2534 Lecture 2: Utility Theory

- ***Tutorial on Bayesian Networks: Weds, Sept.17, 5-6PM, PT266***
- ***LECTURE ORDERING: Game Theory before MDPs? Or vice versa?***

- Preference orderings
- Decision making under strict uncertainty
- Preference over lotteries and utility functions
- Useful concepts
  - Risk attitudes, certainty equivalents
  - Elicitation and stochastic dominance
- Paradoxes and behavioral decision theory
- Multi-attribute utility models
  - preferential and utility independence
  - additive and generalized addition models

Why preferences?

- Natural question: why not specify behavior with *goals*?
- Preferences: *coffee* ≻ *OJ* ≻ *tea*
  - Natural goal: coffee
    - but what if unavailable? requires a 30 minute wait? …
  - allows alternatives to be explored in face of costs, infeasibility,…




Preference Orderings

- Assume (finite) outcome set X (states, products, etc.)
- *Preference ordering* $\succcurlyeq$ over X:
  - $y \succcurlyeq z$ interpreted as: "I (weakly) prefer y to z"
  - $y \succ z$ iff $y \succcurlyeq z$ and $z \not\succcurlyeq y$ (strict preference)
  - $y \sim z$ iff $y \succcurlyeq z$ and $y \succcurlyeq z$ (indifference, *incomparability?*)

- Conditions: $\succcurlyeq$ must be:
  - (a) *transitive*: if $x \succcurlyeq y$ and $y \succcurlyeq z$ then $x \succcurlyeq z$
  - (b) *connected* (orderable): either $y \succcurlyeq z$ or $z \succcurlyeq y$
  - i.e., a total preorder

Preference Orderings

- Total preorder: seems natural, but conditions reasonable?
  - implies (iff) strict relation $\succ$ is asymmetric and neg. transitive*
    - **if a not better than b, b not better than c, then a not better than c*
  - why connected? why transitive? (e.g., money pump)
- Are preference orderings enough?
  - decisions under certainty? under uncertainty?

- Exercise: what properties of $\succcurlyeq$, $\succ$ needed if you desire incomparability?

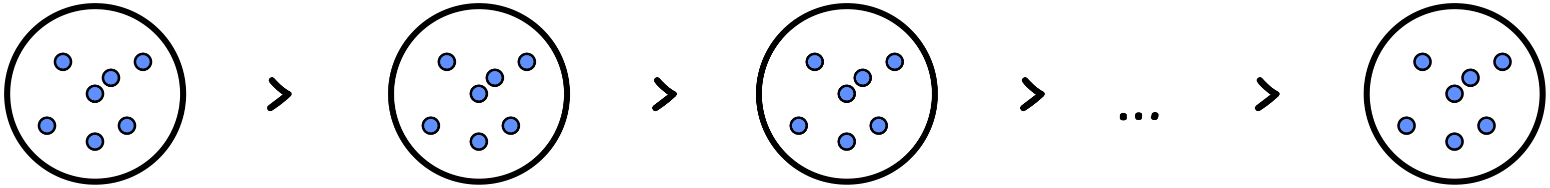

Revealed Preference

- Given a non-empty subset of $Y \subseteq X$, preferences "predict" choice: $c(Y) \in X$ should be a most preferred element
- More general *choice function*: select subset $c(Y) \subseteq Y$

- Given $\succ$, define $c(Y, \succ) = \{y \in Y : \nexists z \in Y \text{ s.t. } z \succ y\}$
  - i.e., the set of "top elements" of $\succ$ (works for partial orders too)
  - Exercise: show that $c(Y, \succ)$ must be non-empty
  - Exercise: show that if $y, z \in c(Y, \succ)$ then $y \sim z$

- CF c is *rationalizable* iff exists $\succ$ s.t. for all Y, $c(Y) = c(Y, \succ)$
  - are all choice functions rationalizable? (give counterexample)

Weak Axiom of Revealed Preference

- Desirable properties of choice functions:
  - (AX1) If $y \in Y$, $Y \subseteq Z$, and $y \in c(Z)$, then $y \in c(Y)$
  - (AX2) If $Y \subseteq Z$, $y,z \in c(Y)$, and $z \in c(Z)$, then $y \in c(Z)$

- **Thm:** (a) given prefs $\succ$, $c(\cdot,\succ)$ satisfies (AX1) and (AX2)

  (b) if c satisfies (AX1) and (AX2), then $c=c(\cdot,\succ)$ for some $\succ$
  - Exercise: prove this

