Decision Making
in Robots and Autonomous Agents

Learning about Preferences
(based on material by C. Boutilier, D. Braziunas)

Subramanian Ramamoorthy
School of Informatics

24 March, 2015

Preference Elicitation in AI

Shopping for a Car:

Luggage Capacity?
Two Door? Cost?
Engine Size?
Color? Options?

The Preference Bottleneck

- *Preference elicitation:*
  the process of determining a user's preferences/ utilities to the extent necessary to make a decision on her behalf

- Why a bottleneck?
  – preferences vary widely
  – large (multiattribute) outcome spaces
  – quantitative utilities (the "numbers") difficult to assess

Automated Preference Elicitation

- Questions:
  - decomposition of preferences
  - what preference information is *relevant* to the task at hand?
  - when is the elicitation effort *worth the improvement* it offers in terms of decision quality?
  - what *decision criterion* to use given partial utility information?

Constraint-based Decision Problems

- Constraint-based optimization (CBO):
  - outcomes over variables $\boldsymbol{X} = \{X_1 \ldots X_n\}$
  - constraints $\boldsymbol{C}$ over $\boldsymbol{X}$ spell out feasible decisions
  - generally compact structure, e.g., $X_1 \,\&\, X_2 \supset \neg X_3$
  - add a *utility function* $u: Dom(\boldsymbol{X}) \rightarrow \mathcal{R}$
  - preferences over configurations

Constraint-based Decision Problems

- Must express *u* compactly like ***C***
  - *generalized additive independence (GAI)*
    - model proposed by Fishburn (1967)
    - nice generalization of additive linear models
  - expressible by graphical model capturing independence

Additive Linear Models of Utility

$$u(\mathbf{x}) = \sum_{i=1}^{n} u_i(x_i) = \sum_{i=1}^{n} \lambda_i v_i(x_i).$$

$\lambda_1 = 0.2$

Color	u_1
red	1.0
blue	0.7
grey	0.0

$\lambda_2 = 0.3$

Drs	u_2
2	1.0
4	0.8
hatch	0.2
wag'n	0.0

$\lambda_3 = 0.5$

Pwr	u_3
350	1.0
280	0.7
150	0.0

u(red,2dr,280hp) = 0.85

Additive Utility

- An additive representation of *u* exists iff decision maker is indifferent between any two lotteries where the marginals over each attribute are identical
  - $l_1(\mathbf{X}) \sim l_2(\mathbf{X})$ whenever $l_1(X_i) = l_2(X_i)$ for all X_i

$$u(\mathbf{x}) = \sum_{i=1}^{n} u_i(x_i) = \sum_{i=1}^{n} \lambda_i v_i(x_i).$$

Factored Utilities: **GAI** Models

- Set of K factors f_k over subset of variables ***X****[k]*
  - "local" utility for each local configuration

$$u(\mathbf{x}) = \sum_{k \le K} f_k(\mathbf{x}[k])$$

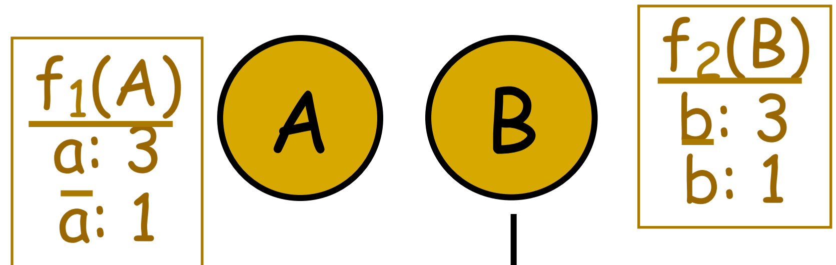

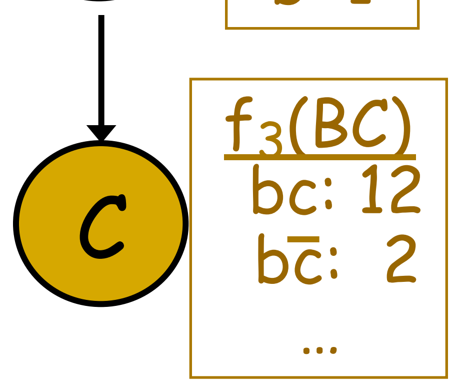

$u(abc) = f_1(a) + f_2(b) + f_3(bc)$

- [Fishburn67] u in this form exists iff
  - lotteries p and q are equally preferred whenever p and q have the same marginals over each ***X****[k]*

