Example:
Planning in High-Stake Decision Situations

Planning with Nonlinear Utility Functions
An Actual High-Stake Decision Problem

50%	10,000,000 dollars	50%	4,000,000 dollars
50%	0 dollars	50%	5,000,000 dollars

Planning with Nonlinear Utility Functions
An Actual High-Stake Decision Problem

Who Wants to be a Millionaire?

The term "computer bug" was coined when an insect of which kind caused a computer to crash

~~a) centipede~~
b) fly
~~c) spider~~
d) moth

100%	$500,000	50%	$32,000
		50%	$1,000,000

Nonlinear Utility Functions

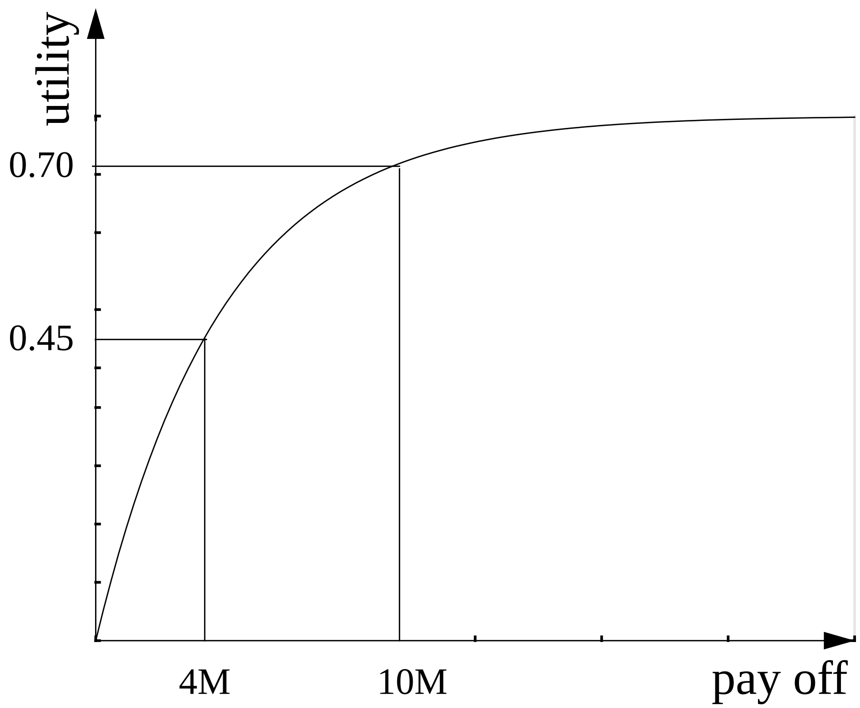

utility function is a strictly monotonically increasing function of the pay-off
[Bernoulli]
[von Neumann and Morgenstern]

here:
risk-averse
decision maker

50%	10,000,000 dollars	50%	4,000,000 dollars
50%	0 dollars	50%	5,000,000 dollars
average	5,000,000 dollars	average	4,500,000 dollars
50%	0.70 (utility)	50%	0.45 (utility)
50%	0.00 (utility)	50%	0.55 (utility)
average	0.35 (utility)	average	0.50 (utility)

Planning with Nonlinear Utility Functions
Nonlinear Utility Functions

the more risk-seeking a decision maker is
the more important is the best-case reward (best-case cost)

50%	10,000,000 dollars
50%	0 dollars

decrease in variance
decrease in average reward (increase in average cost)

50%	4,000,000 dollars
50%	5,000,000 dollars

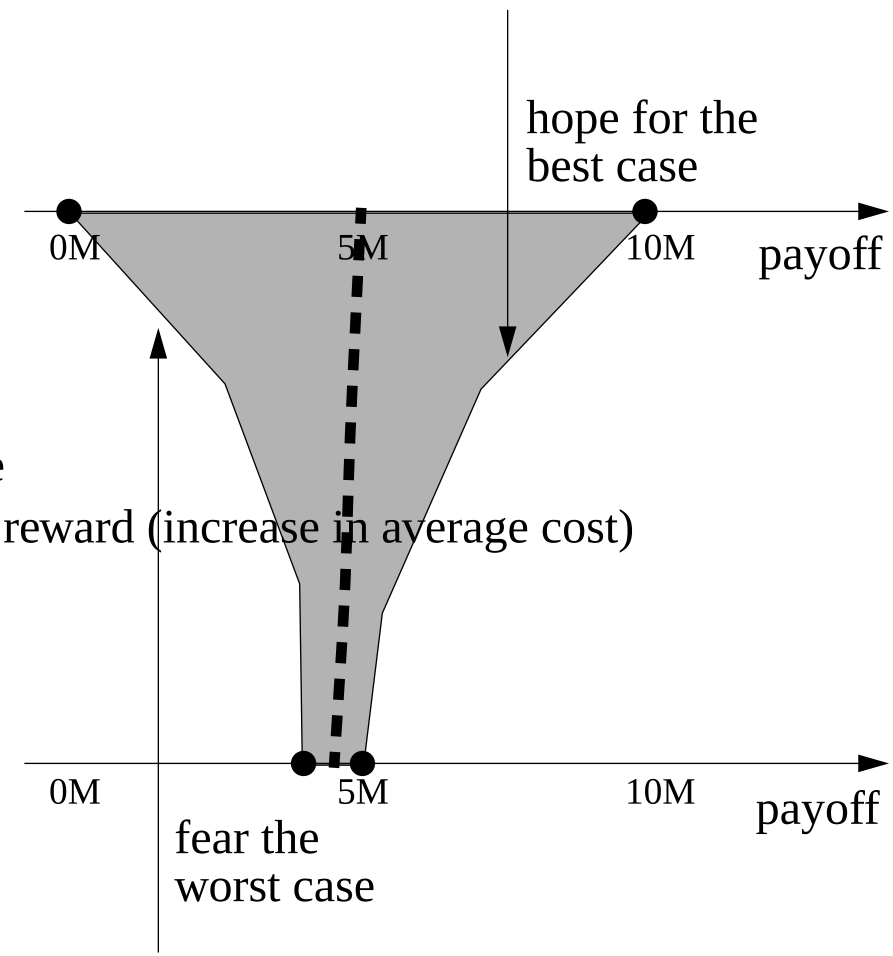

the more risk-averse a decision maker is
the more important is the worst-case reward (worst-case cost)