- Thus: a characterization of rationalizable choice functions
- Weak axiom of revealed preference:
  - (WARP) If $y,z \in Y \cap Z$, $y \in c(Y)$, $z \in c(Z)$, then $y \in c(Z)$ and $z \in c(Y)$
  - Alternative characterization: c satisfies WARP iff (AX1) and (AX2)

Making Decisions: One-shot

- Basic model of (one-shot) decisions:
  - finite set of actions *A*, each leads to set of possible outcomes *X*
  - given preference ordering $\succcurlyeq$, is decision obvious?
- Deterministic actions: $f: A \longrightarrow X$
  - Let $f(A) = \{f(a) \in A\}$ be the set of possible outcomes, choose *a* with *most preferred* outcome: $c(f(A))$
  - preferences more useful than goals: what if *A* is set of *plans*?
- Is it always so straightforward?

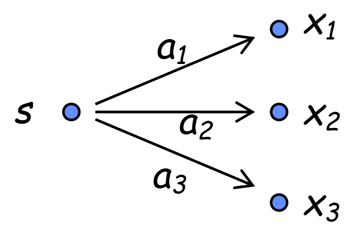

$x_1 \succ x_2 \succ x_3$: then choose a_1

Making Decisions: Uncertainty

- What if a given action has several possible outcomes
  - Nondeterministic actions: $f{:}A \longrightarrow \mathcal{P}(X)$
  - Stochastic actions: $f{:}A \longrightarrow \Delta(X)$
  - Initial state uncertainty (nondeterministic or stochastic)

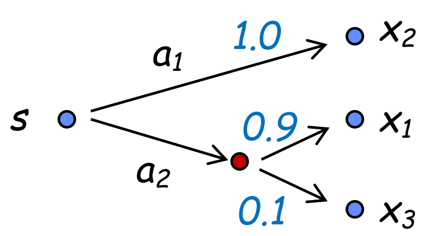

$x_1 \succ x_2 \succ x_3$: choose a_1 or a_2 ?

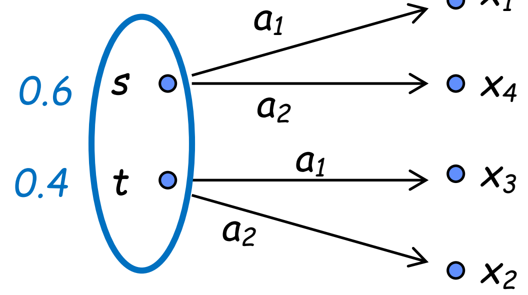

$x_1 \succ x_2 \succ x_3 \succ x_4$: choose a_1 or a_2 ?

Making Decisions: Uncertainty

- Two solutions to this problem:

- Soln 1: Assign values to outcomes
  - decision making under *strict uncertainty* if nondeterministic
  - *expected value/utility theory* if stochastic
  - **Question:** where do values come from? what do they mean?

- Soln 2: Assign preferences to lotteries over outcomes
  - decision making under quantified uncertainty

Making Decisions: Strict Uncertainty

- Suppose you have no way to quantify uncertainty, but each outcome has some "value" to you
  - require the value function respect $\succcurlyeq$: $v(x) \geq v(y)$ iff $x \succcurlyeq y$
- Useful to specify a decision table
  - rows: *actions*; columns: *states of nature*; entries: *values*
  - unknown states of nature dictate outcomes, table has: $v(f(a,\Theta_1))$

	Θ_1	Θ_2	…	Θ_k
a_1	v_{11}	v_{12}	…	v_{1k}
a_2	v_{21}	v_{22}	…	v_{2k}
…				
a_n	v_{n1}	v_{n2}	…	v_{nk}

Strict Uncertainty: Decision Criteria

- *Maximin (Wald):* choose action with *best* worst outcome
  - $\max_a \min_\Theta v(f(a,\Theta))$
  - *a* with max *security level s(a)*
  - very pessimistic
- *Maximax:* choose action with *best* best outcome
  - $\max_a \max_\Theta v(f(a,\Theta))$
  - *a* with max *optimism level o(a)*
- Hurwicz criterion: set $\alpha \in (0,1)$
  - $\max_a \alpha s(a) + (1-\alpha) o(a)$

	Θ_1	Θ_2	Θ_3	Θ_4
a_1	2	2	0	1
a_2	1	1	1	1
a_3	0	4	0	0
a_4	1	3	0	0

- Maximin: a_2
- Maximax: a_3
- Hurwicz: *which decisions are possible?*
- *What if* a_3 = <0.5 3 2 2>?

