MAN course

course page: www.inf.ed.ac.uk/teaching/courses/man

Mondays and Thursdays 5:10-6:00 pm - WRB G.04

next class Monday Sep 27 !

course strictly based on “networks, crowds, markets”:
www.cs.cornell.edu/home/kleinber/networks-book/

coursework (week 5) 30%
exam 70%

background: elementary probabilities & calculus

1-line
summary

socio, techno, eco, bio things happening on/structured by a network social networks (friendship, acquaintance, coboardism, co-affiliation, etc), ecological networks, web pages, citation networks, intra-organisational communication (eg Enron's emails), Internet physical structure, power grids, financial and economical markets, neural systems, intra-celullar networks, etc.

N. Schwartz NYT May 1st 2010

Banks and governments in these five shaky economies owe each other many billions of euros — converted here to dollars — and have even larger debts to Britain, France and Germany. Arrow widths are proportional to debt amounts.

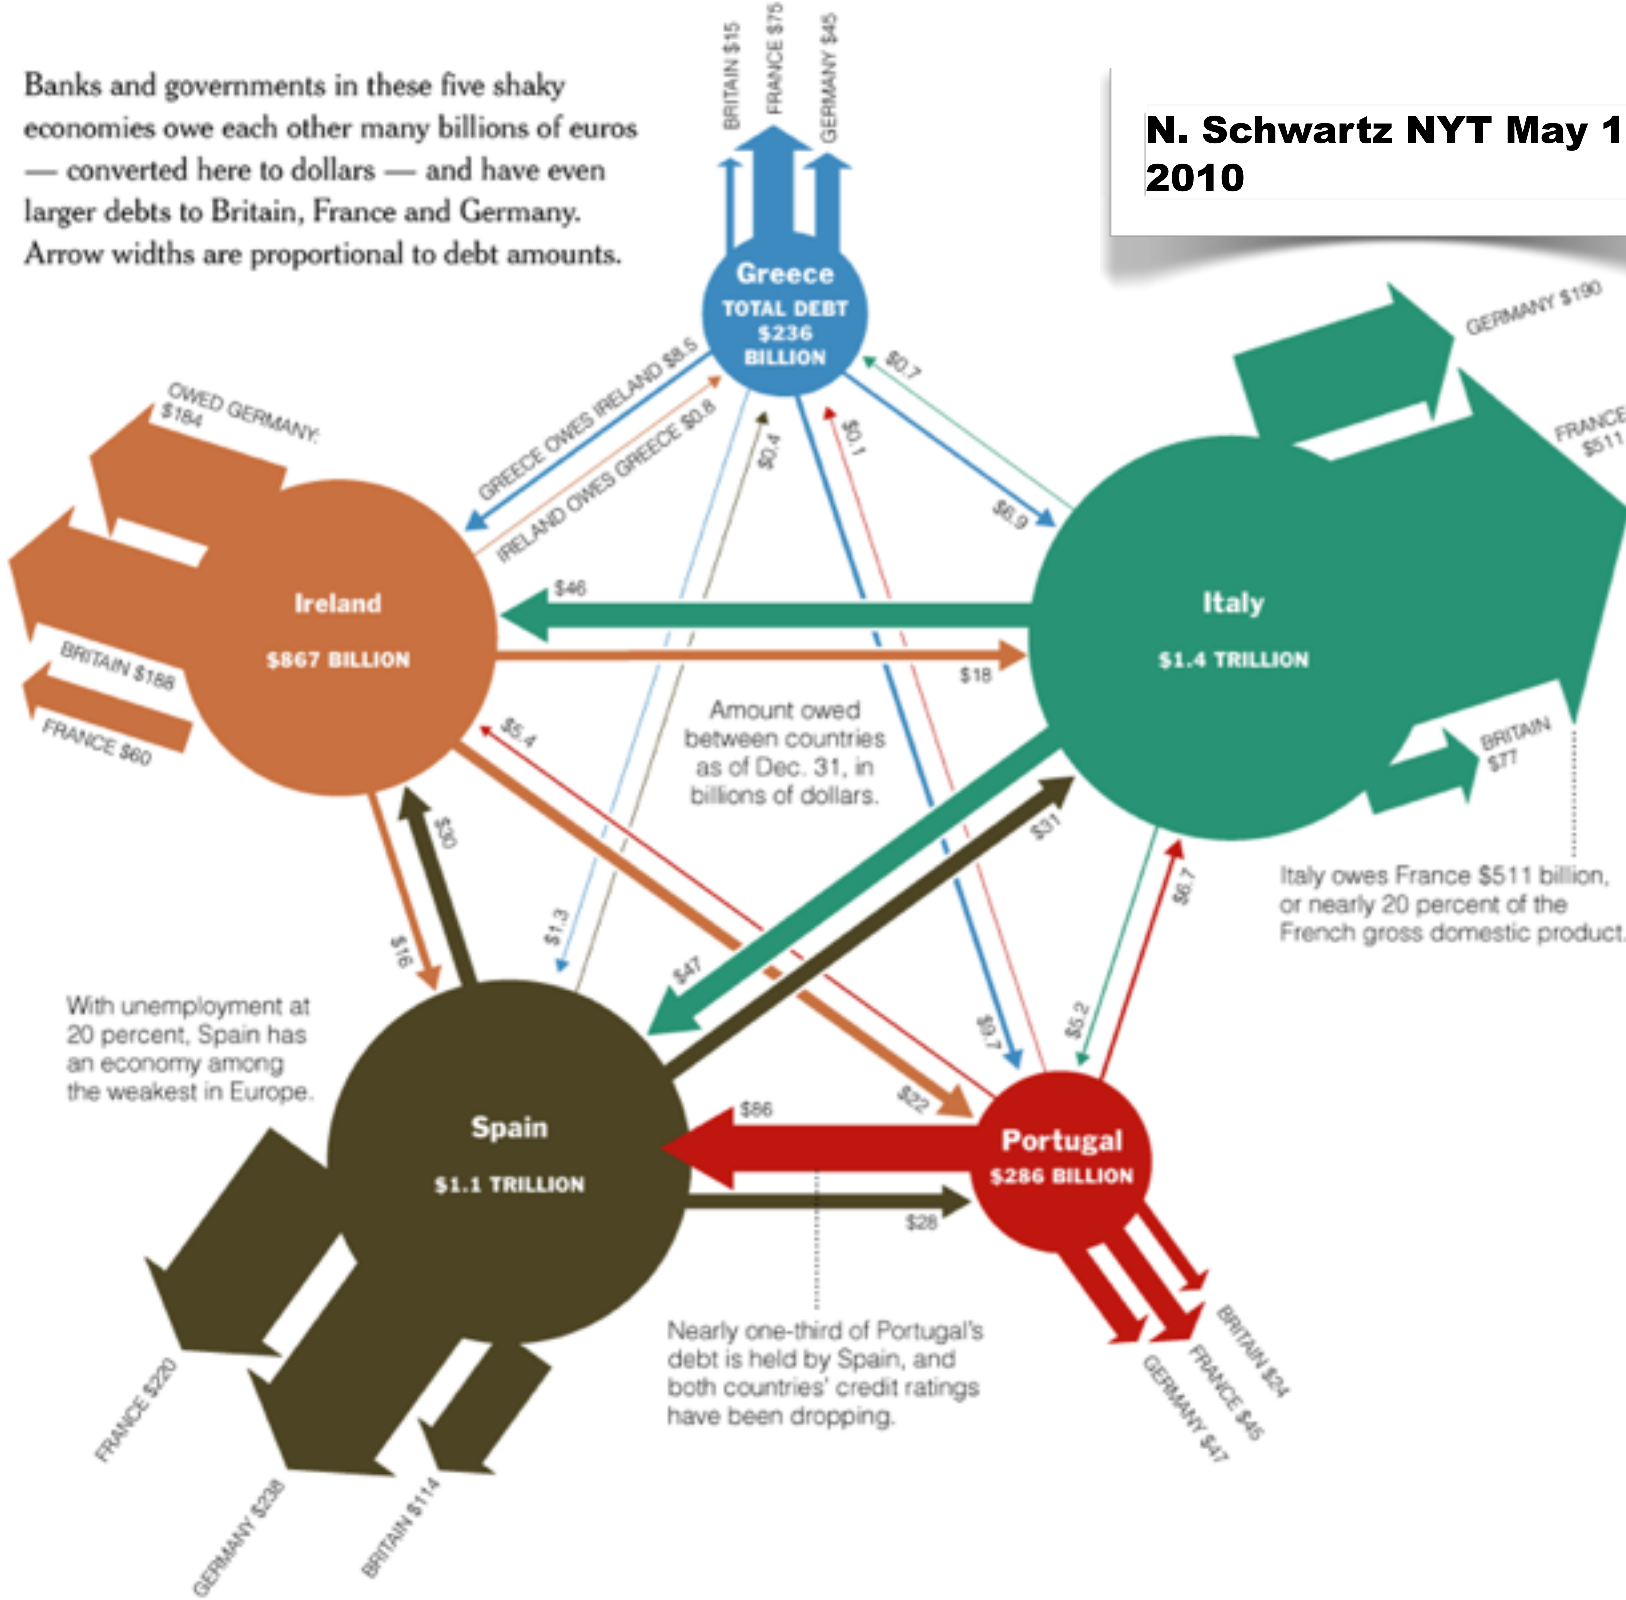

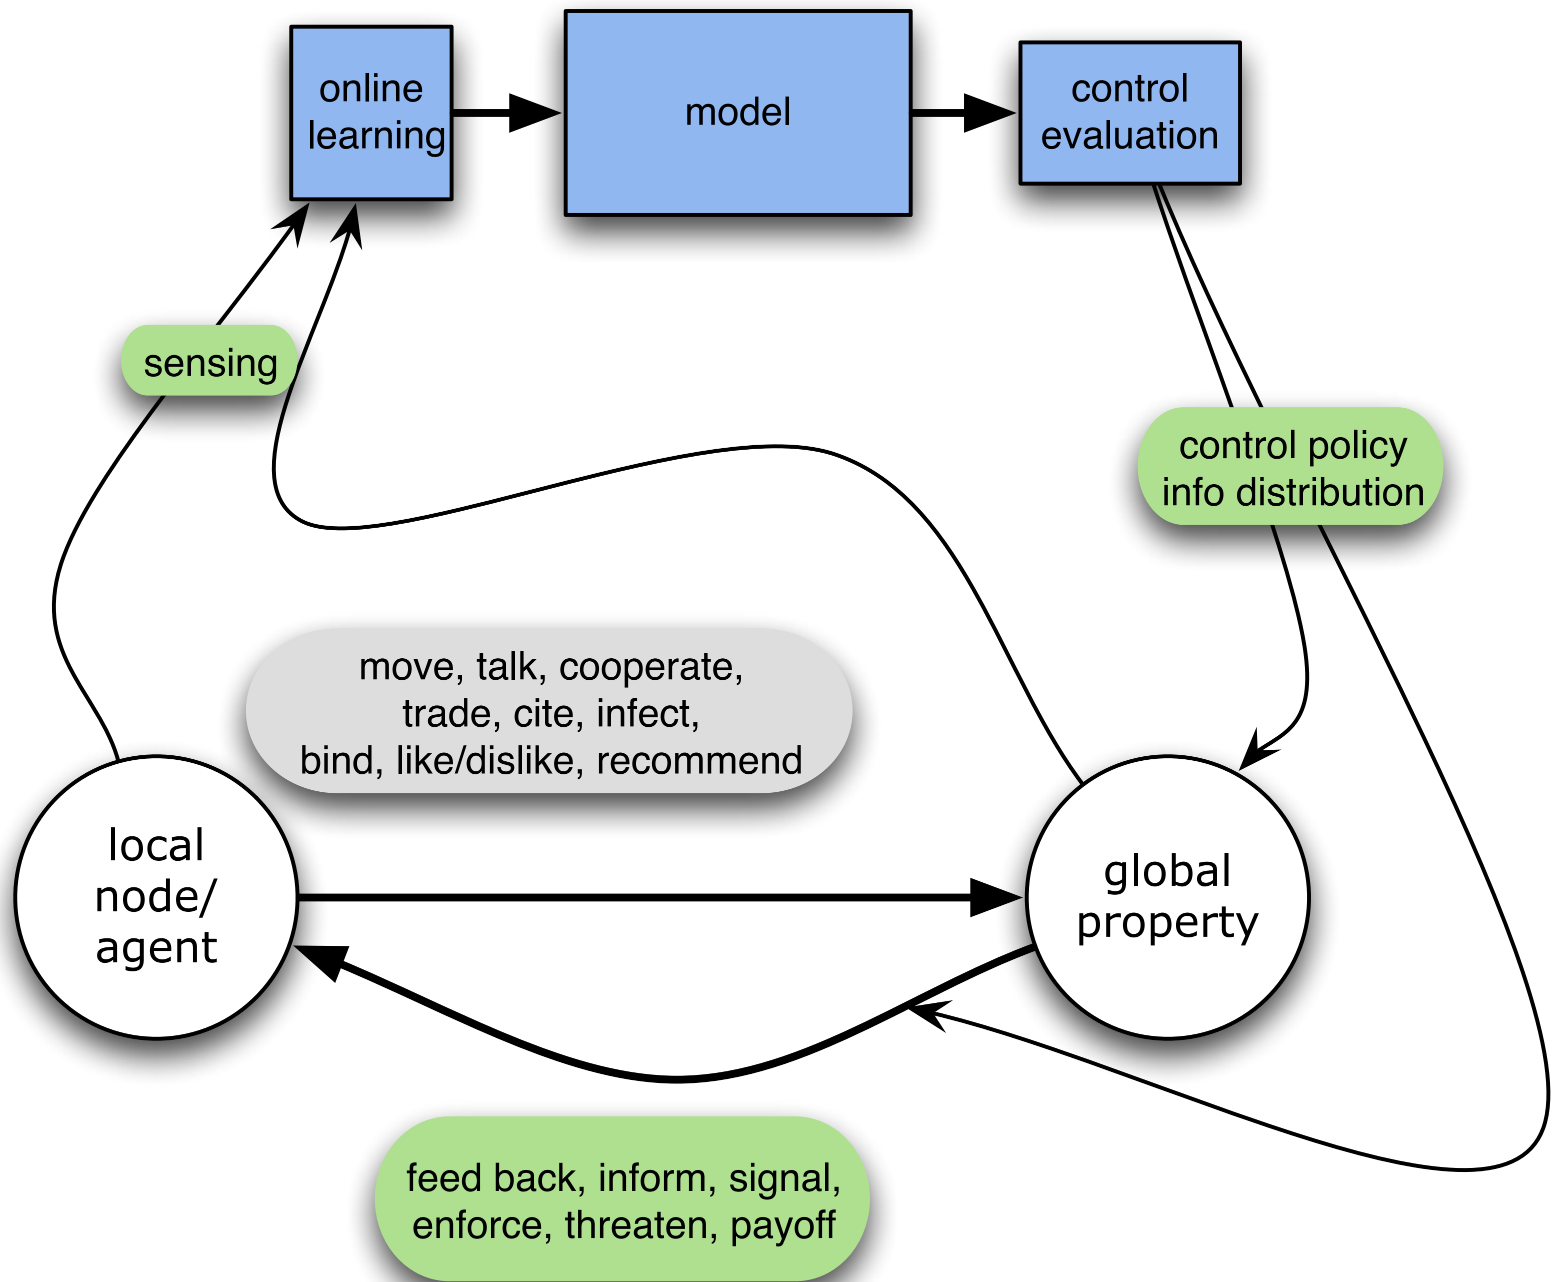
online learning
model
control evaluation
sensing
control policy info distribution
move, talk, cooperate, trade, cite, infect, bind, like/dislike, recommend
local node/ agent
global property
feed back, inform, signal, enforce, threaten, payoff

markets with exogenous events [ncw ch22]

- agents have beliefs (expectations)

- agents take actions under uncertainty about the outcome (bet on A, buy/sell stock)

- decisions are functions of their beliefs **and** of their relation to risk

- a market turns the set of actions into a price and hence a payoff (aggregation)

outcomes are independent of agent choices (ie we assume **exogeneous** outcomes)

2 horse race A and B

- agents have beliefs pA, pB

- agents take actions rA, rB
where rA = fraction of w bet on A
so rA + rB = 1

- decisions are functions of beliefs, payoffs
and relation to risk

- a market turns the set of actions into agents
pay-off using **odds** oA, oB

outcomes are independent of agent choices (ie there is no cheating unlike with eg sumo fighting)

odds?

odds:

oA = 3-to-1 := one gets 3£ for a successful 1£ bet

equivalently, a bet of 1/3£ gets 1£ if successful

1/oA is the price of a contract which is worth 1£ if A wins

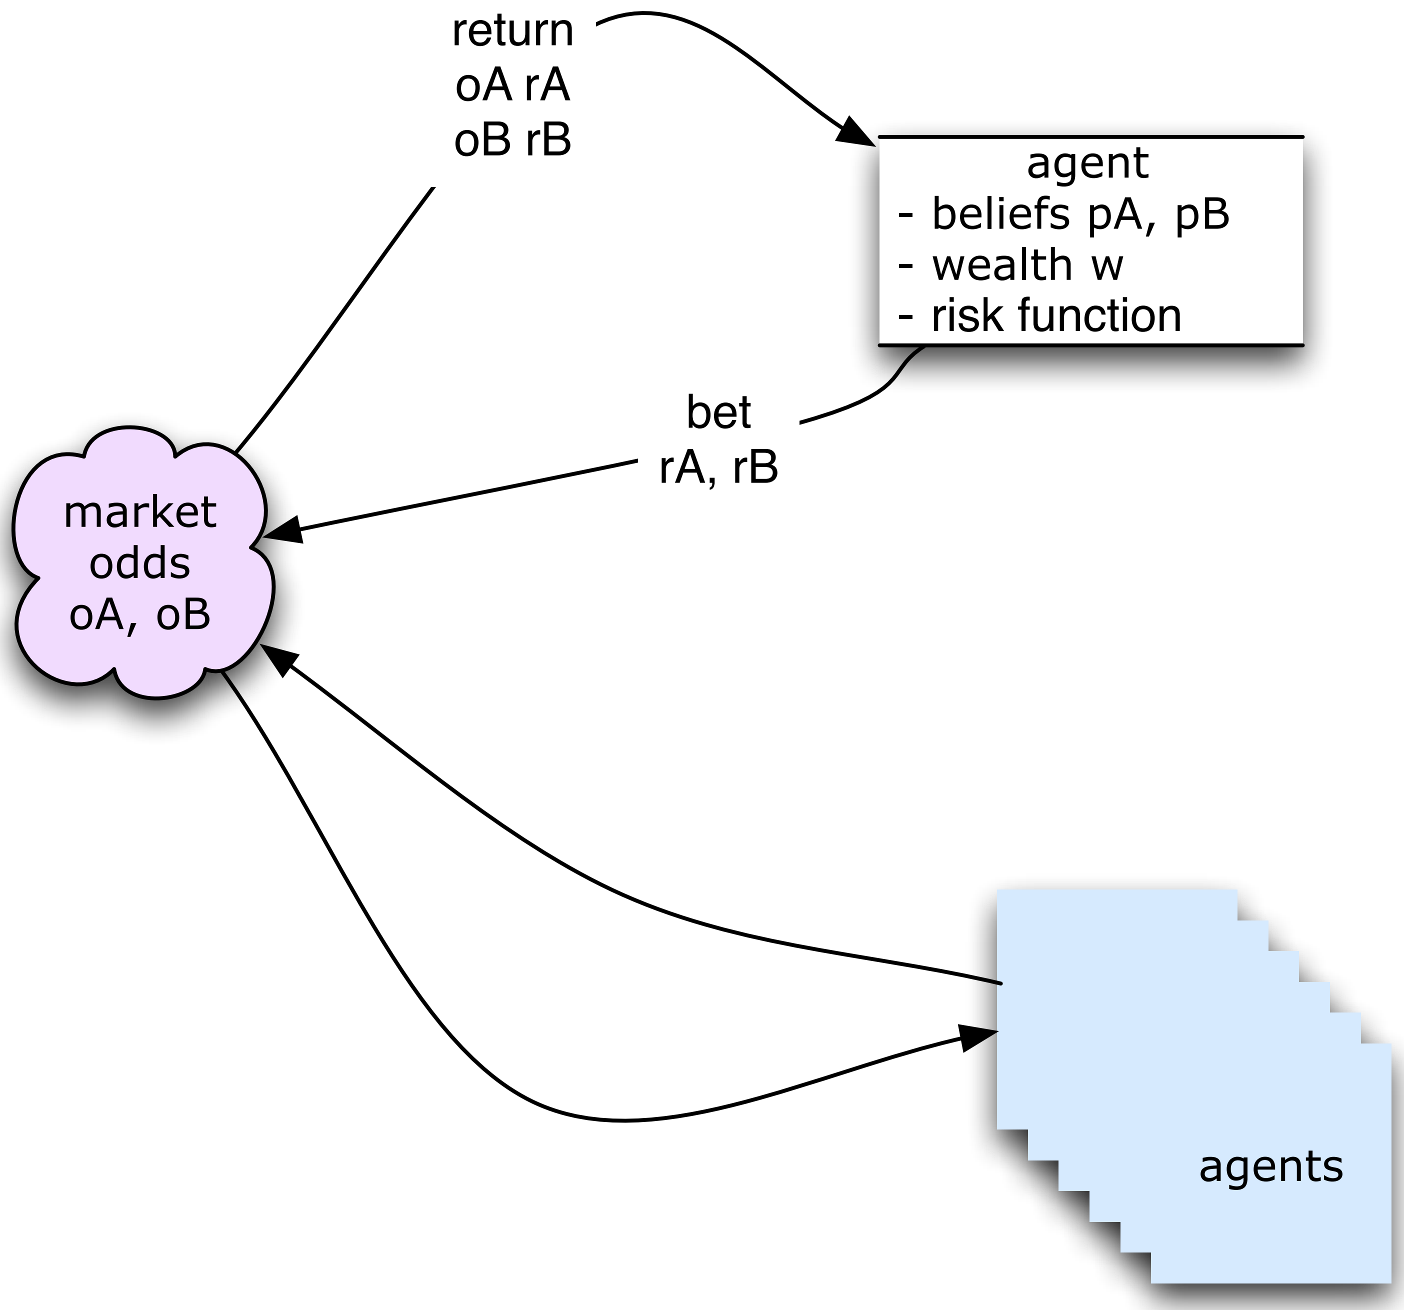
return
oA rA
oB rB
agent
- beliefs pA, pB
- wealth w
- risk function
bet
rA, rB
market
odds
oA, oB
agents

what comes next ....

- how does an agent play?

- how does the market decide the odds?

- what if we repeat the game, what becomes of the wealth distribution?

- then we criticize the model

agent strategy

how does an agent turn a belief into a strategy?

a belief is pA, pB probability on {A,B}

a strategy rA, rB is a function of beliefs, payoffs

reasonable things we can ask of any strategy:
  drA/dpA ≥0
  if pA=1 then rA=1

we introduce a **utility function**
to express how much 1£ is worth, or
how dear 1£ is, to the agent

utility = log

```
why log?
- it is concave: u(x) increases at a decreasing rate
- log (k * x) - log (x) is independent of x
- (often it generalises)
```