Planning with Nonlinear Utility Functions
Planning

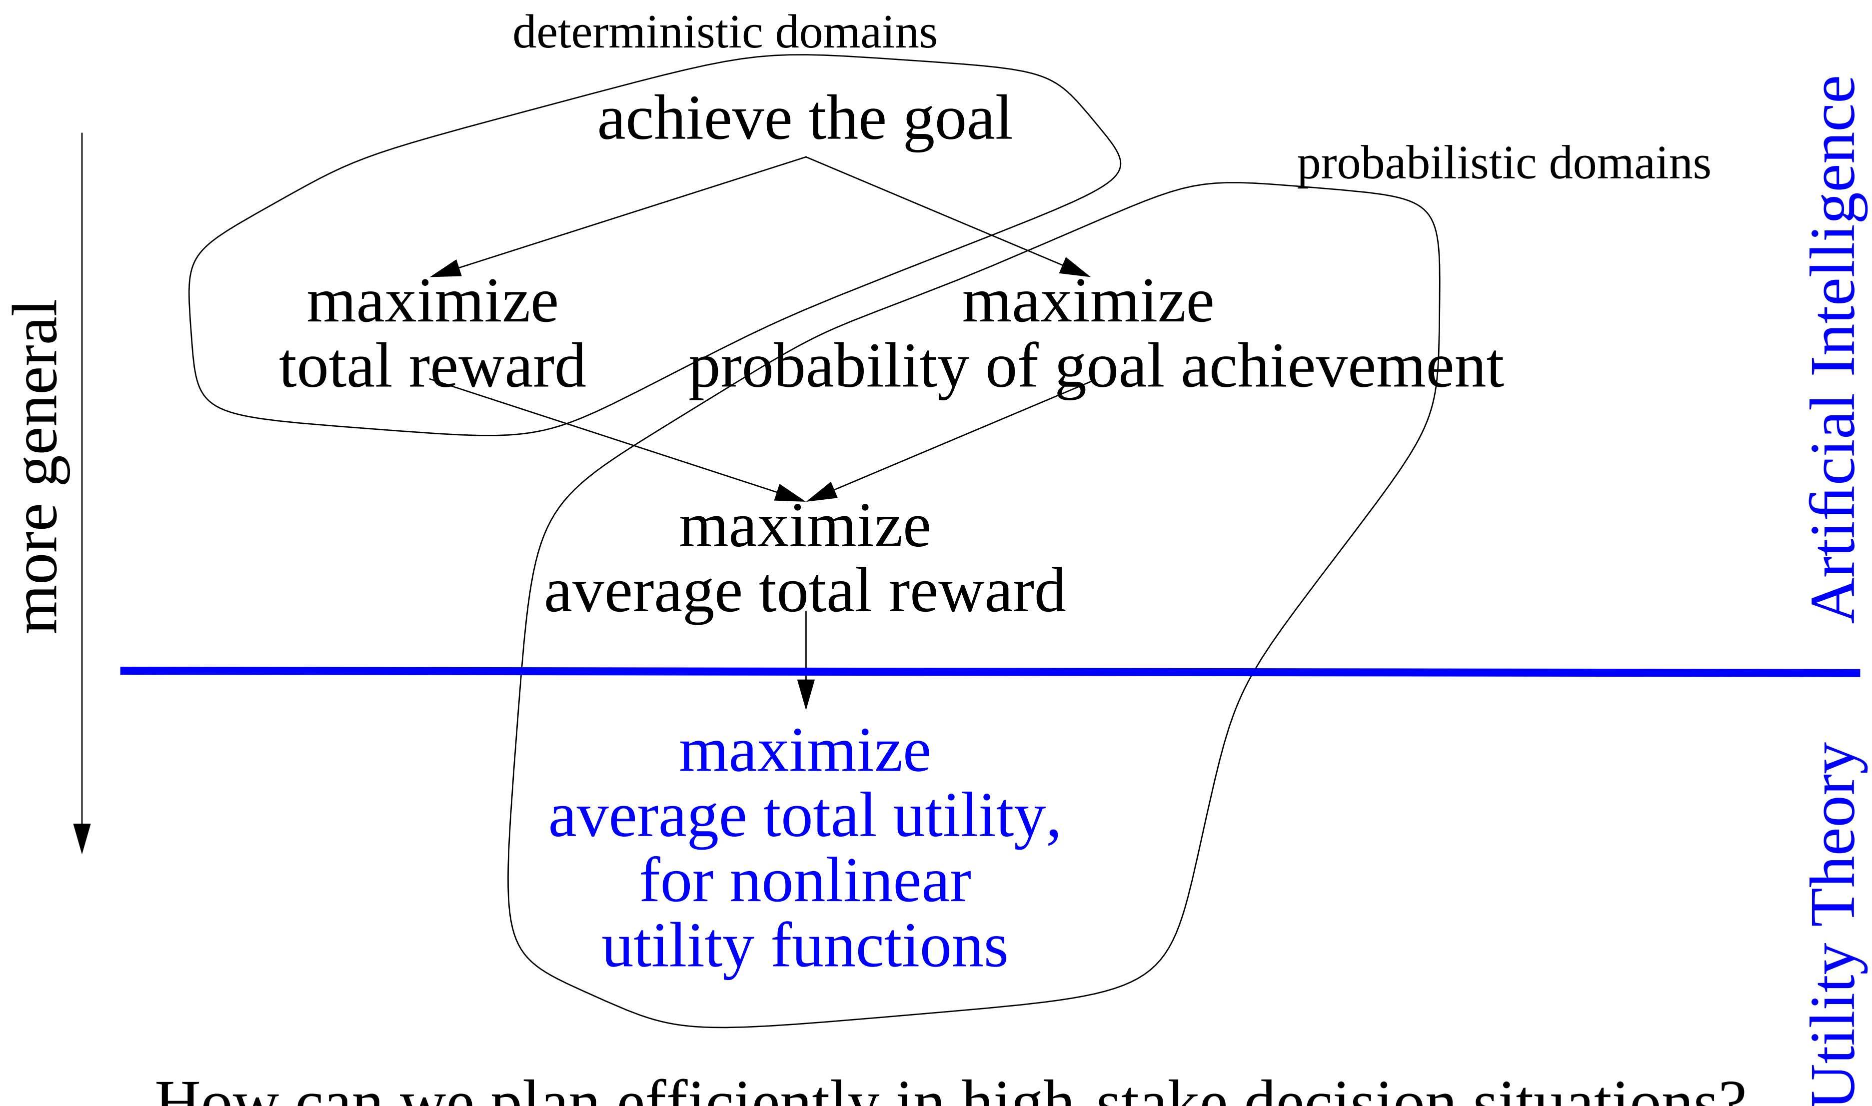

How can we plan efficiently in high-stake decision situations?

Planning with Nonlinear Utility Functions
Plan Evaluation Example

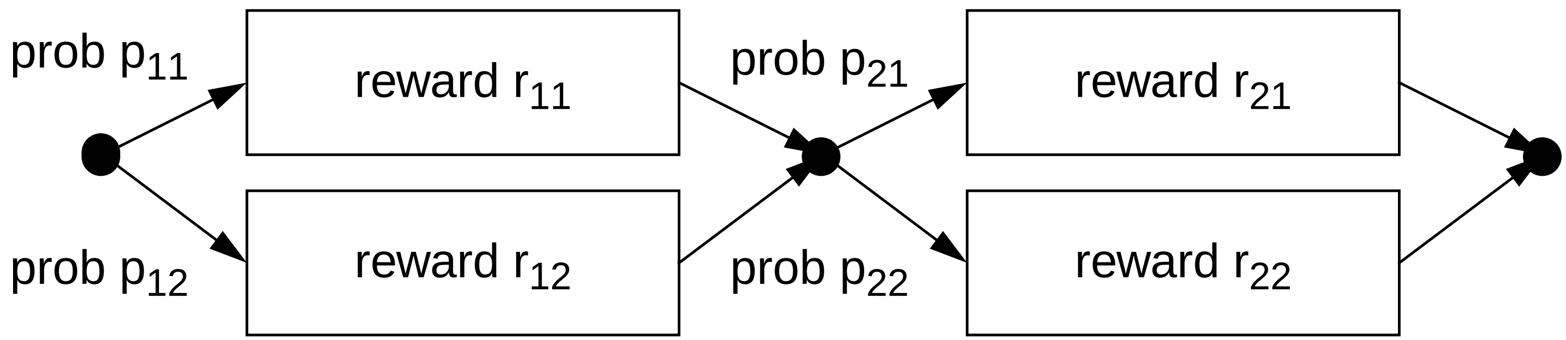

average total reward: $\sum_i \sum_j (p_{1i}\, p_{2j}\, (r_{1i} + r_{2j})) \quad = \quad \sum_i (p_{1i}\, r_{1i}) + \sum_j (p_{2j}\, r_{2j})$

average total utility: $\sum_i \sum_j (p_{1i}\, p_{2j}\, u(r_{1i} + r_{2j}))$ loss of decomposability

[Wellman, Ford, and Larson]
[Haddawy and Hanks]

Planning with Nonlinear Utility Functions
Exponential Utility Functions

= utility functions with the delta property or with constant local risk aversion

[Pratt]

[Howard and Matheson]

- exponential utility functions maintain the decomposability of planning tasks
- exponential utility functions can model a continuum of risk attitudes

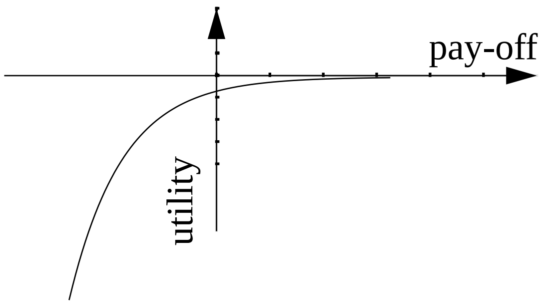

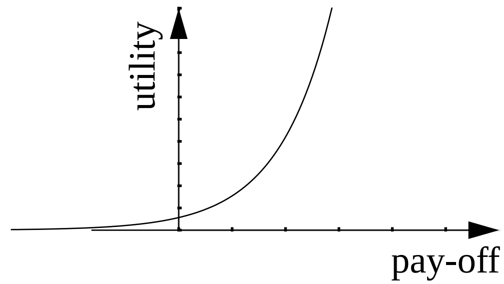

utility(pay-off) = $-\gamma^{\text{pay-off}}$

$0 < \gamma < 1$

insurance holders

utility(pay-off) = $\gamma^{\text{pay-off}}$

$\gamma > 1$

gamblers

$\gamma \cong 0$ | $\gamma \cong 1$ | $\gamma \cong \infty$

extremely risk-averse (minimize worst-case cost) "Murphy's law"

risk-neutral (minimize average cost)

extremely risk-seeking (minimize best-case cost)

Planning with Nonlinear Utility Functions
Plan-Evaluation Example

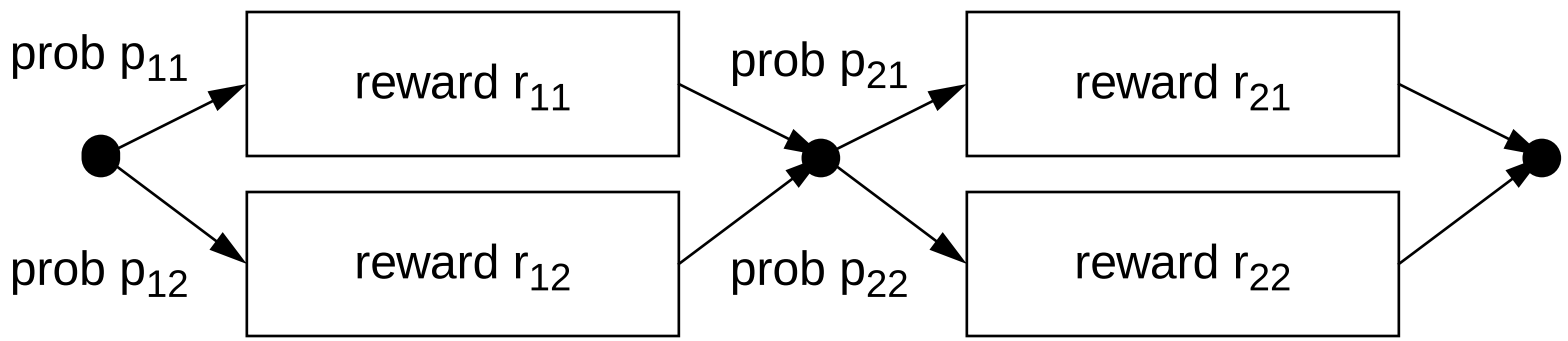

average total reward: $$\sum_i \sum_j (p_{1i}\, p_{2j}\, (r_{1i} + r_{2j})) \quad = \quad \sum_i (p_{1i}\, r_{1i}) + \sum_j (p_{2j}\, r_{2j})$$

average total utility: $$\sum_i \sum_j (p_{1i}\, p_{2j}\, u(r_{1i} + r_{2j})) \quad = \quad \sum_i (p_{1i}\, \gamma^{r_{1i}}) \times \sum_j (p_{2j}\, \gamma^{r_{2j}})$$

[Wellman, Ford, and Larson]
[Haddawy and Hanks]