Minimax Regret (Savage)

- *Regret* of a_i under outcome Θ_j: $r_{ij} = max\ \{v_{kj}\} - v_{ij}$
  - How sorry I'd be doing a_i if I'd known Θ_j was coming
  - Why worry about *worst* outcome: beyond my control
- *Minimax regret:* choose $arg\ min_a\ max_j\ r_{ij}$

	Θ_1	Θ_2	Θ_3	Θ_4	*Max Regret*
a_1	2 / 0	2 / 2	0 / 1	1 / 0	2
a_2	1 / 1	1 / 3	1 / 0	1 / 0	3
a_3	0 / 2	4 / 0	0 / 1	0 / 1	2
a_4	1 / 1	3 / 1	0 / 1	0 / 1	1

*red values are regrets r_{ij}

Qualitative Criteria: Reasonable?

- Criteria all make sense at some level, but not at others
  - indeed, all have “faults”
- *Independence of irrelevant alternatives (IIA):* adding an action to decision problem does not influence relative ranking of other actions

	Θ_1	Θ_2	Θ_3
a_1	6 / 0 / 0	9 / 0 / 0	3 / 1 / 5
a_2	2 / 4 / 4	9 / 0 / 0	4 / 0 / 4
a_3	0 / - / 6	0 / - / 9	8 / - / 0

**red: regrets r_{ij} without a3*

**green: regrets r_{ij} with a3*

- Minimax regret violates IIA
  - a_1 lower MR than a_2 (no a_3)
  - a_2 lower MR than a_1 (with a_3)
- Classic impossibility result:
  - no qualitative decision criterion satisfies all of a set of intuitively reasonable principles (like IIA)

Making Decisions: Probabilistic Uncertainty

- What if:
  - 2% chance no coffee made (30 min delay)? 10%? 20%? 95%?
  - robot has enough charge to check only *one* possibility
  - 5% chance of *damage* in coffee room, 1% at OJ vending mach.

Preference over Lotteries

- If uncertainty in action/choice outcomes, $\succcurlyeq$ not enough
- Each action is a "lottery" over outcomes

- A *simple lottery* over *X* has form:

$$l = [\,(p_1, x_1), (p_2, x_2), \ldots, (p_n, x_n)\,]$$

  where $p_i \geq 0$ and $\sum p_i = 1$
  - outcomes are just trivial lotteries (one outcome has prob 1)

- A *compound lottery* allows outcomes to be lotteries:

$$[\,(p_1, l_1), (p_2, l_2), \ldots, (p_n, l_n)\,]$$

  - restrict to finite compounding

Constraints on Lotteries

- Continuity:
  - If $x_1 \succ x_2 \succ x_3$ then $\exists\, p$ s.t. $[(p,x_1), (1-p,x_3)] \sim x_2$
- Substitutability:
  - If $x_1 \sim x_2$ then $[(p,x_1), (1-p,x_3)] \sim [(p,x_2), (1-p,x_3)]$
- Mononoticity:
  - If $x_1 \succcurlyeq x_2$ and $p \geq q$ then $[(p,x_1), (1-p,x_2)] \succcurlyeq [(q,x_1), (1-q,x_2)]$
- Reduction of Compound Lotteries ("no fun gambling"):
  - $[\ (p, [(q,x_1), (1-q,x_2)]\ ), (1-p, [(q',x_3), (1-q',x_4)])\ ]$

    $\sim [\ (pq,x_1), (p-pq,x_2), (q'-pq',x_3), ((1-p)(1-q'),x_4)\ ]$
- Nontriviality:
  - $x_T \succ x_\perp$

Implications of Properties on $\succcurlyeq$

- Since $\succcurlyeq$ is transitive, connected: representable by ordinal value function *V(x)*

- With constraints on lotteries: we can construct a *utility function* $U(l) \in \mathbf{R}$ s.t. $U(l_1) \geq U(l_2)$ iff $l_1 \succcurlyeq l_2$
  - where $U([\,(p_1, x_1), \ldots, (p_n, x_n)\,]) = \sum_i p_i U(x_i)$
  - famous result of Ramsey, von Neumann & Morgenstern, Savage

- Exercise: prove existence of such a utility function
- Exercise: given any *U* over outcomes *X*, show that ordering $\succcurlyeq$ over lotteries induced by *U* satisfies required properties of $\succcurlyeq$