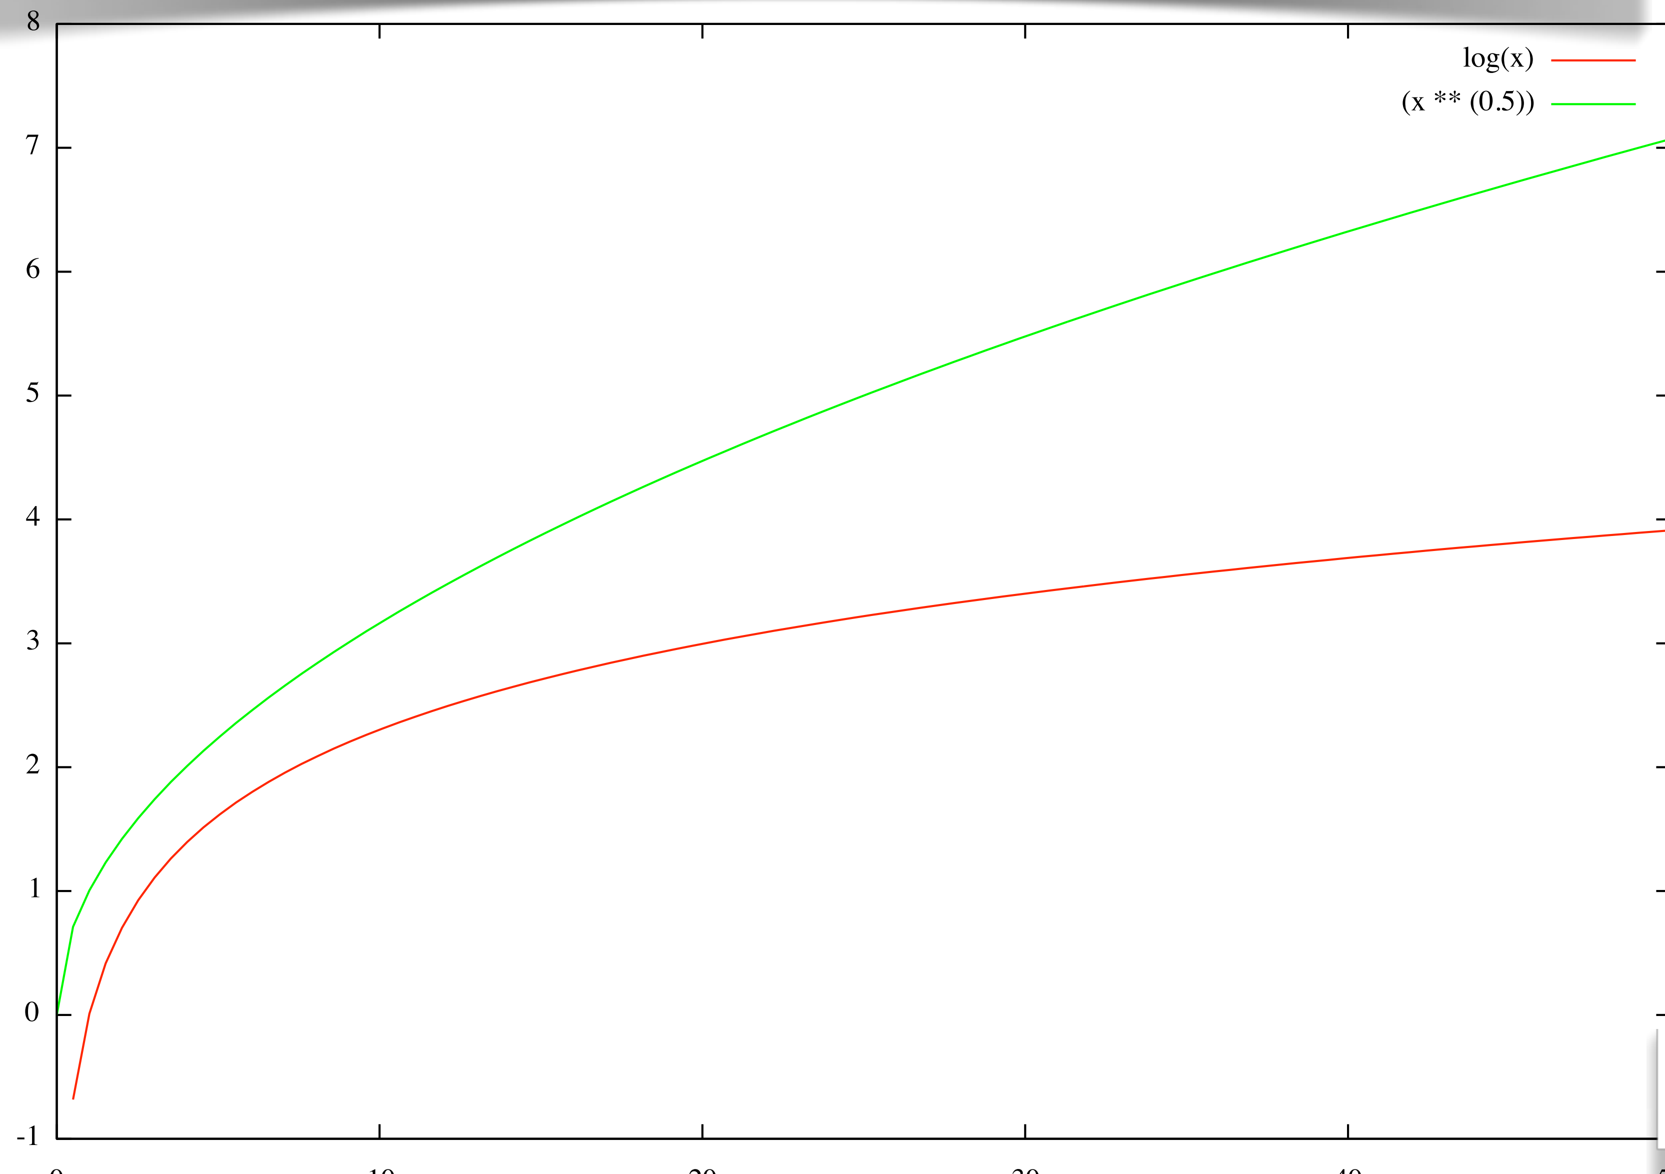

mean (believed) utility

we assume here that agent wants to maximise its mean utility, that is we are looking for:

```
        argmax(rA,rB).(pA * log(rA)+ pB * log(rB))
```

which (as we will see) does not depend on w or oA,oB

```
              payoff = oA rA w if A wins
                       oB rB w if B wins
mean utility
    = mean log(payoff)
    = pA * log(rA oA w)+ pB * log(rB oB w)
    = pA * log(rA)+ pB * log(rB) +
      pA * log(oA)+ pB * log(oB) + log w
```

in the second equation the *italicized terms* are independent of the agent strategy rA, rB; we need to max the first part pA * log(rA)+ pB * log(rB)