Optimization with GAI Models

$$u(\mathbf{x}) = \sum_{k \le K} f^k(\mathbf{x}[k])$$

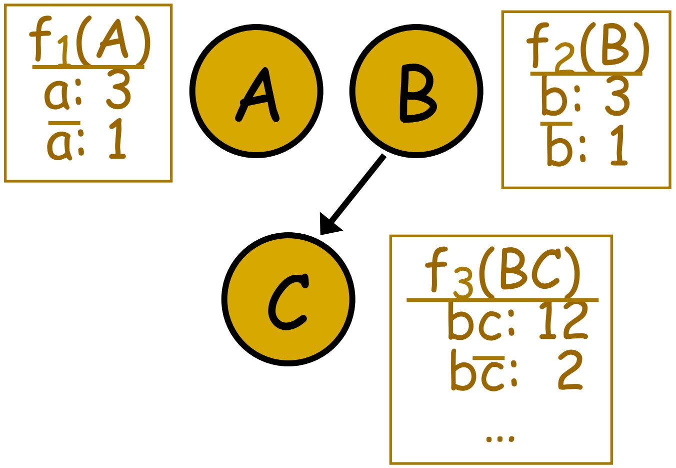

- Optimize using Integer Programming (or, e.g., Variable Elimination)
  - number of variables *linear* in size of GAI model

$$\max_{\{I_{\mathbf{x}[k]}, X_i\}} \sum_k \sum_{\mathbf{x}[k] \in Dom(\mathbf{X}[k])} u_{\mathbf{x}[k]} I_{\mathbf{x}[k]} \quad \text{subj. to } A, C$$

Difficulties in CBO

- **Utility elicitation:** how do we assess individual user preferences?
  - need to elicit GAI model structure (independence)
  - need to elicit (constraints on) GAI parameters
  - need to make decisions with imprecise parameters

Strict Utility Function Uncertainty

- User's actual utility *u* unknown
- Assume *feasible set* $F \subseteq U = [0,1]^n$
  - allows for unquantified or "strict" uncertainty
  - e.g., **F** a set of linear constraints on GAI terms

$$u(red,2door,280hp) > 0.4$$
$$u(red,2door,280hp) > u(blue,2door,280hp)$$

- How should one make a decision? elicit info?

Strict Uncertainty Representation

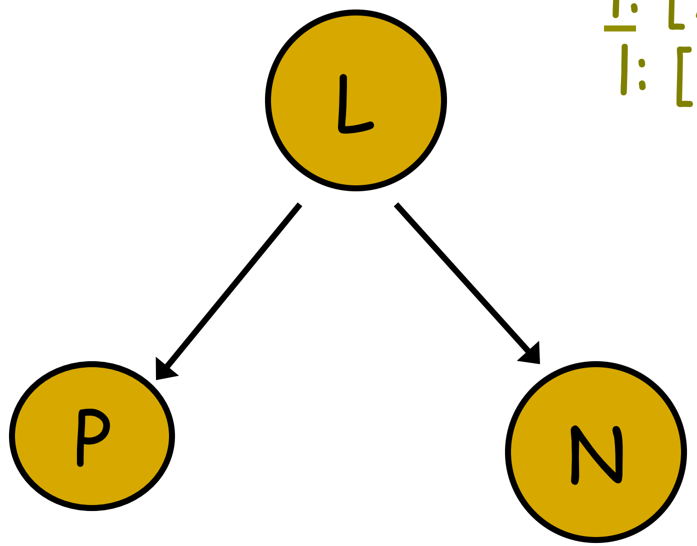

$f_1(L)$
$\underline{l}$: [7,11]
l: [2,5]