Implications of Properties on ≽

- Assume some collection of actions/choices at your disposal

- Knowing *$U(x_i)$* for each *outcome* allows tradeoffs to be made over uncertain courses of action (lotteries)
  - simply compute expected utility of each course of action

- **Principle of Maximum Expected Utility (MEU)**
  - utility of choice is a expected utility of its outcome
  - appropriate choice is that with *maximum expected utility*
  - *Why? Action (lottery) with highest EU is the action (lottery) that is most preferred in ordering ≽ over lotteries!*

Some Discussion Points

- Utility function existence: proof is straightforward
  - Hint: set $U(x_T) = 1$; $U(x_\perp) = 0$; find a p s.t. $x \sim [(p, x_T), (1-p, x_\perp)]$

- Utility function for > over lotteries is not unique:
  - any positive affine transformation of *U* induces same ordering >
  - normalization in range [0,1] common
- Ordinal preferences "easy" to elicit (if *X* small)
  - cardinal utilities trickier for people: an "art form" in decision anal.
- Outcome space often factored: exponential size
  - requires techniques of multi-attribute utility theory (MAUT)
- Expected utility accounts for risk attitudes: inherent in preferences over lotteries
  - see utility of money (next)

Risk profiles and Utility of money

- What would you choose?
  - (a) \$100,000 or (b) [(.5, \$200,000), (.5, 0)]
  - what if (b) was \$250K, \$300K, \$400K, \$1M; p = .6, .7, .9, .999, …
  - generally, *U(**EMV**(lottery)) > U(lottery)* *EMV = expected monetary value*
- Utility of money is nonlinear: e.g., *U(\$100K) > .5U(\$200K)+.5U(\$0)*
- Certainty equivalent of *l*: $U(CE) = U(l); CE = U^{-1}(EU(l))$

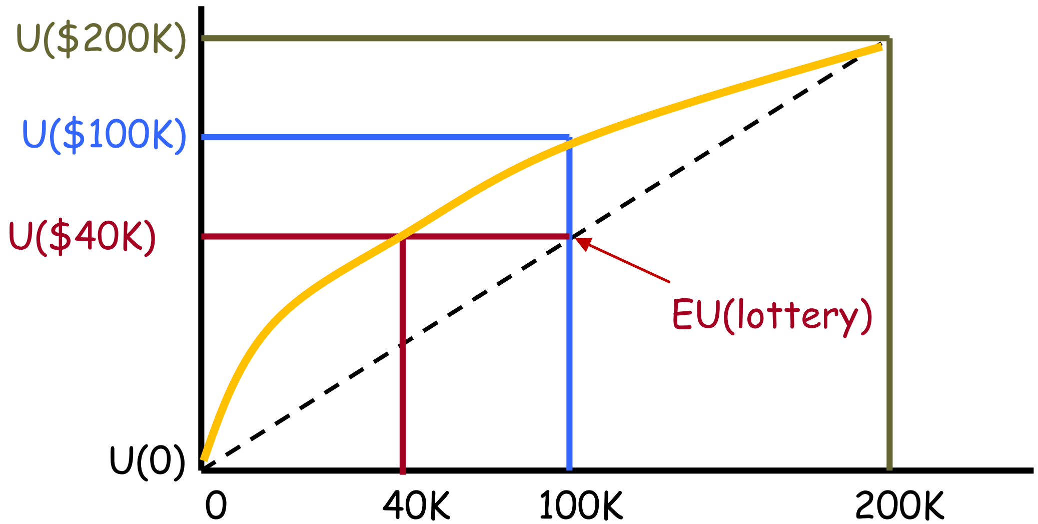

For many people, CE ~ \$40K
Note: 2nd \$100K "worth less" than 1st \$100K

Risk attitudes

- *Risk Premium: EMV(l) – CE(l)*
  - how much of EMV will I give up to remove risk of losing
- Risk averse:
  - decision maker has positive risk premium; *U(money)* is concave
- Risk neutral:
  - decision maker has zero risk premium; *U(money)* is linear
- Risk seeking:
  - decision maker has negative risk premium; *U(money)* is convex
- Most people are risk averse
  - this explains insurance
  - often risk seeking in negative range
  - linear a good approx in small ranges