Planning with Nonlinear Utility Functions
Planning

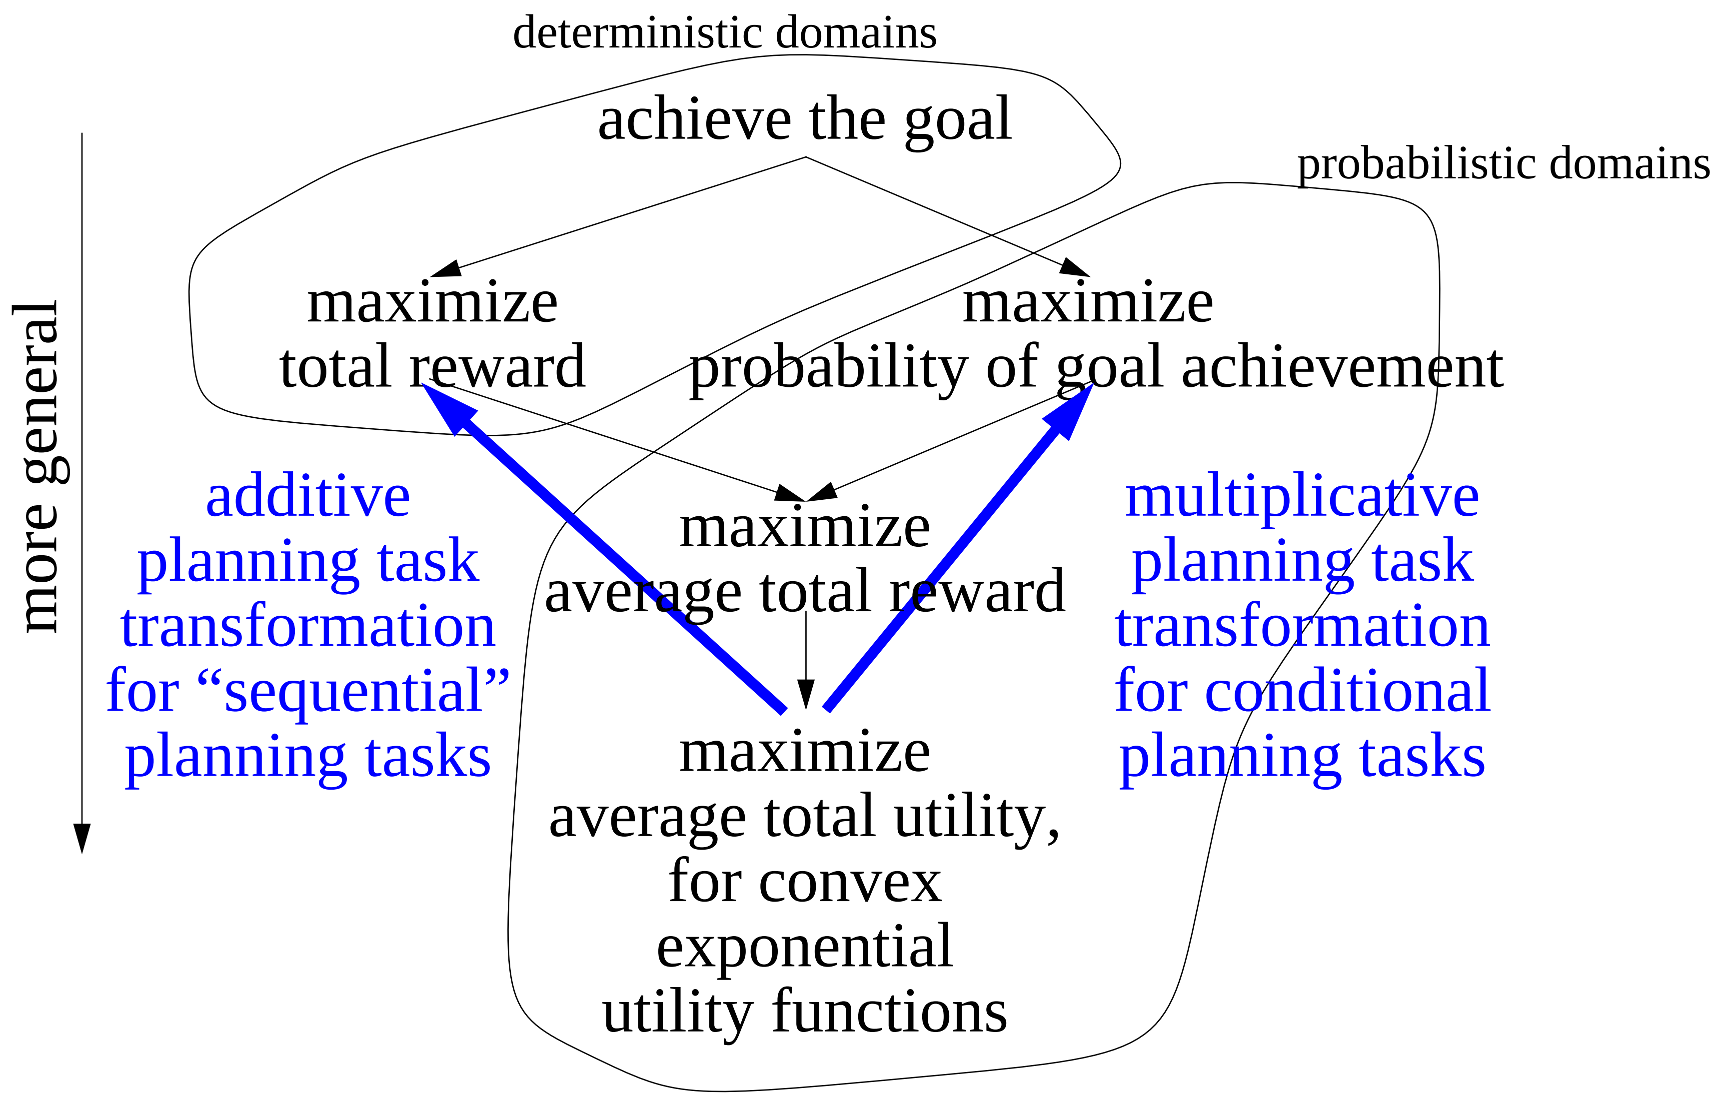

Planning with Nonlinear Utility Functions
Advantages of Planning Task Transformations

- simple representation changes
- can be performed on a variety of planning task representations
- can easily be integrated into agent architectures
- extend the functionality of existing planners
- make planning with exponential utility functions as fast as these planners

Planning with Nonlinear Utility Functions
Information Gathering - A Realistic Planning Task

when (and what) to sense

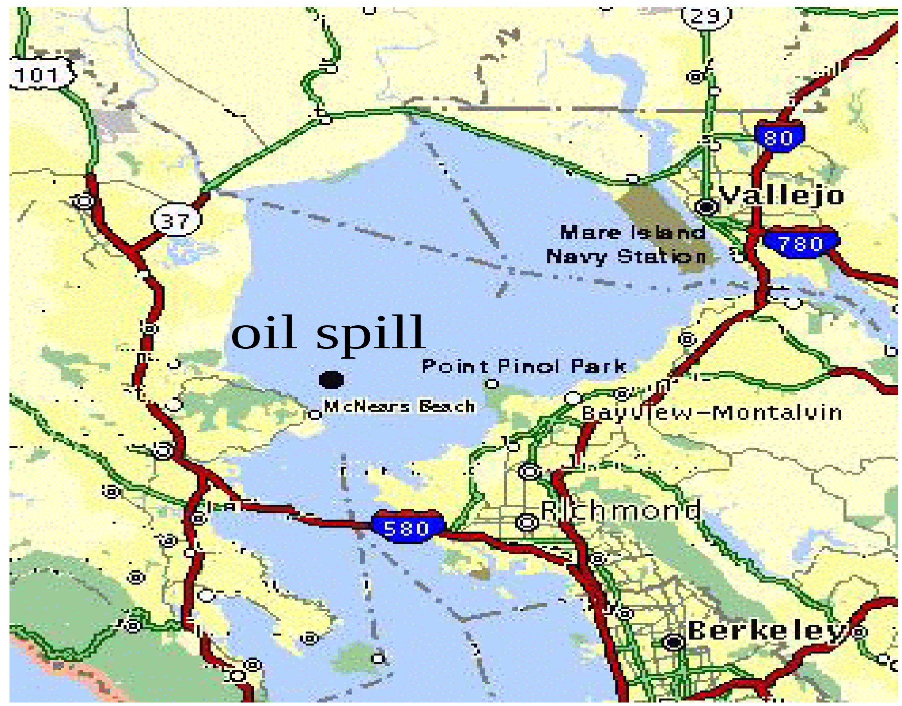

Can we characterize
how the behavior of decision makers changes with their risk attitude?

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

when (and what) to sense

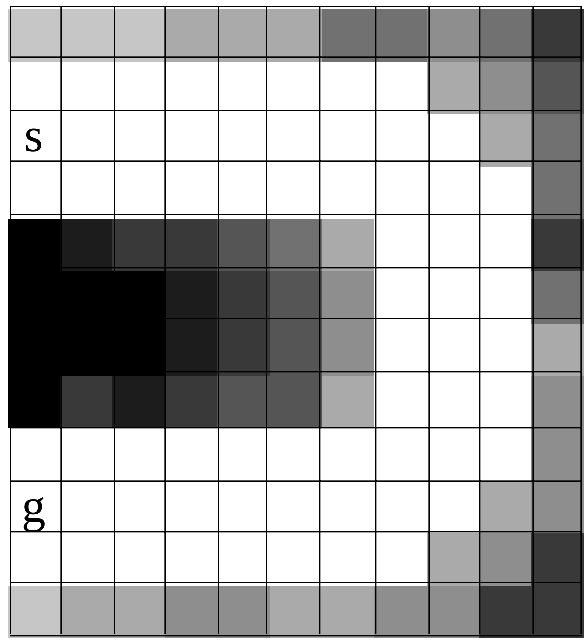

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

move north	(N)
move east	(E)
move south	(S)
move west	(W)

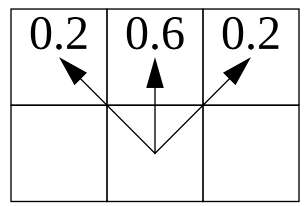

does not provide the robot with any information about its current location

cost: 1 (road) ... 10 (mud)

---

sense (O)