$f_2(L,N)$
l,n: [2,4]
l,$\bar{n}$: [1,2]

Utility Function

Bayesian Utility Function Uncertainty

- User's actual utility *u* unknown
- Assume *density P* over $U = [0,1]^n$
- Given *belief state* *P*, EU of decision *x* is:

$$EU(x,P) = \int_{u} \vec{p}_x \, \vec{u} \, P(\vec{u})$$

- Decision making is easy, but elicitation harder?
  – query must assess *expected value of information*

Bayesian Representation

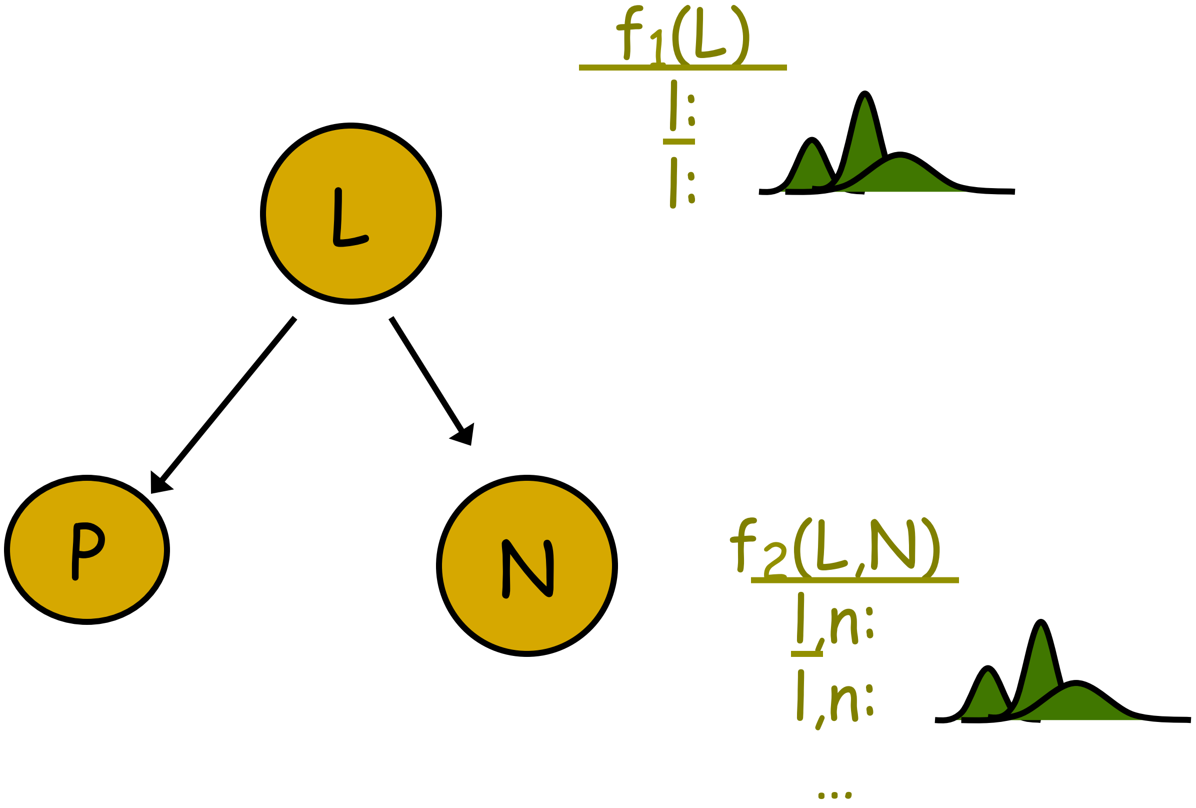

Utility Function

Query Types

- *Comparison queries* (is $\boldsymbol{x}$ preferred to $\boldsymbol{x'}$?)
  - impose linear constraints on parameters
    - $\Sigma_k f_k(\boldsymbol{x}[k]) > \Sigma_k f_k(\boldsymbol{x'}[k])$
  - Interpretation is straightforward