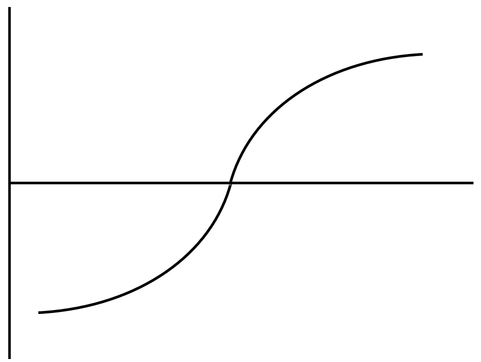

St. Peterburg Paradox

- How much would you pay to play this game?
  - A coin is tossed until it falls heads. If it occurs on the Nth toss you get \$$2^N$

$$EMV = \sum_{n=1}^{\infty} \left(\frac{1}{2}\right)^n 2^n = \sum_{n=1}^{\infty} 1 = \infty$$

  - Most people will pay about \$2-\$20
- Not a paradox *per se*… doesn’t contradict utility theory

A Game

- Situation 1: choose either
  - (1) $1M, Prob=1.00
  - (2) $5M, Prob=0.10; $1M, Prob=0.89; nothing, Prob=0.01

Another Game

- Situation 2: choose either
  - (3) $1M, Prob=0.11; nothing, Prob=0.89
  - (4) $5M, Prob=0.10; nothing, Prob=0.90

Allais' Paradox

- Situation 1: choose either
  - (1) \$1M, Prob=1.00
  - (2) \$5M, Prob=0.10; \$1M, Prob=0.89; nothing, Prob=0.01
- Situation 2: choose either
  - (3) \$1M, Prob=0.11; nothing, Prob=0.89
  - (4) \$5M, Prob=0.10; nothing, Prob=0.90
- Most people: (1) > (2) and (4) > (3)
  - e.g., in related setups: 65% (1) > (2); 25% (3) > (4)
- Paradox: no way to assign utilities to monetary outcomes that conforms to expected utility theory and the stated preferences (violates substitutability)
  - possible explanation: *regret*

Allais' Paradox: The Paradox

- Situation 1: choose either
  - (1) \$1M, Prob=1.00
    - *equiv: (\$1M 0.89; \$1M 0.11)*
  - (2) \$5M, Prob=0.10; \$1M, Prob=0.89; nothing, Prob=0.01
  - **So if (1)>(2), by subst: \$1M > (\$5M 10/11; nothing 1/11)**

- Situation 2: choose either
  - (3) \$1M, Prob=0.11; nothing, Prob=0.89
  - (4) \$5M, Prob=0.10; nothing, Prob=0.90
    - *equiv: nothing 0.89; \$5M 0.10; nothing 0.01*
  - **So if (4)>(3), by subst: (\$5M 10/11; nothing 1/11) > \$1M**

…and the Fall 2014 survey says

- Situation 1:
  - (1)>(2): a (x%)
  - (2)>(1): b (y%)
- Situation 2:
  - (3)>(4): c (w%)
  - (4)>(3): d (z%)
- *The 2534 class of 2014 is ____________________*
  - *many people who take a class on decision theory tend to think in terms of expected monetary value (so 2534 surveys tend to be consistent than more standard empirical results; however, if there was real money on the line, my guess is the proportions would be somewhat more in line with experiments)*

Ellsberg Paradox

- Urn with 30 red balls, 60 yellow or black balls; well mixed
- Situation 1: choose either
  - (1) $100 if you draw a red ball
  - (2) $100 if you draw a black ball
- Situation 2: choose either
  - (3) $100 if you draw a red or yellow ball
  - (4) $100 if you draw a black or yellow ball
- Most people: (1) > (2)   and (4) > (3)
- Paradox: no way to assign utilities (all the same) and beliefs about yellow/black proportions that conforms to expected utility theory
  - possible explanation: *ambiguity aversion*

Utility Representations

- Utility function $u: X \rightarrow [0,1]$
  - decisions induce distribution over outcomes
  - *or* we simply choose an outcome (no uncertainty), but constraints on outcomes
- If *X* is combinatorial, sequential, etc.
  - representing, eliciting *u* difficult in explicit form