NB: this depends on the believed probability pA, pB

risk/utility optimization

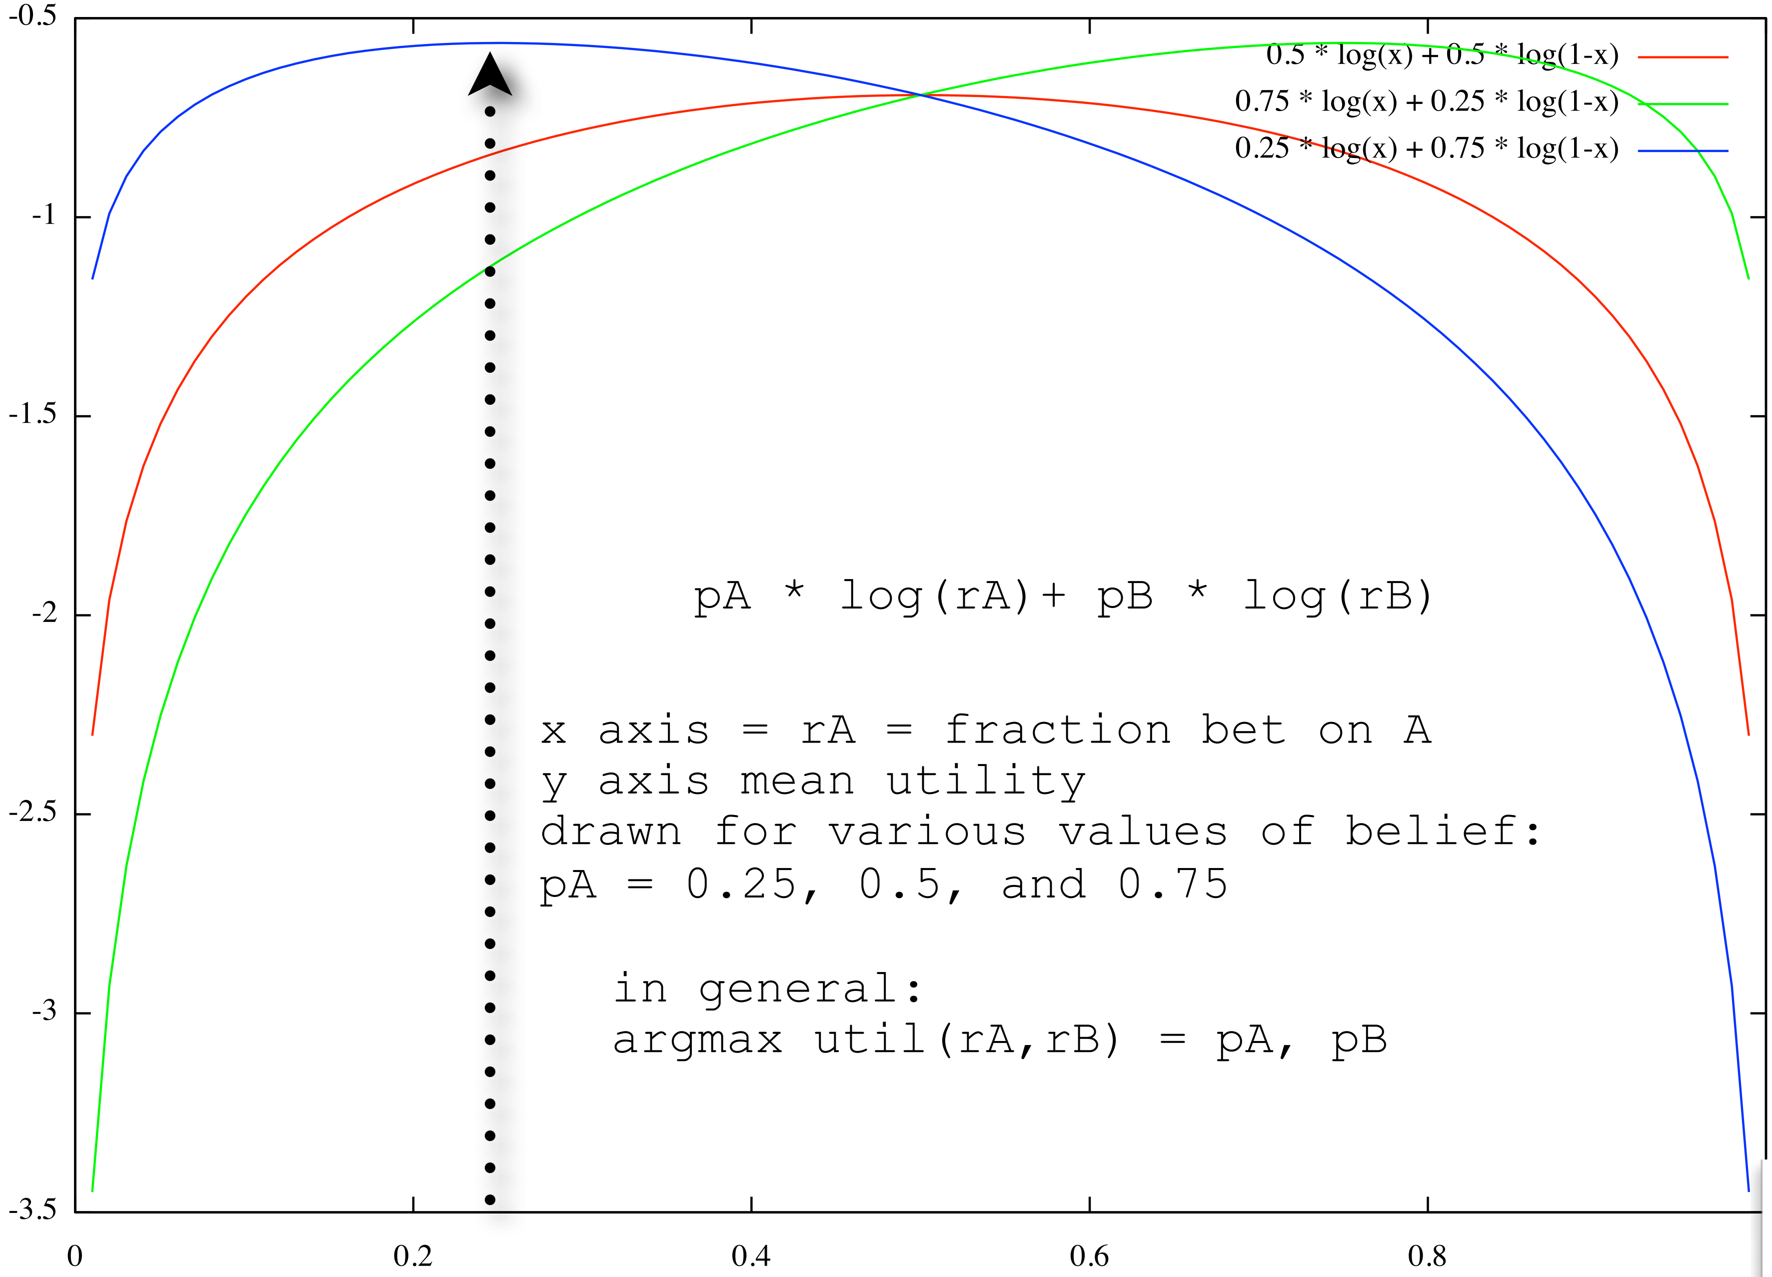

the bettor bets his beliefs

d/d rA (pA * log(rA)+ pB * log(rB)) =
pA/rA - pB/rB

so the optima strategy is:

argmax = pA, pB

and max believed mean utility difference is

pA * log(pA*oA)+ pB * log(pB*oB)

we have subtracted the initial utility log(w)

NB: as expected, we do have
drA/dpA ≥ 0
if pA=1 then rA=1

multi-agent vs the market

we now assume N agents with:

- wealth wn
- beliefs pAn, pBn
- all agents with the same utility function: log

how does the market turn the bets into odds?

market: what are the odds?

```
the market receives the total bet
    w = sum wn
of which on A, B:
   wA = sum rAn * wn
   wB = sum rBn * wn
   wA + wB = w
```

```
total due:
    oA wA = oA sum pAn * wn  if A wins
    oB wB = oB sum pBn * wn  if B wins
subject to (supposing the market is free):

                 oA wA = oB wB = w
```

which we can also write in terms of price-of-1£:

```
1/oA = wA/w = sum rAn * (wn/w)
1/oB = wB/w = sum rBn * (wn/w)
```

a risk-free strategy

$$1/oA = wA/w$$
$$1/oB = wB/w$$

$$1/oA + 1/oB = 1$$

it follows that the strategy

$$rA, rB = 1/oA, 1/oB$$

guarantees a risk-free, 1-to-1, payoff

so the assumption that the agents bet all their wealth w is not a constraint

what are the prices 1/oA, 1/oB?

assuming the optimal strategy pA_n, pBn for agent n:

$$1/oA = \mathrm{sum}_n\ pA_n * w_n/w$$
$$1/oB = \mathrm{sum}_n\ pB_n * w_n/w$$

define the wealth fraction $\mathbf{f_n}$:= w_n/w

$$1/oA = \mathrm{sum}_n\ pA_n * f_n$$
$$1/oB = \mathrm{sum}_n\ pB_n * f_n$$

- everyone shares the same belief pA: 1/oA = pA
- agent n dominates, ie $f_n \sim 1$: $1/oA \sim pA_n$

reconsider:

$$1/oA = sum_n \; pA_n * w_n/w$$

the price is the weighted average of the “market beliefs”, or the “market prediction” about the outcome

caveat ......

this is only true with «loggy» agents; else it also depends on the agents’ utilities/risk functions

wealth dynamics

what if the game is repeated?

Bayesian learning: believing

- X a finite set (say)
- $p \in \mathbf{G}X$ a **hidden** probability on X
- $P = \oplus_n f_n\, p_n \in GGX$ a belief represented as a probability on GX
- s an observation on multisets over X

$P(p_n) = f_n$ - or more rigorously $P(\{p_n\}) = f_n$

NB: a belief is a prob on a prob now!

By multiplication, we have

$$\mu P(A) = \mathrm{sum}_n\, f_n\, p_n(A)$$