does not move the robot

provides the robot with certainty about its current location

cost: 0.2

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

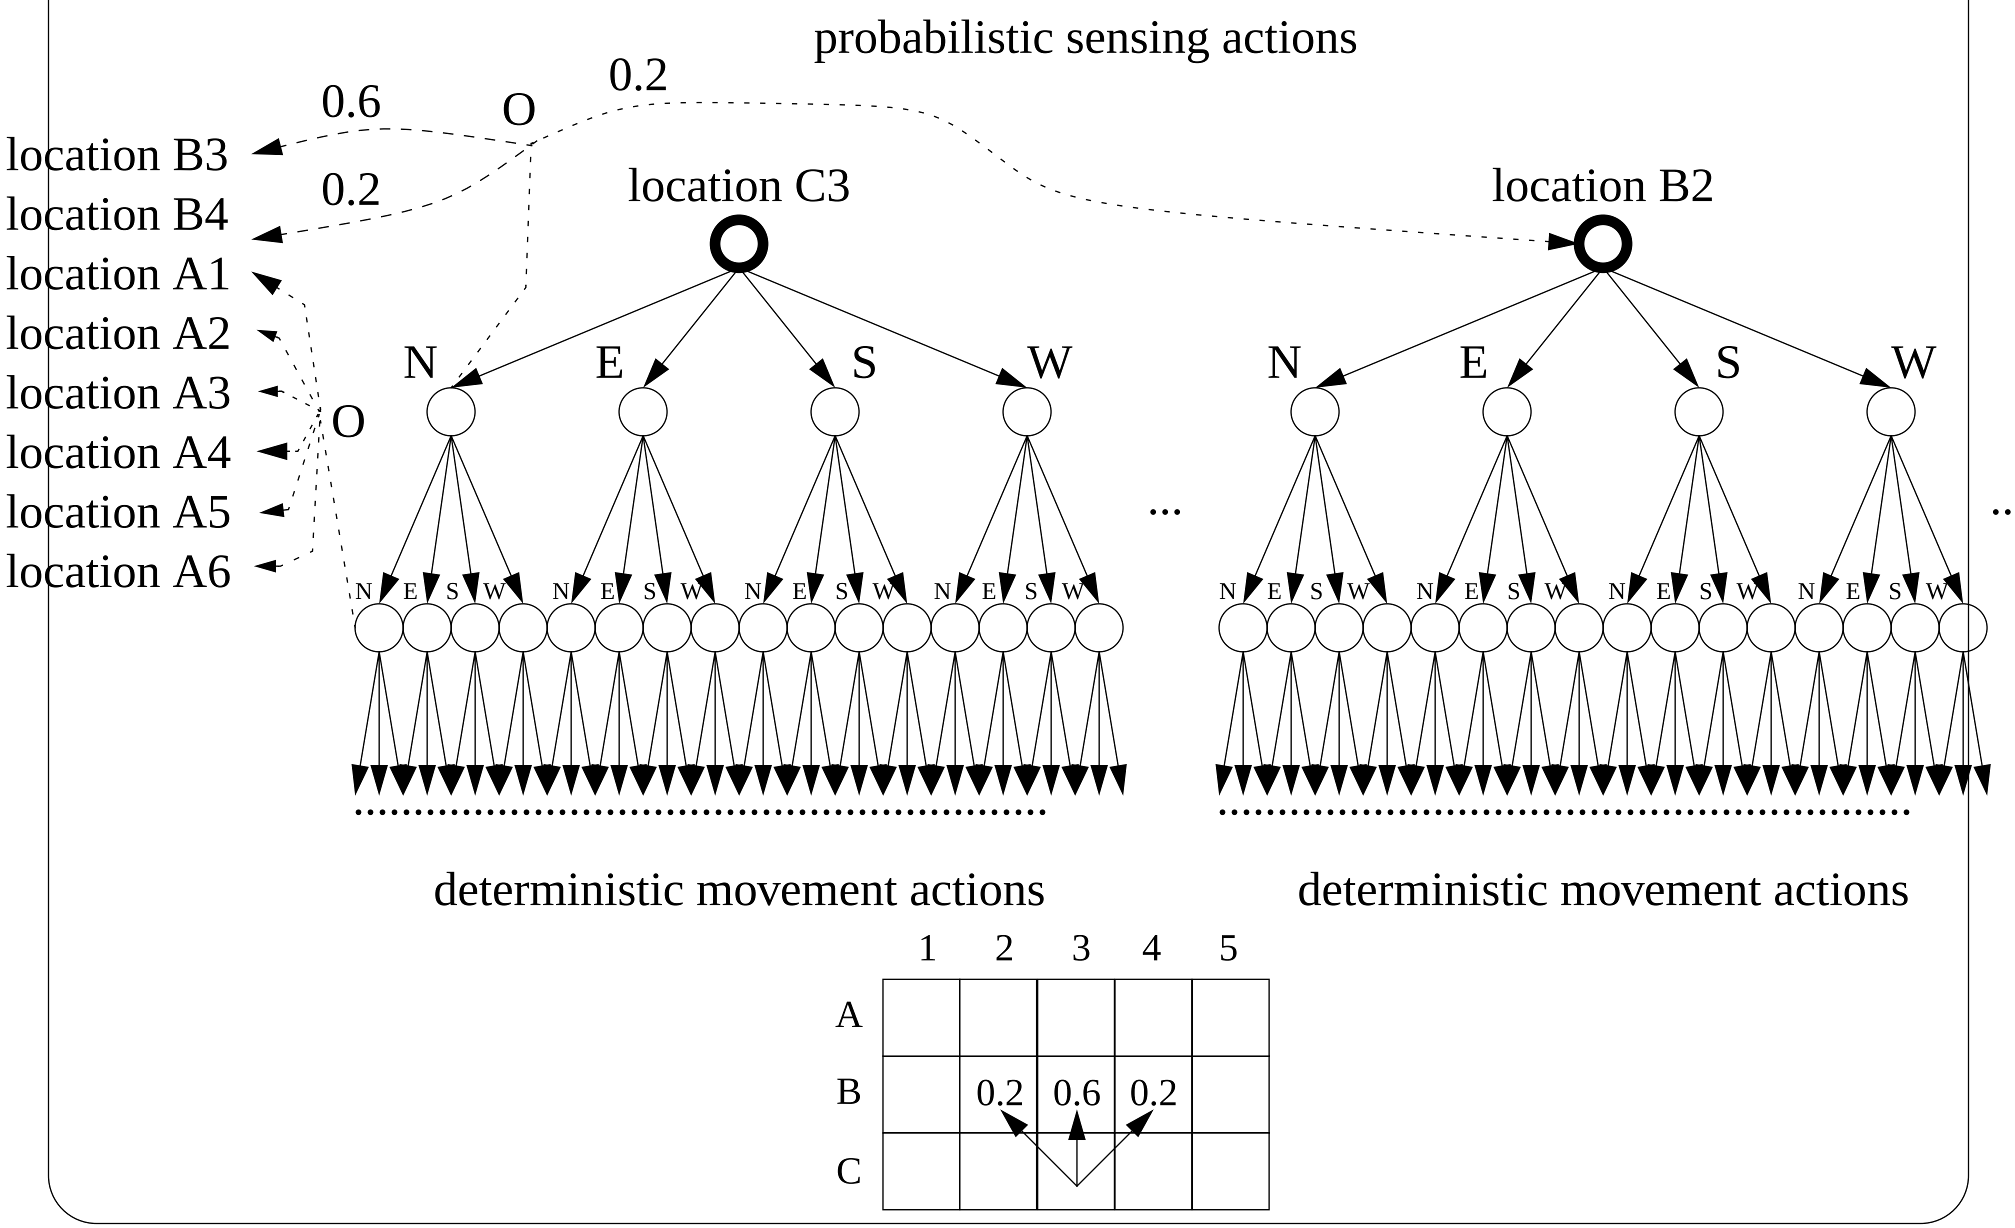

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

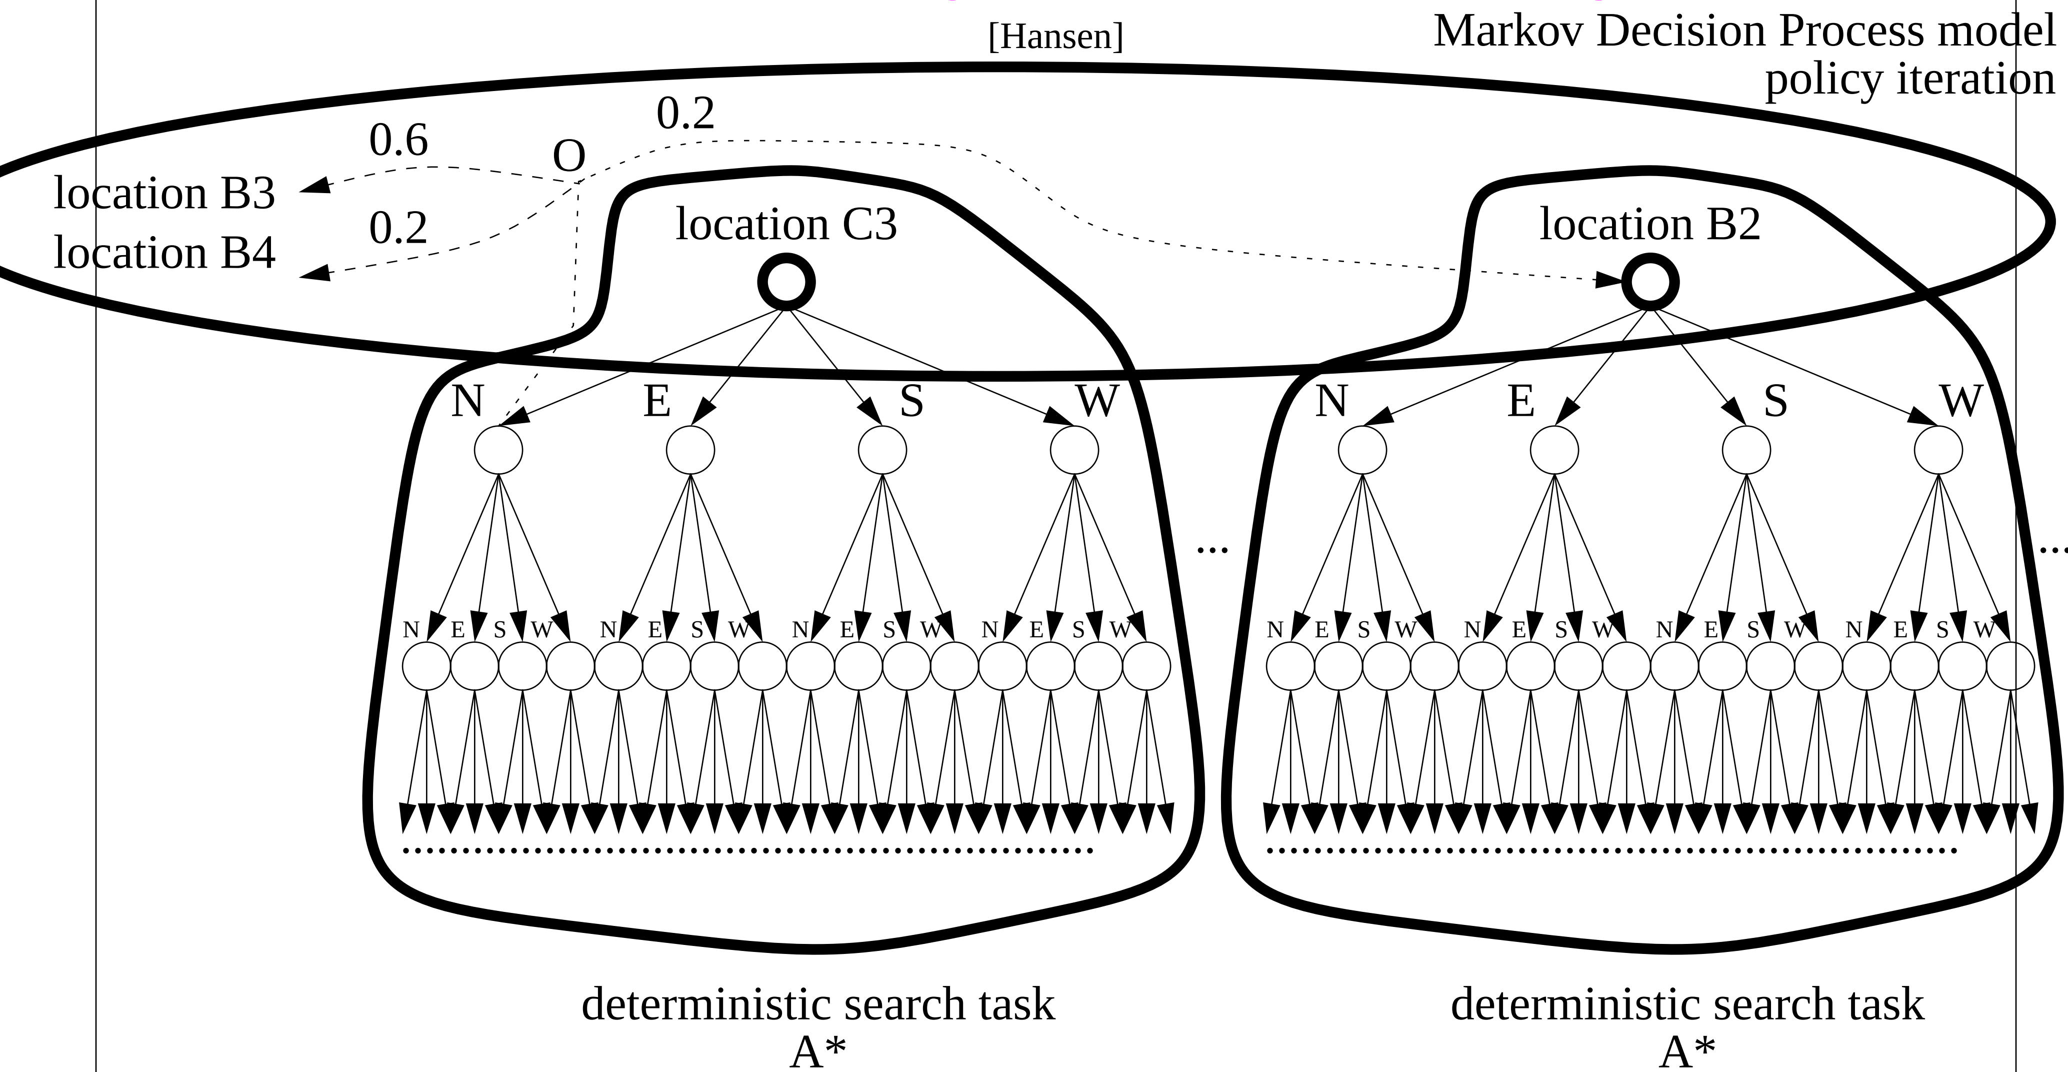

combine this sensor-planning method with the multiplicative planning-task transformation