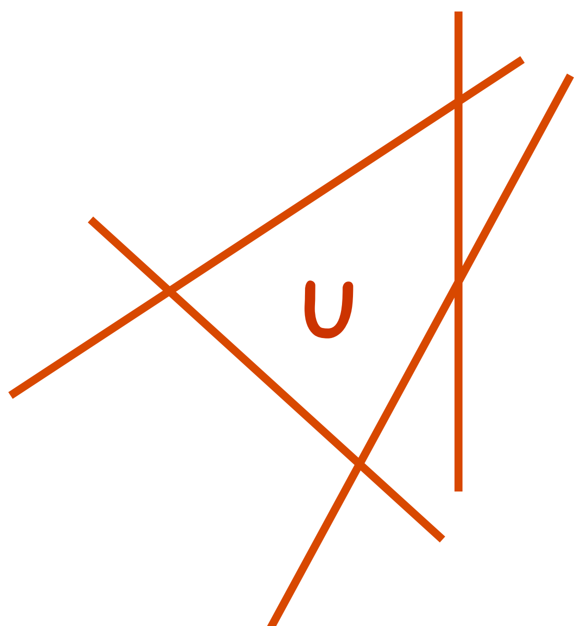

Query Types

- *Bound queries* (is $f_k(\boldsymbol{x}[k]) > v$?)
  - response tightens bound on specific utility parameter
  - can be phrased as a *local standard gamble query*

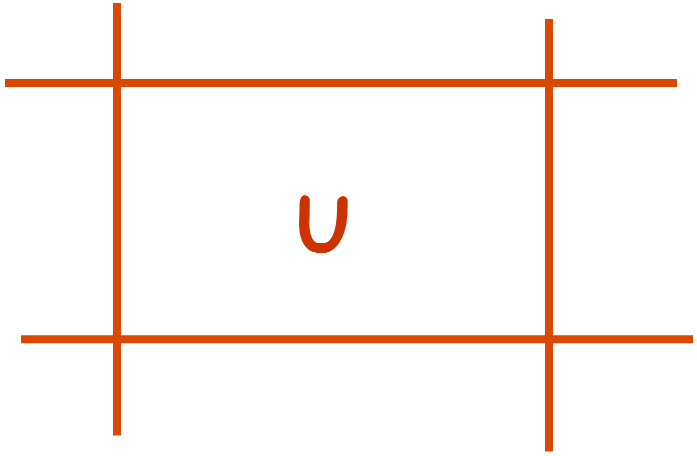

Difficulties with Bound Queries

- Bound queries focus on *local* factors
  - but values cannot be fixed without reference to others!
  - seemingly "different" local prefs correspond to same *u*

$$u(Color,Doors,Power) = u_1(Color,Doors) + u_2(Doors,Power)$$

$$\underset{10}{u(red,2door,280hp)} = \underset{\not{6}\ \ 1}{u_1(red,2door)} + \underset{\not{4}\ \ 9}{u_2(2door,280hp)}$$

$$\underset{6}{u(red,4door,280hp)} = \underset{3}{u_1(red,4door)} + \underset{3}{u_2(4door,280hp)}$$

Local Queries [BB05]

- We wish to avoid queries on whole outcomes
  - can't ask purely local outcomes
  - but can condition on a *subset* of default values

- *Conditioning set* $C(f)$ for factor $f_i(\boldsymbol{X}_i)$:
  - variables that share factors with $\boldsymbol{X}_i$
  - setting default outcomes on $C(f)$ renders $\boldsymbol{X}_i$ independent of remaining variables
  - enables local calibration of factor values