a **majority** vote where f_n is the weight accorded to p_n in the prediction

Bayesian learning: learning

- we sample repeatedly from the hidden p, which gives us the observation s above

- we modify the weights in the majority vote of P in order to get closer to the real p:

$$s \cdot f_n / f_n = p_n(s) / \mu P(s) \qquad \textbf{(1)}$$

this defines a new or updated:

$$\mathbf{P} = \oplus_n f_n p_n \Rightarrow \mathbf{s} \cdot \mathbf{P} = \oplus_n (s \cdot f_n)\, p_n$$

NB: the support remains unchanged by the update

P is called the prior, $s \cdot P$ the posterior.

$s \cdot s' \cdot P = (s's) \cdot P$ - ie chunking does not matter

NB:

$$s \cdot f_n / f_n = p_n(s) / \mu P(s)$$
$$s \cdot f_n / s \cdot f_m = p_n(s) / p_m(s) * f_n / f_m$$

in both formulas we are abusing notation

p or µP are not really defined on multisets, but we can promote/extend them using GX → G(multiset(X))

$$p(s) = \mathrm{prod}_{x \text{ in } X}\ p(x)^{s(x)}$$

where s(x) is the number of occurrences of x in s

belief $P = \oplus_n f_n p_n$
outcome s
updating:
$s \cdot f_n / f_n =) p_n(s) / \mu P(s)$ **(1)**

one can rewrite **(1)** - equivalently as **(2)**

$$s \cdot f_n / s \cdot f_m = p_n(s) / p_m(s) * f_n / f_m$$

the invariance under permutation of the observation s, say ABABAB -> AAABBB follows from (2)

$$\mathbf{s \cdot f_n / s \cdot f_m = p_n(s)/p_m(s) * f_n/f_m}$$

since $p_{\mathbf{n}}(s)$ and $p_{\mathbf{m}}(s)$ are invariant under permutation (because we assume that the successive outcomes are independent)

Similarly the invariance under rechunking is easy to see with (2) as

$$\begin{aligned} & s1s2 \cdot f_n / s1s2 \cdot f_m \\ &= pf_n(s1s2)/pf_m(s1s2) * f_n/f_m \\ &= pf_n(s1)/pf_m(s1) * pf_n(s2)/pf_m(s2) * f_n/f_m \end{aligned}$$

Bayesian learning: converging

This defines a Markov chain (MC) on GGX defined as

$$Q(P, s\cdot P) = p(s)$$

that is to say we are 'walking' randomly on GGX, so the kernel Q ∈ [GGX;GGGX] might have a steady state in GGGX - but in fact the interesting limit is a "point-mass" in GGX

assuming $p=p_n$ is the real probability

$$s\cdot P \to p \text{ as } |s| \to \infty$$

as

$$\log(s\cdot f_n/s\cdot f_m) \sim |s| \times KL(p, p_m) \geq 0$$

where KL is the relative entropy of p and p_m (aka the **Kullback-Leibler** divergence)

KL

$$KL(p,q) = \sum_x p(x)\log(p(x)/q(x))$$

$KL(p,q) \geq 0$ and $KL(p,q) = 0$ only if $p = q$.