see the paper for a complete algorithm and proofs

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

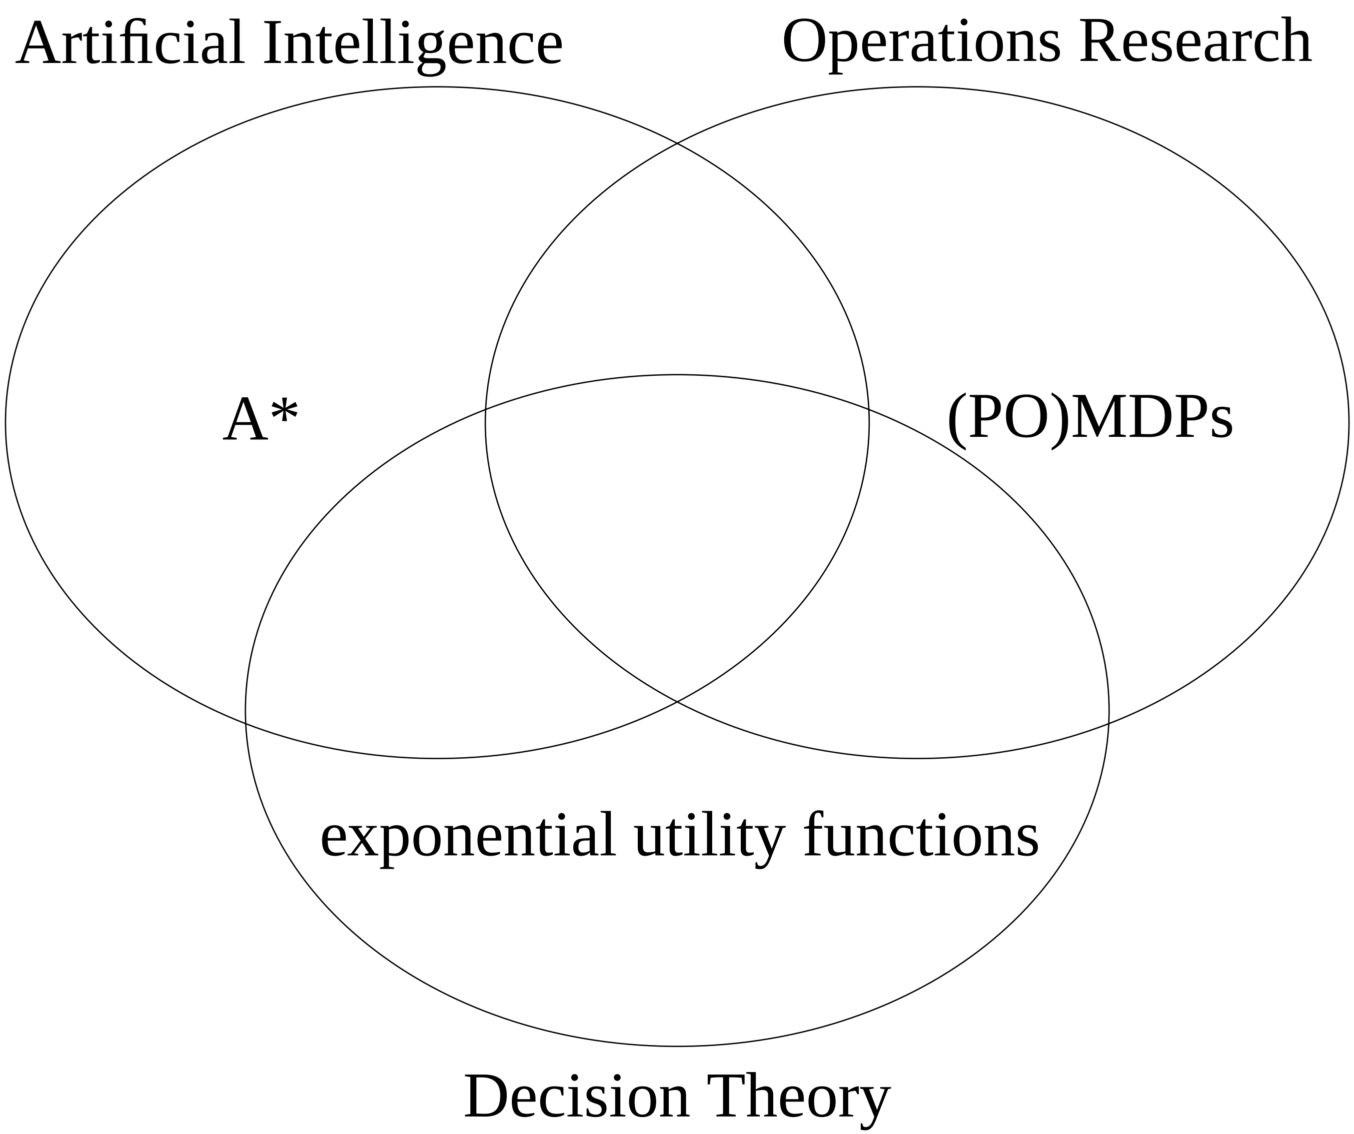

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

policy with fewer sensing actions

policy with more sensing actions

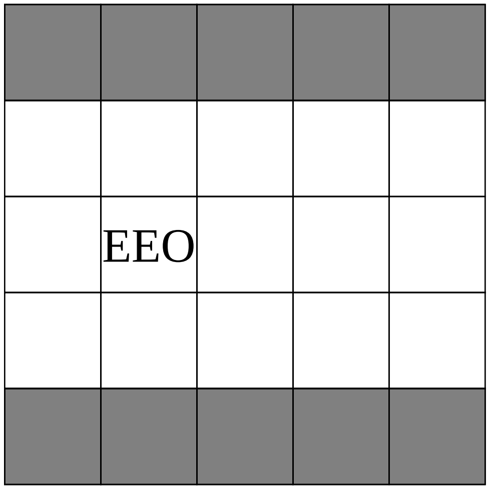

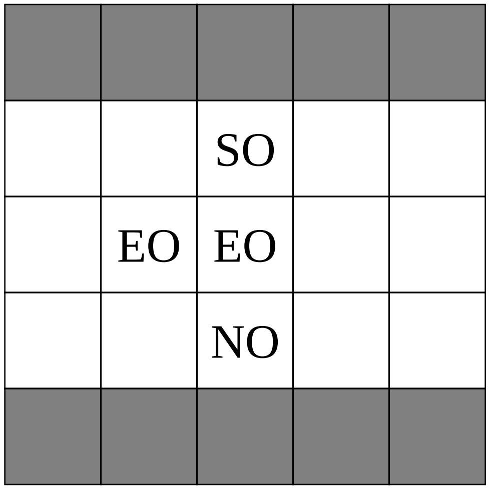

policy

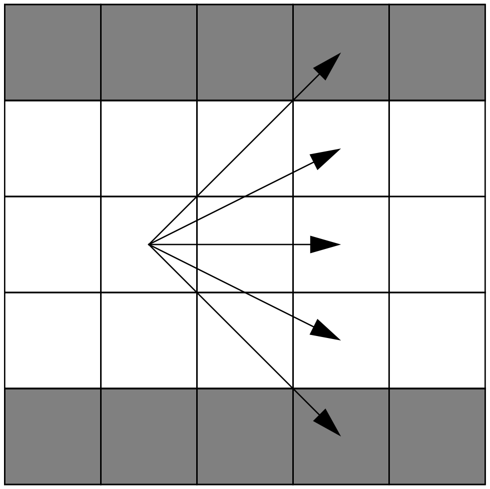

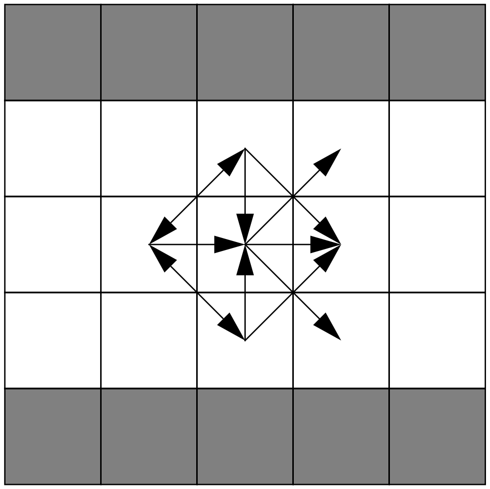

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