Local Standard Gamble Queries

- Local std. gamble queries
  - use "best" and "worst" (anchor) local outcomes
    -- conditioned on default values of conditioning set
  - bound queries on other parameters relative to these
  - gives *local value function v(x[i])* (e.g., *v(ABC)*)
- Hence we can legitimately ask *local* queries:

$$(\mathbf{x}_i, \mathbf{x}^0_{C_i}) \sim \langle p, (\mathbf{x}_i^\top, \mathbf{x}^0_{C_i}); 1-p, (\mathbf{x}_i^\perp, \mathbf{x}^0_{C_i})\rangle$$

- But local Value Functions not enough:
  - must calibrate: requires global scaling

Global Scaling

- Elicit utilities of anchor outcomes *with respect to* global best and worst outcomes
  - the 2*m "best" and "worst" outcomes for each factor
  - these require *global* standard gamble queries *(note: same is true for pure additive models)*

Bound Query Strategies

- Identify conditioning sets C_i for each factor f_i
- Decide on "default" outcome
- For each f_i identify top and bottom *anchors*
  - e.g., the most and least preferred values of factor i
  - given default values of C_i
- Queries available:
  - local std gambles: use anchors for each factor, C-sets
  - global std gambles: gives bounds on anchor utilities

Partial preference information
Bayesian uncertainty

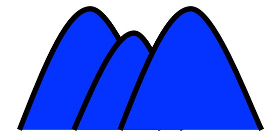

- Probability distribution p over utility functions
- Maximize expected (expected) utility

MEU decision $x^* = \arg\max_x E_p [u(x)]$

- Consider:
  - elicitation costs
  - values of possible decisions
  - optimal tradeoffs between elicitation effort and improvement in decision quality

Query selection

- At each step of elicitation process, we can
  - obtain more preference information
  - make or recommend a terminal decision