Product Configuration*

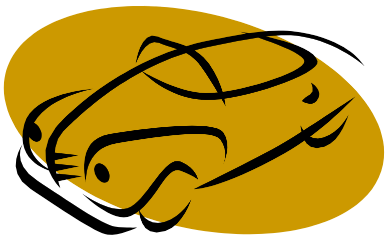

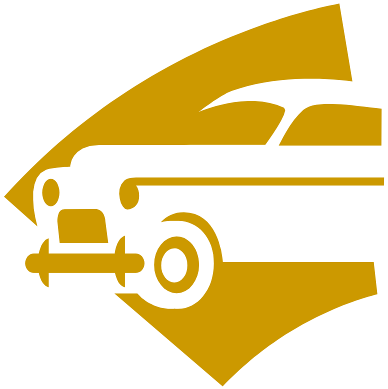

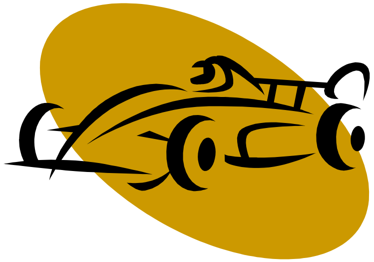

Luggage Capacity?
Two Door? Cost?
Engine Size?
Color? Options?

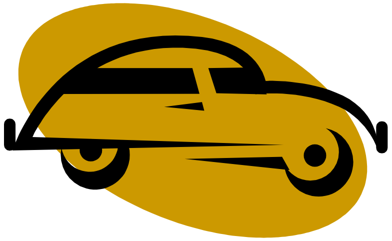

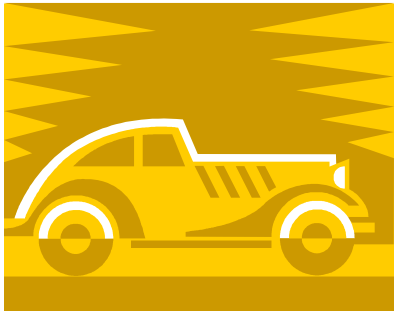

COACH*

- POMDP for prompting Alzheimer's patients
  - solved using factored models, value-directed compression of belief space
- Reward function (patient/caregiver preferences)
  - indirect assessment (observation, policy critique)

Winner Determination in Combinatorial Auctions

- *Expressive bidding* in auctions becoming common
  - expressive languages allow: combinatorial bids, side-constraints, discount schedules, etc.
  - direct expression of utility/cost: economic efficiency
- Advances in *winner determination*
  - determine least-cost allocation of business to bidders
  - new optimization methods key to acceptance
  - applied to large-scale problems (e.g., sourcing)

Non-price Preferences

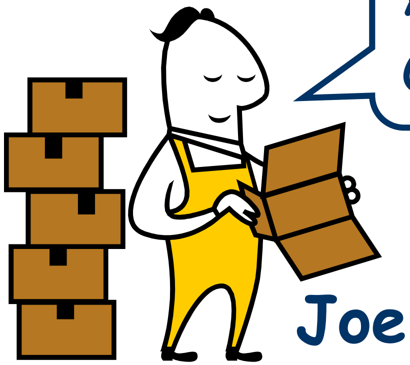

A and B for $12000.
C and D for $5000...

A for $10000.
B and D for $5000 if A;
B and D for $7000 if not A...

etc...

A, C to Fred.
B, D, G to Frank.
F, H, K to Joe...
Cost: $57,500.

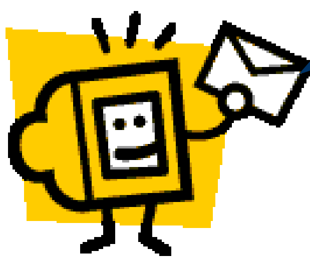

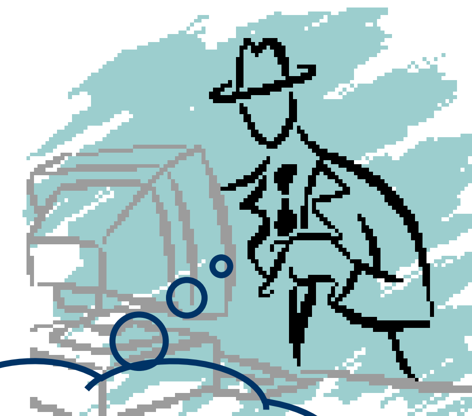

That gives too much business to Joe!!

Non-price Preferences

- WD algorithms *minimize cost alone*
  - but preferences for *non-price attributes* play key role
  - Some typical attributes in sourcing:
    - *percentage volume business to specific supplier*
    - *average quality of product, delivery on time rating*
    - *geographical diversity of suppliers*
    - *number of winners (too few, too many), …*
- Clear utility function involved
  - difficult to articulate precise tradeoff weights
  - "What would you pay to reduce *%volumeJoe* by 1%?"