location occupancy

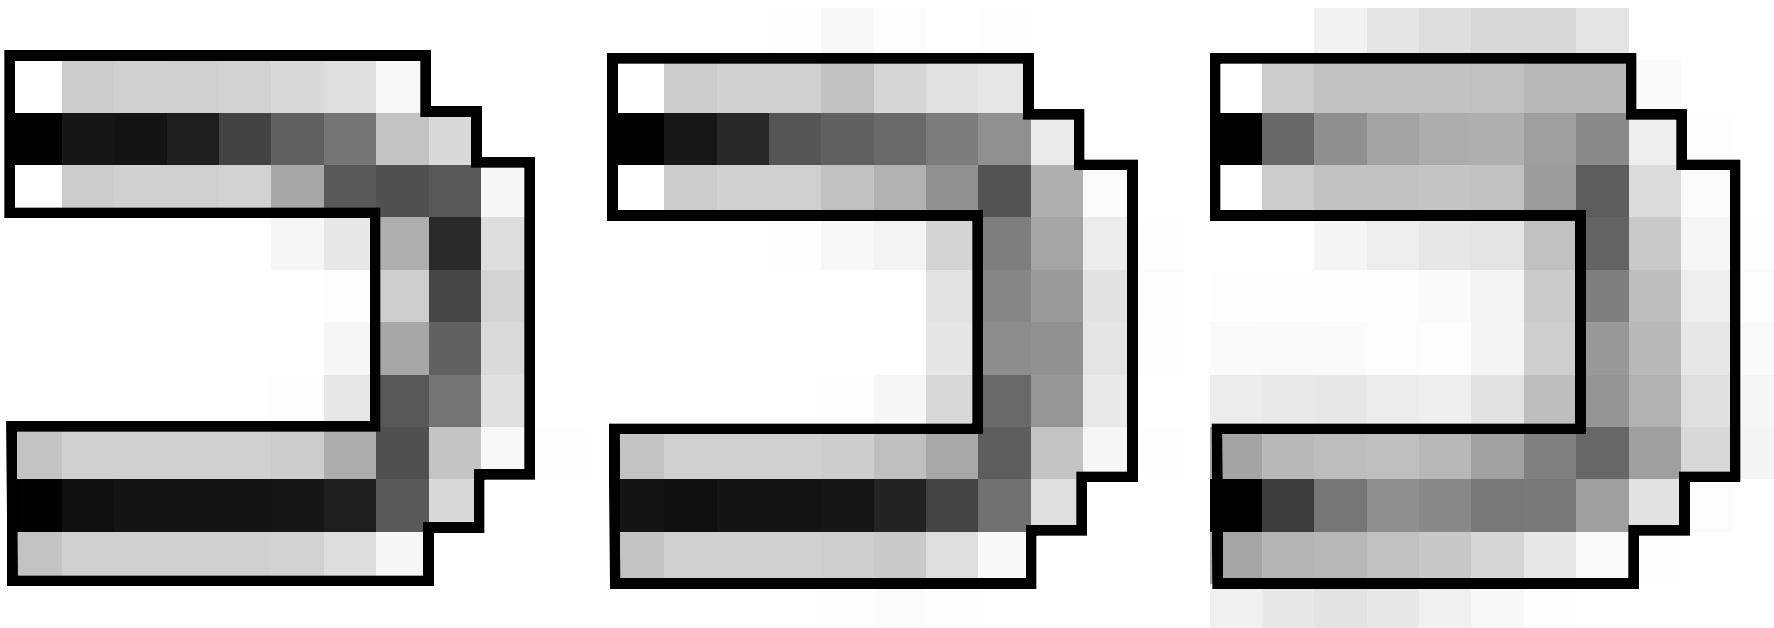

$\gamma = 0.86$ $\gamma = 1.00$ $\gamma = 1.40$

more risk-averse

minimize average cost

standard for
decision-theoretic planners

more risk-seeking

the more risk-seeking a decision maker is
the more important is the best-case reward (best-case cost)

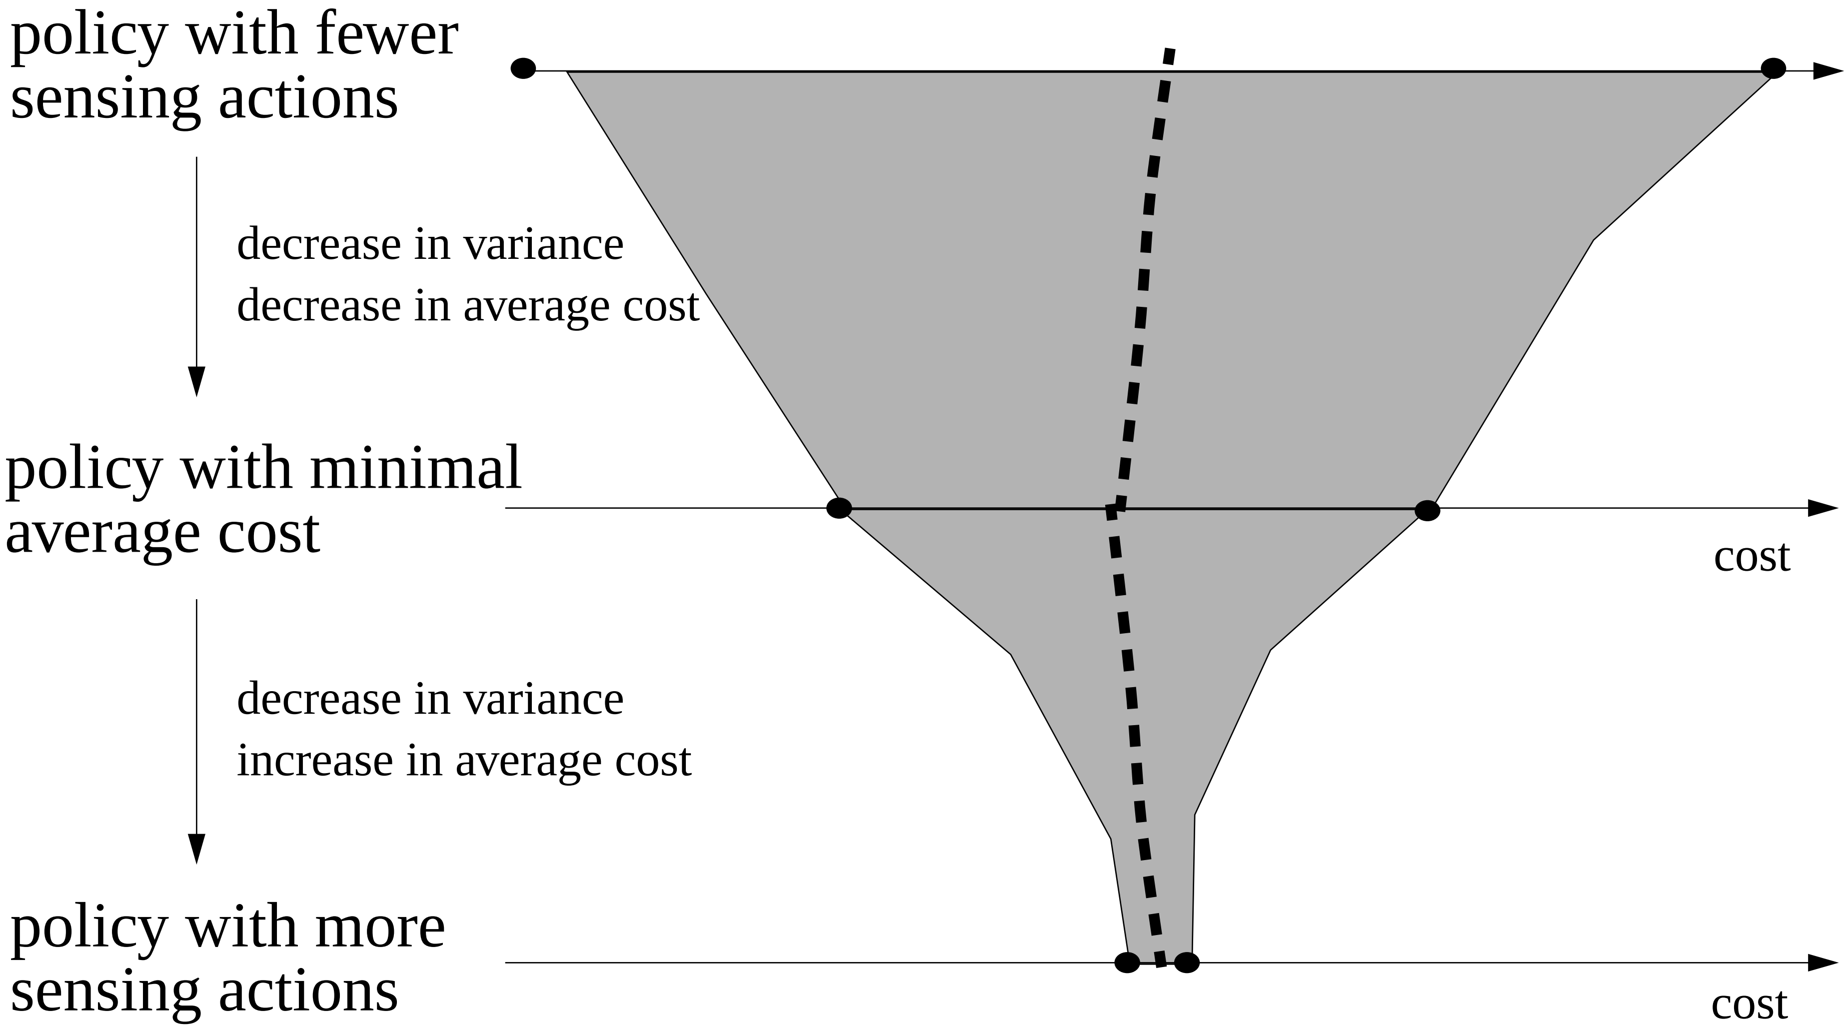

the more risk-averse a decision maker is
the more important is the worst-case reward (worst-case cost)

