to maximize social welfare?

Metadeliberation auction
[Cavallo & Parkes, 08]

- Given optimal policy, dynamic-VCG mechanism can be applied to deal with incentives.
- Computing optimal deliberation/allocation policy is tractable (reduction to multi-armed bandits problem).
- **Note**: even in this one-time allocation scenario, a realistic analysis of the problem reveals the need for *dynamic* solution.

Computation

Beyond bandits: heuristics for special cases

Self-correcting dynamic multi-unit auctions
[Constantin & Parkes, here]

- When computing optimal policy is infeasible...
- Propose heuristic method that is strategyproof, yet achieves social-welfare ~90% of optimal.
- See talk tomorrow for details.

Auctions with online supply

[Babaioff, Blumrosen, & Roth, workshop here]

- Dynamically arriving items – unknown total quantity.
- Approximate mechanisms.
- Nonetheless truthful – possible due to the restricted setting.

Open problems, future directions

Interdependent values

- Interestingly, sequential nature of problem kind of *helps* here: ex post payments become natural.
  - Version of the team mechanism is still within-period ex post efficient.
  - But no apparent way to extend dynamic-VCG...
  - Can we achieve no-deficit, IR, and efficienct in interdependent settings?

Computation

- The general case looks hopeless.
- Continue to identify tractable special cases?
- Adopt more realistic equilibrium notions?

References

- [Vickrey, 61] William Vickrey. Counterspeculations, auctions, and competitive sealed tenders. Journal of Finance, 16:8–37, 1961.
- [Clarke, 71] Edward Clarke. Multipart pricing of public goods. Public Choice, 8:19–33, 1971.
- [Groves, 73] Theodore Groves. Incentives in teams. Econometrica, 41:617–631, 1973.
- [Green & Laffont, 77] Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. Econometrica, 45:427–438, 1977.
- [Holmstrom, 79] Bengt Holmstrom. Groves' scheme on restricted domains. Econometrica, 47(5):1137–1144, 1979.

- [Arrow, 79] Kenneth J. Arrow. The property rights doctrine and demand revelation under incomplete information. In M. Boskin, editor, Economics and Human Welfare. Academic Press, 1979.

- [d'Aspremont & Gerard-Varet, 79] C. D'Aspermont and L.A. Gerard-Varet. Incentives and incomplete information. Journal of Public Economics, 11:25–45, 1979.

- [Bailey, 97] Martin J. Bailey. The demand revealing process: To distribute the surplus. Public Choice, 91:107–126, 1997.

- [Cavallo, 06] Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'06), pages 882–889, 2006.

- [Guo & Conitzer, 07] Mingyu Guo and Vincent Conitzer. Worst-case optimal redistribution of VCG payments. In Proceedings of the 8th ACM Conference on Electronic Commerce (EC-07), San Diego, CA, USA, pages 30–39, 2007.

- [Moulin, 2007] Herv´e Moulin. Efficient, strategy-proof and almost budget-balanced assignment. unpublished, 2007.

- [Guo & Conitzer, 08] Mingyu Guo and Vincent Conitzer. Optimal-in-expectation redistribution mechanisms. In Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-08), 2008.

- [Hartline & Roughgarden, 2008] Jason D. Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In Proceedings of the 40th annual ACM symposium on Theory of Computing (STOC'08), 2008.

- [de Clippel, Naroditskiy, Greenwald, here]. Geoffroy de Clippel, Victor Naroditskiy, and Amy Greenwald. Destroy to Save. In Proceedings of the 10th ACM Conference on Electronic Commerce (EC-09), 2009.

- [Myerson, 1986] Roger Myerson. Multistage games with communication. Econometrica, 54(2):323–358, 1986.

- [Athey & Segal, 07] Susan Athey and Ilya Segal. An efficient dynamic mechanism. Working paper, http://www.stanford.edu/~isegal/agv.pdf, 2007.

- [Cavallo, Parkes, & Singh, 06] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Optimal coordinated planning amongst self-interested agents with private state. In Proceedings of the Twenty-second Annual Conference on Uncertainty in Artificial Intelligence (UAI'06), 2006.

- [Cavallo, 08] Ruggiero Cavallo. Efficiency and redistribution in dynamic mechanism design. In Proceedings of the 9th ACM Conference on Electronic Commerce (EC-08) (to appear), 2008.
- [Bergemann & Valimaki, 08] Dirk Bergemann and Juuso Valimaki. Efficient dynamic auctions. Cowles Foundation Discussion Paper 1584, http://cowles.econ.yale.edu/P/cd/d15b/d1584.pdf, 2006.
- [Cavallo, Parkes, & Singh, 07] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Online mechanisms for persistent, periodically inaccessible self-interested agents. In DIMACS Workshop on the Boundary between Economic Theory and Computer Science, 2007.
- [Friedman & Parkes, 03] E. Friedman and D. C. Parkes. Pricing WiFi at Starbucks– issues in online mechanism design. In Proc. Fourth ACM Conference on Electronic Commerce (EC'03), pages 240–241, 2003.
- [Parkes & Singh, 03] David C. Parkes and Satinder Singh. An MDP-based approach to Online Mechanism Design. In Proceedings of the 17th Annual Conf. on Neural Information Processing Systems (NIPS'03), 2003.
- [Lavi & Nisan, 04] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. Theoretical Computer Science, 310:159–180, 2004. Earlier version in ACMEC 2000.

- [Porter, 04] Ryan Porter. Mechanism design for online real-time scheduling. In Proceedings of the ACM Conference on Electronic Commerce (EC'04), pages 61–70, 2004.
- [Gershkov and Moldovanu, 09] Alex Gershkov and Benny Moldovanu. Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach. Forthcoming in American Economic Journal: Microeconomics.
- [Gittins & Jones, 74] J. C. Gittins and D. M. Jones. A dynamic allocation index for the sequential design of experiments. In In Progress in Statistics, pages 241–266. J. Gani et al., 1974.
- [Cavallo & Parkes, 08] Ruggiero Cavallo and David C. Parkes. Efficient metade-liberation auctions. In Proceedings of the 26th Annual Conference on Artificial Intelligence (AAAI-08), 2008.
- [Constantin & Parkes, here] Florin Constantin and David C. Parkes. Self-Correcting Sampling-Based Dynamic Multi-Unit Auctions. In the 10th ACM Electronic Commerce Conference (EC'09), 2009
- [Babaioff, Blumrosen, & Roth, workshop here] Moshe Babaioff, Liad Blumrosen, and Aaron Roth. Auctions with Online Supply. In Fifth Workshop on Ad Auctions, 2009.