Manual Scenario Navigation*

- Current practice: manual *scenario navigation*
  - impose constraints on winning allocation
    - ***not a hard constraint!***
  - re-run winner determination
  - new allocation satisfying constraint: higher cost
  - assess tradeoff and repeat (often hundreds of times) until satisfied with some allocation

Utility Representations

- Utility function $u: X \rightarrow [0,1]$
  - decisions induce distribution over outcomes
  - *or* we simply choose an outcome (no uncertainty), but constraints on outcomes
- If X is combinatorial, sequential, etc.
  - representing, eliciting u difficult in explicit form
- Some structural form usually assumed
  - so u parameterized compactly (weight vector w)
  - e.g., linear/additive, generalized additive models

- *Representations for qualitative preferences, too*
  - *e.g., CP-nets, TCP-nets, etc.* ***[BBDHP03, BDS05]***

Flat vs. Structured Utility Representation

- Naïve representation: vector of values
  - e.g., *car7:1.0, car15:0.92, car3:0.85, …, car22:0.0*
- Impractical for combinatorial domains
  - e.g., can't enumerate exponentially many cars, nor expect user to assess them all (choose among them)
- Instead we try to exploit independence of user preferences and utility for different attributes
  - the relative preference/utility of one attribute is independent of the value taken by (some) other attributes
- Assume $X \subseteq Dom(X_1) \times Dom(X_2) \times \ldots Dom(X_n)$
  - e.g., *car7: Color=red, Doors=2, Power=320hp, LuggageCap=0.52m*^{3}

Preferential, Utility Independence

- $\boldsymbol{X}$ and $\boldsymbol{Y} = \boldsymbol{V}$-$\boldsymbol{X}$ are *preferentially independent* if:
  - $\boldsymbol{x}_1\boldsymbol{y}_1 \geq \boldsymbol{x}_2\boldsymbol{y}_1$ iff $\boldsymbol{x}_1\boldsymbol{y}_2 \geq \boldsymbol{x}_2\boldsymbol{y}_2$ (for all $\boldsymbol{x}_1, \boldsymbol{x}_2, \boldsymbol{y}_1, \boldsymbol{y}_2$)
  - e.g., *Color: red>blue* regardless of value of *Doors, Power, LugCap*
  - conditional P.I. given set $\boldsymbol{Z}$: definition is straightforward

- $\boldsymbol{X}$ and $\boldsymbol{Y} = \boldsymbol{V}$-$\boldsymbol{X}$ are *utility independent* if:
  - $l_1(\boldsymbol{X}\boldsymbol{y}_1) \geq l_2(\boldsymbol{X}\boldsymbol{y}_1)$ iff $l_1(\boldsymbol{X}\boldsymbol{y}_2) \geq l_2(\boldsymbol{X}\boldsymbol{y}_2)$ (for all $\boldsymbol{y}_1, \boldsymbol{y}_2$, all distr. l_1, l_2)
  - e.g., preference for *lottery(Red,Green,Blue)* does not vary with value of *Doors, Power, LugCap*
    - implies existence of a "utility" function over local (sub)outcomes
  - conditional U.I. given set $\boldsymbol{Z}$: definition is straightforward

Question

- Is each attribute PI of others in preference relation 1? 2?