Because $\log x \leq x - 1$ so $-\sum_i p_i \log(q_i/p_i) \geq -\sum_i p_i(q_i/p_i - 1) = 0$.

Besides $\log x = x - 1$ iff $x = 1$.

Input interpretation:

plot	$-1 + x - \log(x)$	$x = 0$ to 10

Plot:

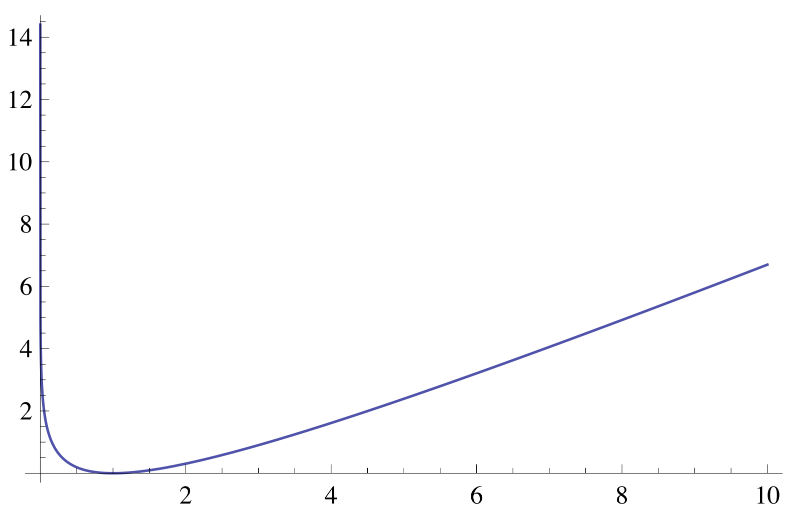

http://www.wolframalpha.com/input/?i=Plot[x-1-Log[x],{x,+0,+10}]

convergence proof

Compare the density updates, we have $s \cdot f_i / s \cdot f_j = f_i p_i(s) / f_j p_j(s)$, so in log form:

$$\log(s \cdot f_i / s \cdot f_j) \quad = \quad \log(f_i / f_j) + \log(p_i(s) / p_j(s))$$

then for $|s| \to +\infty$:

$$\begin{array}{rcll} 1/|s| \log(s \cdot f_i / s \cdot f_j) & \sim & \sum_{x \in X} (s(x)/|s|) \log(p_i(x)/p_j(x)) & \text{by independence of trials} \\ & \sim & \sum_{x \in X} p(x) \log(p_i(x)/p_j(x)) & \text{by SLN} \end{array}$$

where $s(x)$ is the number of x in s.

Supposing $p_i = p$ is the hidden real probability:

$$\log(s \cdot f_i / s \cdot f_j) \quad \sim \quad |s| \times KL(p, p_j) \geq 0$$

Then if $i \neq j$, $KL(p, p_j) > 0$ which implies $s \cdot f_j \to 0$; and hence $\lim s \cdot f_i \to 1$.

So $s \cdot P \to \delta_p$ as $|s| \to \infty$ and we learn eventually the true probability.

```
- justifies the update rule (1), as it does
eventually find the solution

- KL is a natural tool to assess convergence;
there is more to say here ...
```

market payoffs is formally identical to learning!

updated wealth per agent:

$$w_n' = oA\ pA_n\ w_n \quad \text{if A wins}$$
$$w_n' = oB\ pB_n\ w_n \quad \text{if B wins}$$

so the new wealth ratios for agents m and n is

$$f_m'/f_n' = pA_m/pA_n\ f_m/f_n \quad \text{if A wins}$$
$$f_m'/f_n' = pB_m/pB_n\ f_m/f_n \quad \text{if B wins}$$

which exactly as in the Bayesian update formula **(2)**
with $P = \oplus\ f_n\ p_n$ and s = A wins or B wins
which implies that $f_n \to 1$ for the agent that knows the true pA

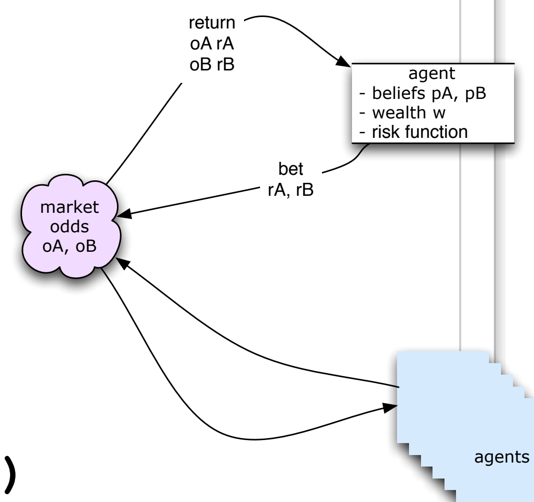

what about the updated price-of-1£?

$$1/oA' = \text{sum}_n\ pA_n * f_n' = \mu P'(A)$$

so $1/oA \to pA$ the true price

more generally, the market is selecting for agents with more accurate beliefs (in the KL sense)

the true p does not need to be in the support of P (ie no player needs to know the true probability)

you can think of the betting market as an interpretation of Bayesian learning as well - let your beliefs bet concurrently ...

reflections on the model

why utility is a log - see above

why maximising mean utility?

why belief is a probability?

how are the odds fixed in advance?
market microstructure - does not matter with "loggy" agents
but in general?

where do beliefs come from? information? do not agents derive their beliefqs also from looking at other agents?

what if the market has a fee?

how does that compare with stock markets?