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

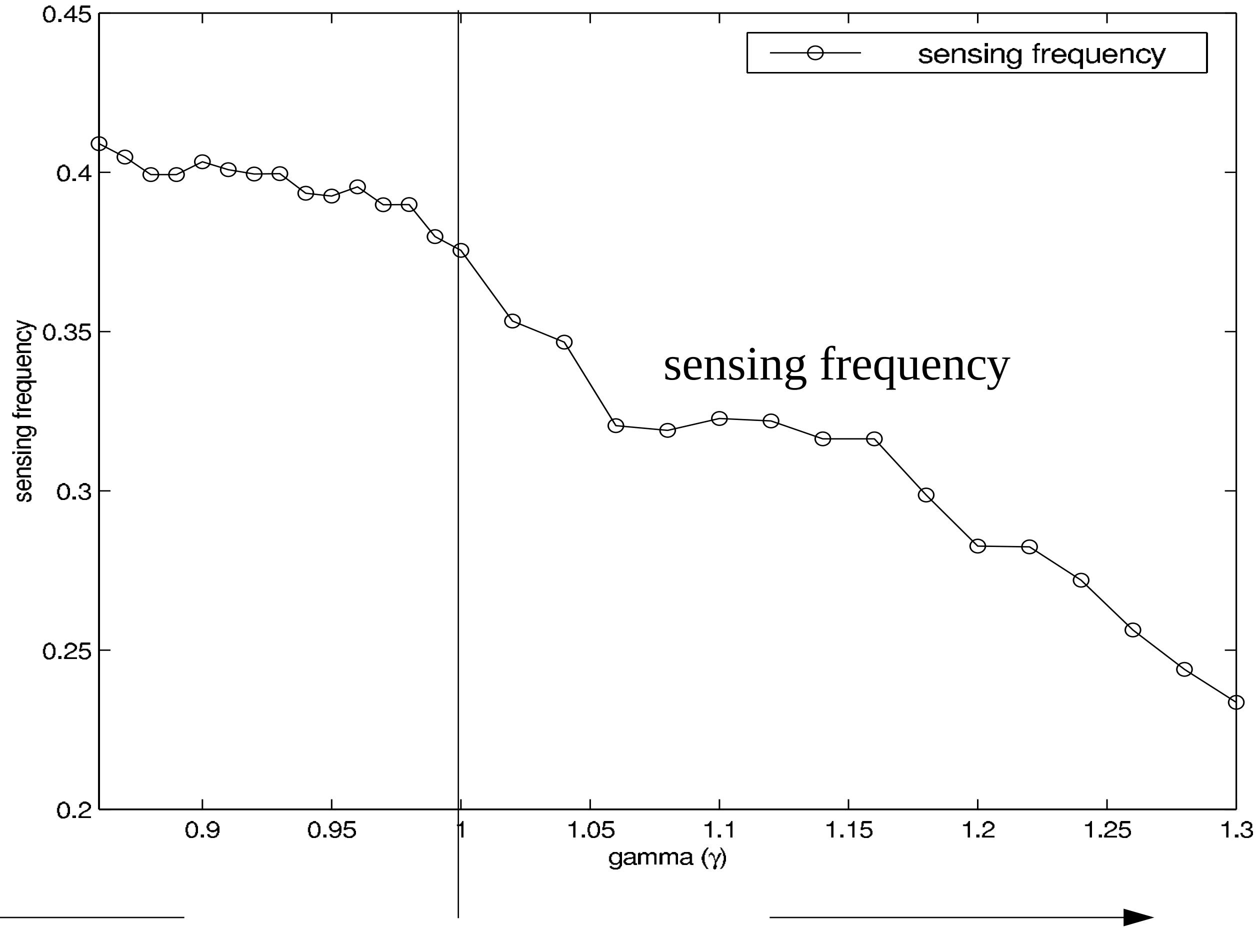

more risk-averse

minimize average cost

more risk-seeking

standard for
decision-theoretic planners

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

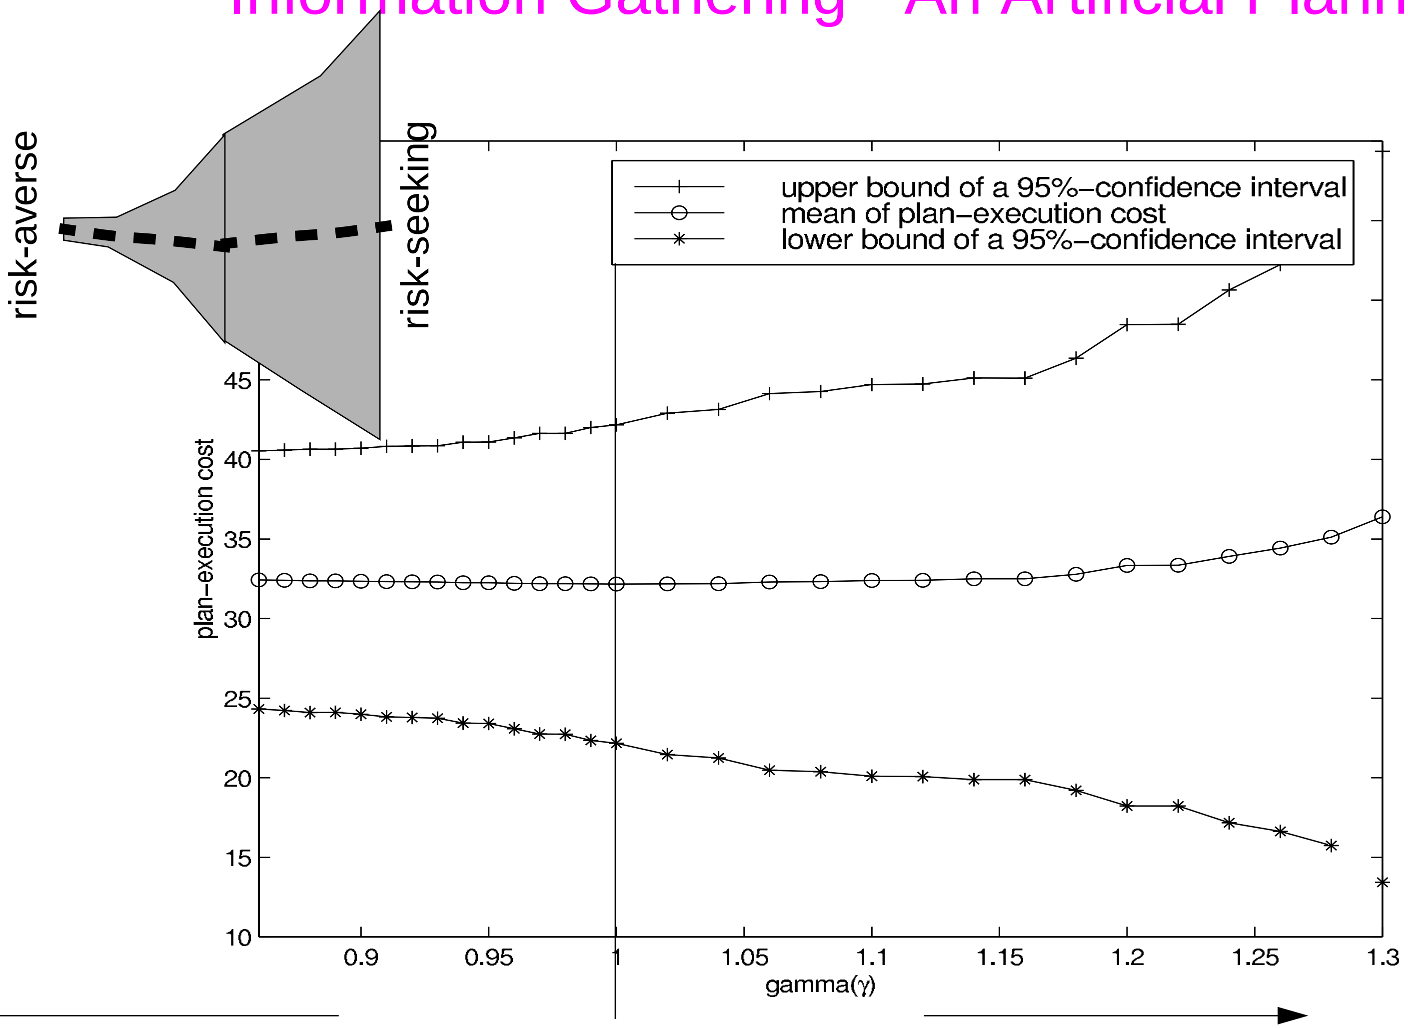

average cost +
2x standard deviation
("worst-case cost")

average cost

average cost -
2x standard deviation
("best-case cost")

more risk-averse | minimize average cost | more risk-seeking

standard for
decision-theoretic planners

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

$\gamma = 0.86$ - risk averse

$\gamma = 0.86$ (pessimistic robots)

	1	2	3	4	5	6	7	8	9	10	11
A								SSSO			
B		SEO	SEO	SEO	SEO	SEO	SO	SSO			
C	EO	EO	EO	EO	EO	EO	SEO	SO	SO	WSO	
D		NEO	NEO	NEO	EEEO	EEO	EO	ESO	SO	WSO	WWSO
E							EESO	ESO	SO	WSO	
F								SSSO	SO	WSO	
G								SSO	SO	WSO	
H						SSWO	SSWO	SO	WSO	WO	
I	SO	SWO	SWO	SWO	SWO	SWO	SWO	SWO	WO	WWO	WWWO
J	goal	WO	WO	WO	WO	WO	WO	WO			
K	NO	NWO	NWO	NWO	NWO	NWO	NWO	NWO	NWWO		
L							NNWO	NNWO			

$\gamma = 1.40$ - risk seeking

$\gamma = 1.40$ (optimistic robots)

	1	2	3	4	5	6	7	8	9	10	11
A	SSEEEO				SSEO	SSEO	SSSO	SSSO			
B	SEEEEO			SEEO	SEEO	SEO	SSO	SSO	SSSSSSSWO		
C	EEEEEEESO			EEEESO	EEESO	EESO	SO	SO	SSSSSSWO		
D	EEEEEEESO			EEEESO	EEESO	EESO	ESO	SSSSSWO	SSSSSWO	WSSO	WWSSO
E	NNEEEEO	NEEEEEO	NEEEEO	NEEEO	NEEO	EESSSO	ESSSO	SSSSWO	SSSSWO	WSSO	WWSSO
F	SSSSO	SSSSO	SSSSO	SSSSWO	SSSSWWO	EESSO	ESSSO	SSSWO	SSSWO	WSSO	WWSO
G	SSSO	SSSO	SSSO	SSSWO	SSSWWO	SSSWWWO	SSSWWWWO	SSWO	SSWO	WSO	WWSO
H	SSO	SSO	SSWO	SSWWO	SSWWO	SSWWWO	SSWWWWO	SSWWWWWO	SWO	WWO	WWWO
I	SO	SO	SWO	SWWO	SWWWO	SWWWWO	SWWWWWO	SWWWWWWO	WO	WWO	WWWO
J	goal	WO	WWO	WWWO	WWWWO	WWWWWO	WWWWWWO	WWWWWWWO	WWWWWWWWO	WWWWWWWWWO	WWWO
K	NO	NO	NWO	WWWO	WWWWO	WWWWWO	WWWWWWO	WWWWWWWO	WWWWWWWWO	WWWWWWWWWO	WWWWWWWWWWO
L	NNO	NNO	NNWO	NWWO	NWWWO	NWWWWO		NWWWWWWO			

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

computation time

	more risk-averse ←	minimize average cost	more risk-seeking →
	g = 0.86	g = 1.00	g = 1.40
A* node expansions	2,071	2,815	5,808
millisecs / node exp	4.6	2.0	5.2
		(Hansen's method) standard for decision-theoretic planners	

Planning with Nonlinear Utility Functions
Information Gathering - An Artificial Planning Task