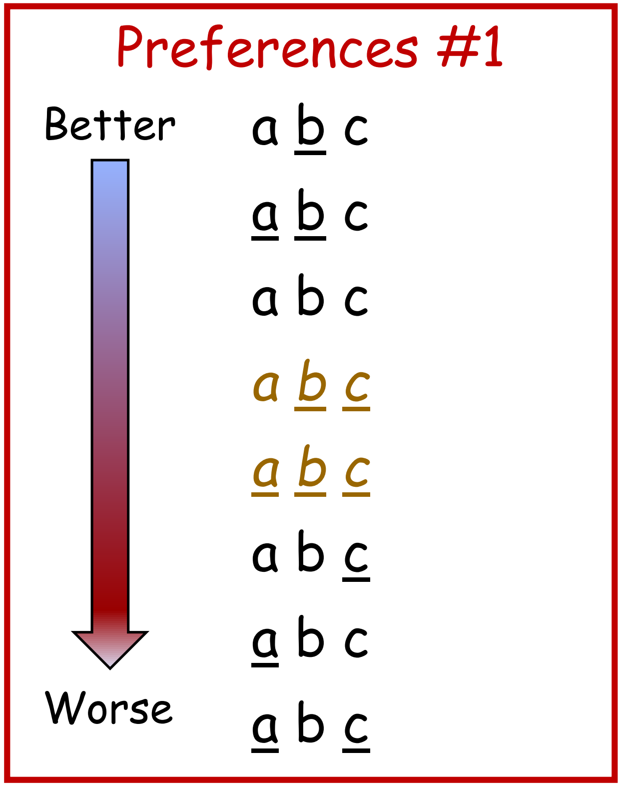

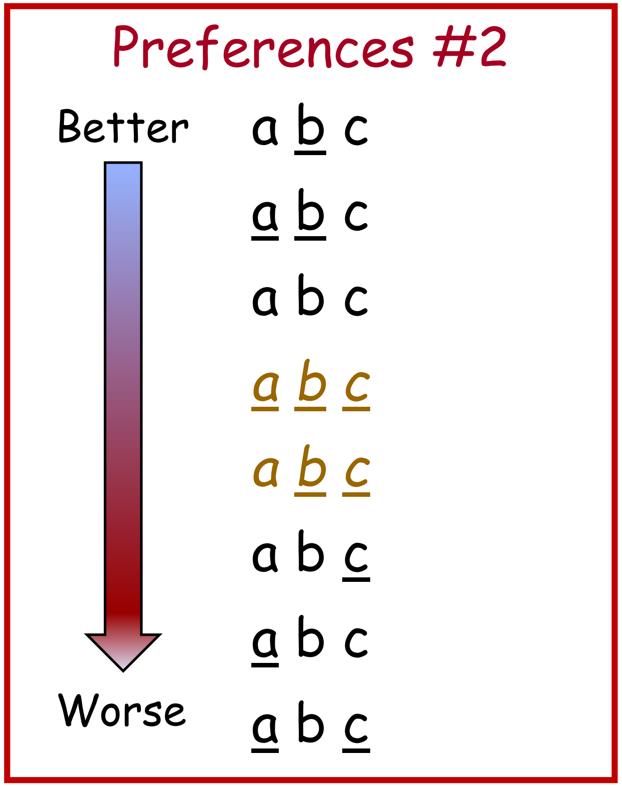

- Does UI imply PI? Does PI imply UI?

Additive Utility Functions

- *Additive representations* commonly used **[KR76]**
  - breaks exponential dependence on number of attributes
  - use sum of *local utility functions* u_i over attributes
  - or equivalently *local value functions* v_i plus scaling factors λ_i

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

  - e.g., *$U(Car) = 0.3\ v_1(Color) + 0.2\ v_2(Doors) + 0.5\ v_3(Power)$* and *$v_1(Color)$: cherryred:1.0, metallicblue:0.7, …, grey:0.0*
- This will make elicitation much easier (more on this next time)
- It *can* also make optimization more practical (more next time)

Additive Utility Functions

- 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

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

Generalized Additive Utility

- *Generalized additive models* more flexible
    - *interdependent value additivity* **[Fishburn67]**, GAI **[BG95]**
  - assume (overlapping) set of *m* subsets of vars ***X[j]***
  - use sum of *local utility functions* u_j over attributes

$$u(\mathbf{x}) = \sum_{j=1}^{m} u_j(\mathbf{x}_j)$$

  - e.g., *U(Car) = 0.3* v_1*(Color,Doors) + 0.7* v_2*(Doors,Power)* with

    v_1*(Color,Door) : blue,sedan:1.0; blue,coupe:0.7;blue,hatch:0.1,*
    red, sedan: 0.8, red,coupe:0.9; red,hatch:0.0

- This will make elicitation much easier (more on this next time)
- It *can* also make optimization more practical (more next time)

GAI Utility Functions

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

$$u(\mathbf{x}) = \sum_{j=1}^{m} u_j(\mathbf{x}_j)$$

Further Background Reading

- John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior. Princeton University Press, Princeton, 1944.*
- L. Savage. *The Foundations of Statistics. Wiley, NY, 1954.*
- R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-offs. Wiley, NY, 1976.*
- P. C. Fishburn. Interdependence and additivity in multivariate, unidimensional expected utility theory. *International Economic Review, 8:335–342, 1967.*
- Peter C. Fishburn. *Utility Theory for Decision Making. Wiley, New York, 1970.*
- F. Bacchus , A. Grove. Graphical models for preference and utility. *UAI-95, pp.3–10, 1995.*
- S. French, *Decision Theory*, Halsted, 1986.