Bayesian approach
Myopic EVOI

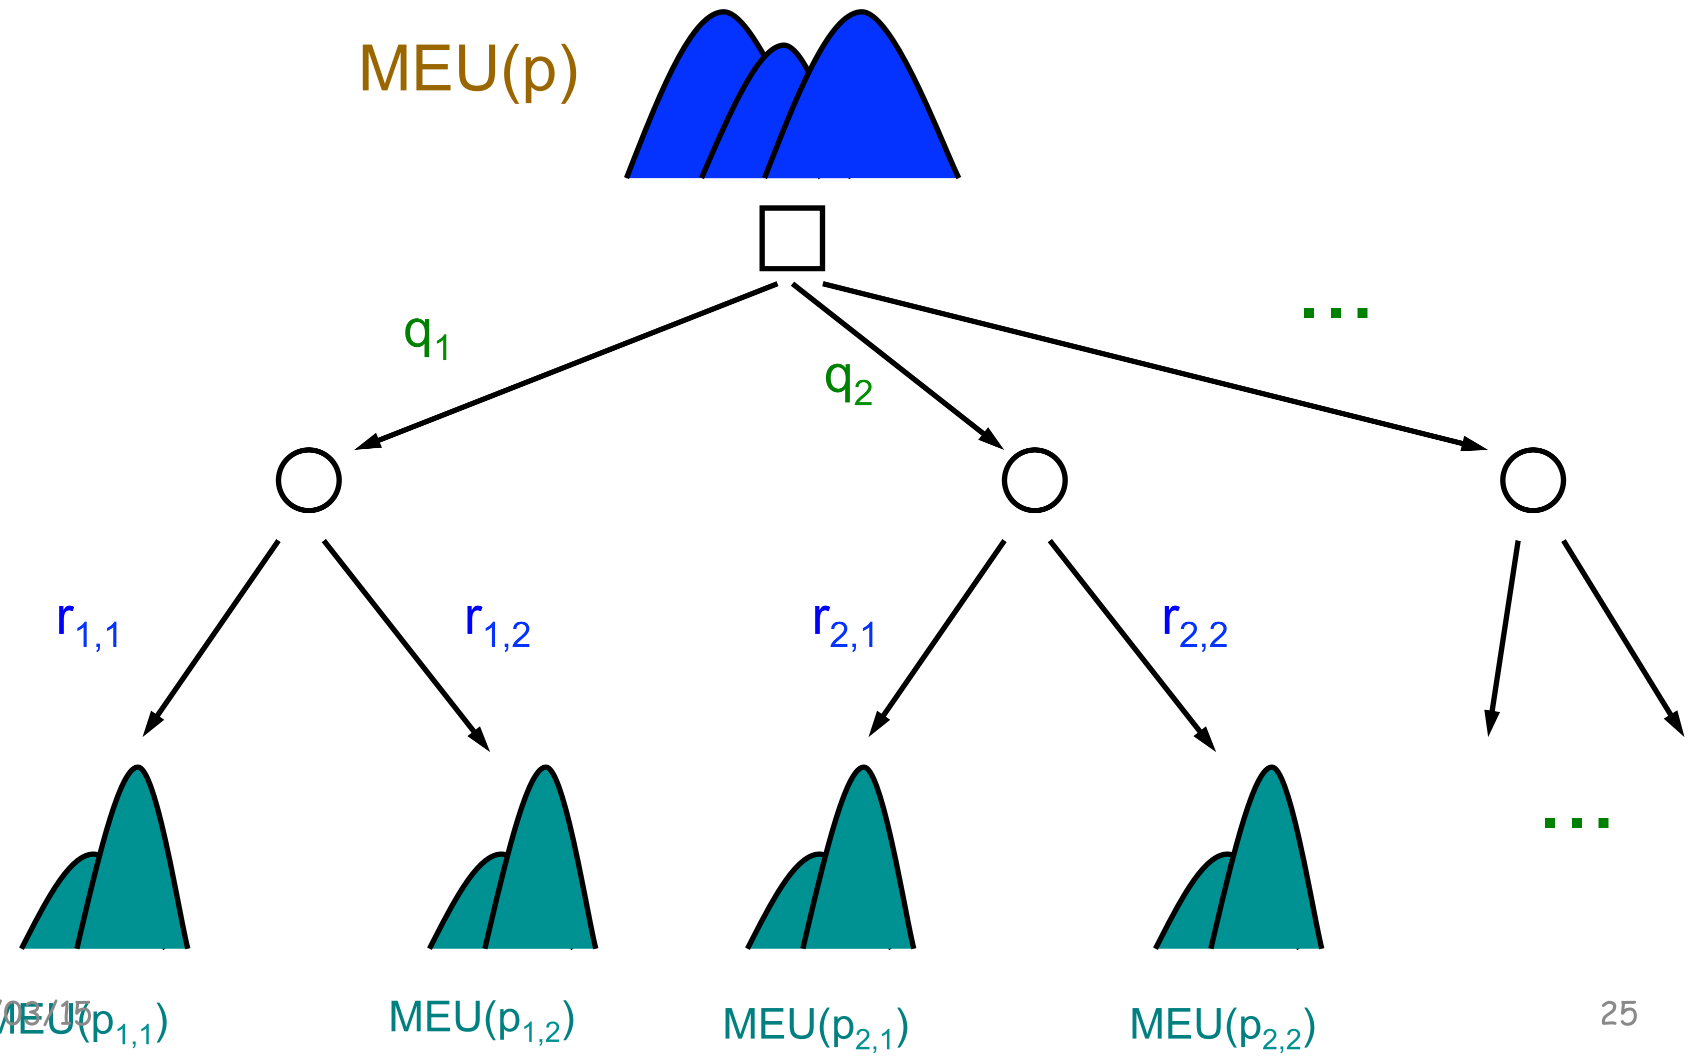

Expected Value of Information

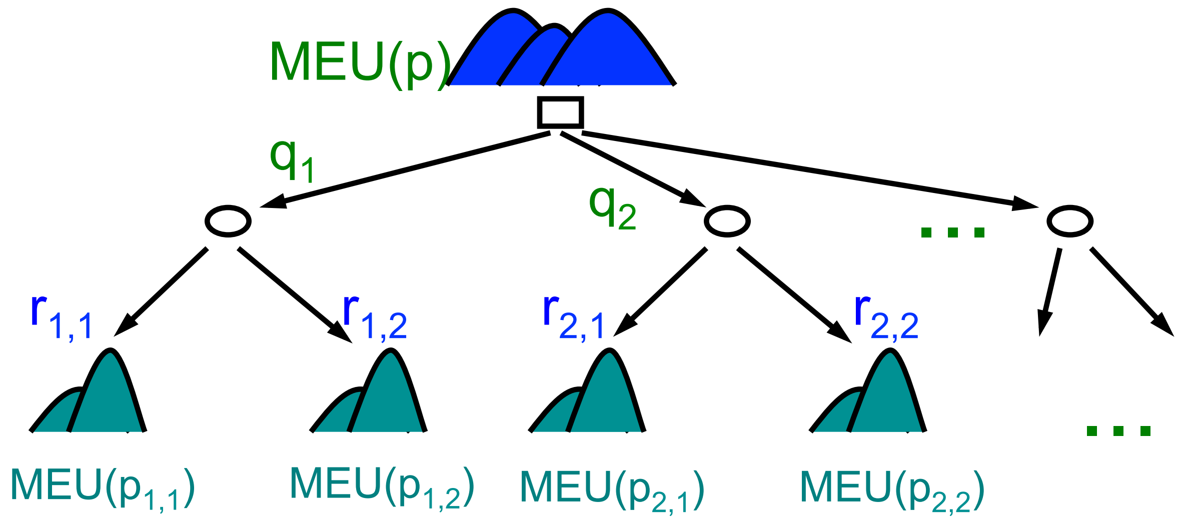

- $MEU(p) = E_p\ [u(x^*)]$
- Expected posterior utility: $EPU(q,p) = E_{r|q,p}\ [MEU(p^r)]$
- Expected value of information of query q:

$$EVOI(q) = EPU(q,p) - MEU(p)$$

Bayesian approach
Myopic EVOI

- Ask query with highest EVOI - cost
- [Chajewska et al ’00]
  - Global standard gamble queries (SGQ)

    “Is $u(o_i) > l$?”
  - Multivariate Gaussian distributions over utilities

- [Braziunas and Boutilier ’05]
  - Local Std. Gamble Query over utility factors
  - Mixture of uniforms distributions over utilities

Local elicitation in GAI models

[Braziunas and Boutilier '05]

$$u(\mathbf{x}) = u_1(\mathbf{x}_{I_1}) + \ldots + u_m(\mathbf{x}_{I_m})$$

- Local elicitation procedure
  - Bayesian uncertainty over local factors
  - Myopic EVOI query selection $\langle \mathbf{x}_i, l \rangle$
- Local comparison query

  "Is local value of factor setting x_i greater than l"?
  - Binary comparison query
  - Requires yes/no response
  - query point l can be optimized analytically

Preference Elicitation as POMDP

- States?
  - Utility function, U
- Belief states?
  - Probability distributions over U (which is n-dim continuous space)
- Actions?
  - Queries Q and decisions D
  - Queries induce no change in the underlying system state u , but do provide information
  - Each decision d is a terminal action
- There is a cost for asking questions, $c(q)$

Preference Elicitation as POMDP, contd.

- $Reward(d, u) = EU(d, u)$
- Sensor Model?
  - Response Model $P(r_q | q, u)$
- Value expressions:

$$V^*(P) = \max_{a \in A} Q_a^*(P)$$

$$Q_i^*(l, P) = c(q_i(l)) + \gamma \sum_{r \in R} \Pr(r | q_i(l), P) V^*(P_r)$$