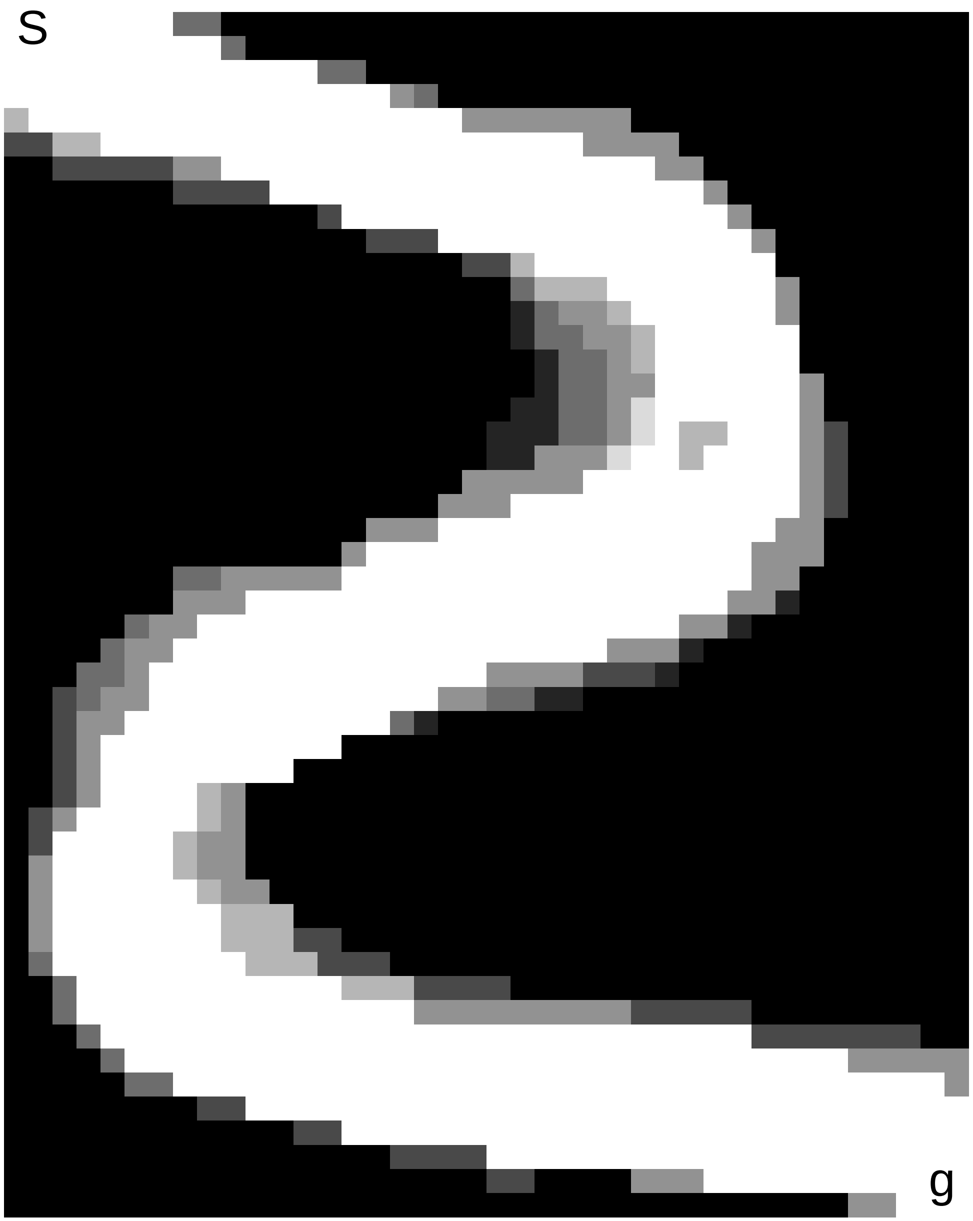

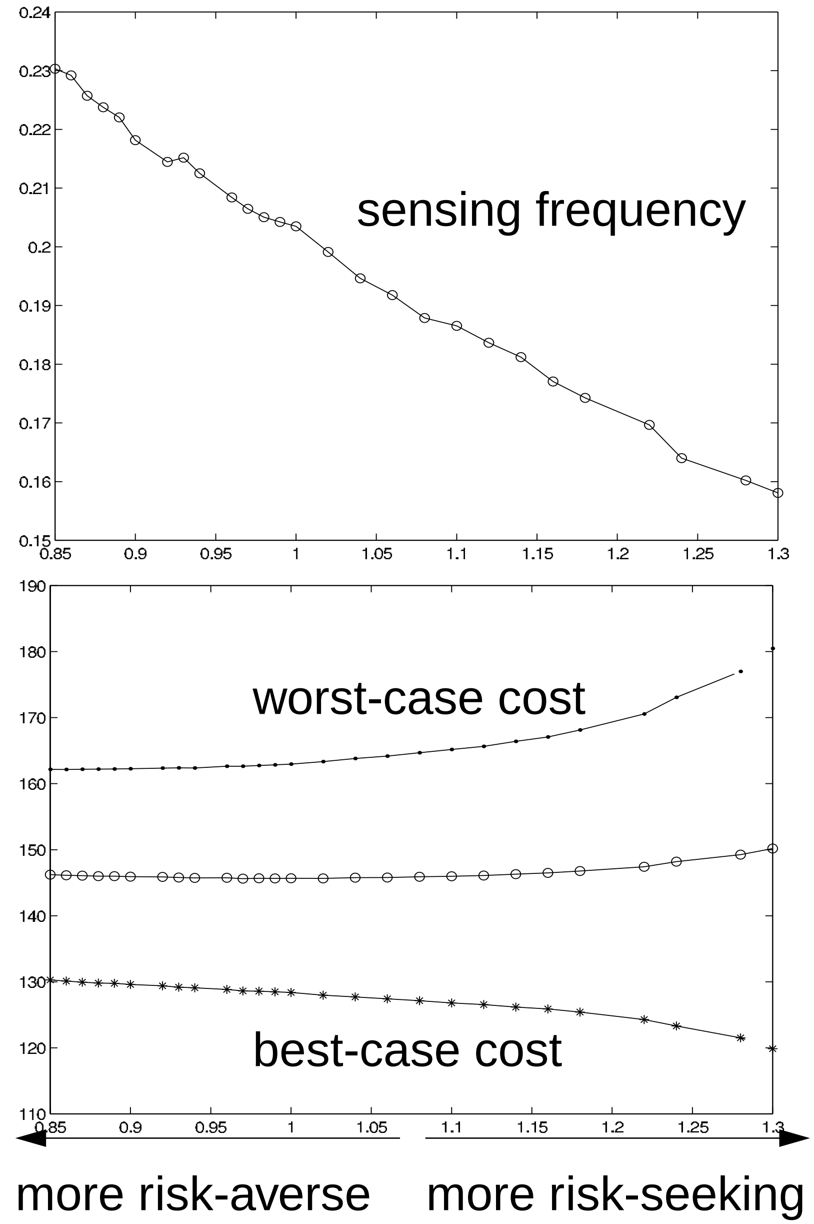

Planning with Nonlinear Utility Functions
Auction Planning - An Application

50%	10,000,000 dollars	50%	4,000,000 dollars
50%	0 dollars	50%	5,000,000 dollars

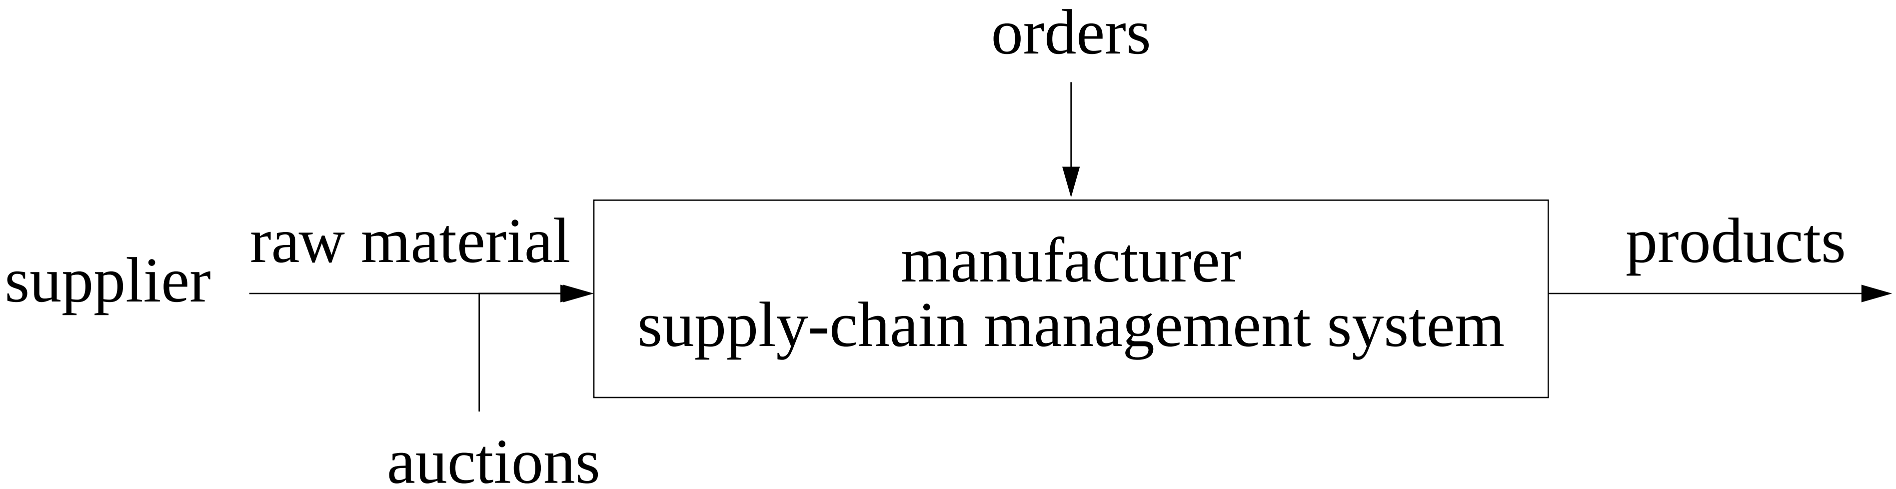

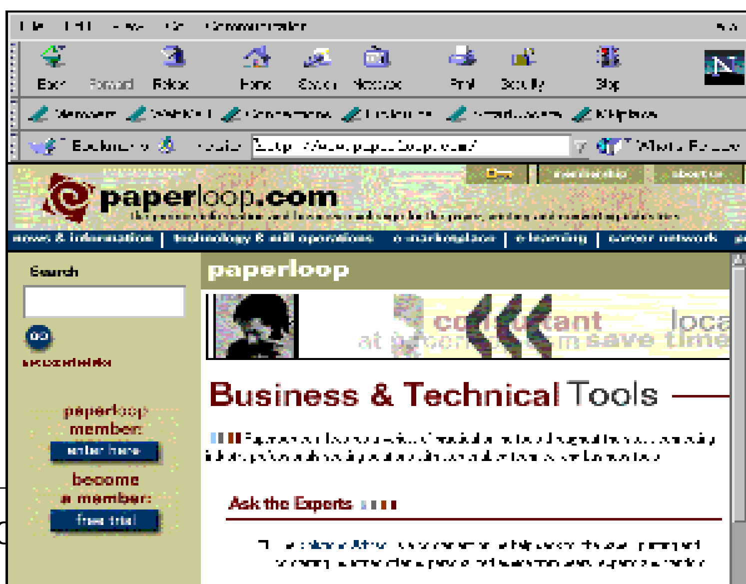

Planning with Nonlinear Utility Functions
Auction Planning - An Application

Symmetric Independent Private Values Model

- only one item is for sale
- the item will be sold to the highest bidder for any positive price
- the number of bidders N is known to all bidders
- each bidder knows their own valuation v_i for the auctioned item (= the difference in profit between owning and not owning it)
- no bidder knows the other bidder's valuations for the auctioned item but these valuations are independent random variables drawn from a given continuous distribution $F(v)$ with density $f(v)$ over the nonnegative real-values bids, and this distribution is known
- the bidders are indistinguishable

Sealed-Bid Model

- the bidders submit secret bids

First-Price Model

- the winner of the auction pays what they bid

Planning with Nonlinear Utility Functions
Auction Planning - An Application

The optimal bidding function for risk-neutral decision makers that participate in first-price sealed bid auctions in the symmetric independent private values (SIPV) model is

$$b(v) = v - \frac{\int_0^v F(t)^{N-1} dt}{F(v)^{N-1}}$$ [McAfee and McMillan]

Planning with Nonlinear Utility Functions
Auction Planning - An Application

Theorem:

The optimal bidding function for risk-averse decision makers with concave exponential utility functions that participate in first-price sealed bid auctions in the symmetric independent private values (SIPV) model is

$$b(v) = (N-1) \cdot \log_{\gamma} F(v) - \log_{\gamma} \int_0^v \gamma^{-t} \, dF(t)^{N-1}$$

- v: valuation of item
- N: number of bidders
- F: distribution over valuations of item for other bidders
- γ: utility functions $u(v) = -\gamma^v$

Planning with Nonlinear Utility Functions
Auction Planning - An Application