The Actuarial Profession
making financial sense of the future

GIRO Convention 2009

What is the appropriate framework to describe and understand risk?

The "intelligent agents" paradigm for non-life insurance

Pietro Parodi
Aon Benfield
e: pietro.parodi@aonbenfield.com

Agenda

The Actuarial Profession
making financial sense of the future

Setting the scene

The Actuarial Profession
making financial sense of the future

A bit of risk epistemology

- Understanding risk requires *building an effective model of the environment*
  - At least as difficult as finding the *true* theory of the physical world…
  - … and the world changes constantly…
  - … and so do the rules of the game
- The problem of understanding risk is an "ecological" problem rather than a mathematical or a scientific one
  - Players must survive and thrive in an uncertain environment
  - The environment is a mathematically sophisticated one
  - Plenty of knowledge which can't be either rigorously treated nor ignored

Understanding risk in non-life insurance…

- … involves concretely
  - Making predictions based on data ("learning from data"), e.g. selecting rating factors
  - Dealing with uncertain and soft/expert knowledge, e.g. individual loss estimates
  - Dealing with risk that changes with time, e.g. reserving
  - Making successful decisions in a risky environment, e.g. on pricing
  - Modelling collective behaviour, e.g. to design regulation on capital requirements
- These are typical problems of computational intelligence
  - Computational intelligence attempts to design **intelligent agents** that deal with the problems above

I. Making data-based predictions („learning from data“)

Learning from data – *An overview*

- Many actuarial problems require *learning* the characteristics of a model from a set of data, allowing to make *predictions:*
  - Pricing (frequency/severity model)
  - Selection of rating factors
  - Reserving
  - Capital modelling
- The appropriate framework for prediction is machine learning (*aka* statistical learning)
  - Supervised learning
  - Unsupervised learning

A simple example: rating factors selection

- Objective: predict reinsurance premium Y based on insurer's profile
- Factors: age profile, sex profile, average *direct* premium, etc
- Given: a *dictionary* $\{\psi_\gamma\}_{\gamma\in\Gamma}$ of functions ("features"), select the features that are needed to *predict* the regression function:

$$f_\beta(x_1,x_2...x_n)=\sum_{\gamma\in\Gamma}\beta_\gamma\psi_\gamma(x_1,x_2...x_n)$$

- Feature selection criterion:

Minimise $\mathrm{EPE}(f)=E(L(Y,f(X)))$ on an independent sample

EXPECTED PREDICTION ERROR **LOSS FUNCTION**

- **Example of loss function:** $L(Y,f(X))=(Y-f(X))^2$ **(squared loss)**

The Actuarial Profession
making financial sense of the future

Model selection – Three main issues

1. Prediction accuracy (on an independent sample!)
   - Bias/variance trade-off
2. Interpretation: keep only relevant variables
3. Efficiency
   - Best subset selection is computationally intractable

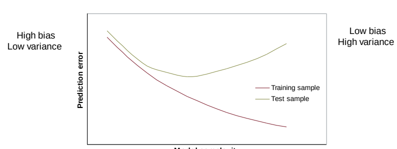

The Actuarial Profession
making financial sense of the future

Model validation protocols

- The objective in all cases is estimating the prediction error
- Ideally one should divide the database randomly into three data sets:
  - Training set (50%) → to fit the model
  - Validation set (25%) → to estimate prediction error for model selection
  - Test set (25%) → to estimate the prediction error of the selected model
- When there is insufficient data, EPE(f) can be calculated approximately:
  - By using K-fold cross-validation
  - By using analytical methods such as AIC, BIC, MDL
  - By using bootstrap (randomised samples with replacement)
- None of these methods can obviously assess the prediction error on new data from a changing/changed risk environment!

The Actuarial Profession
making financial sense of the future

The industry standard for feature selection is GLM

- The model is of the form $Y = g^{-1}(\Sigma\, a_j \psi_j(x_1, x_2, \ldots x_n))$
- Loss function: $L(Y, f(X)) = -2 \log \Pr_{f(X)}(Y)$
- Main ingredients:
  - An error structure (exponential family)
  - A link function (g)
  - A dictionary of functions $\{\psi_j\}$ (often implicit)
- Model selection and validation ("standard" approach):
  - Greedy approach, e.g. forward/backward stepwise selection
  - Include/exclude features based on *t*-test, F-statistic, AIC, BIC, MDL…

$$AIC = -\frac{2}{N} \cdot \text{loglik} + 2 \cdot \frac{d}{N}$$

N = no of points,
d = no of parameters,
loglik = log-likelihood @ max

The Actuarial Profession
making financial sense of the future

GLM – Results on our example

- A multivariate Gaussian model is sufficient in this case
- Forward selection yields Y ~ AvgDirPremium + %Young as the winning model

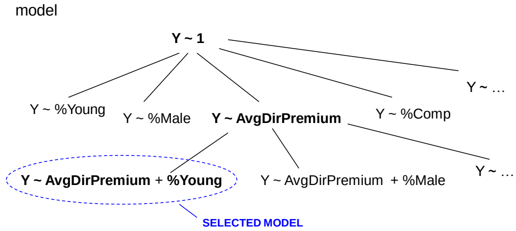

The Actuarial Profession
making financial sense of the future

An alternative approach: regularised regression

- Main idea: to minimise $\mathrm{EPE}(f) = \|Y - f(X)\|_{l_2}^2$ on an independent set, minimise a *regularised* functional:

$$\mathrm{EPE}(f) = \|Y - f(X)\|_{l_2}^2 + \lambda g_\beta(X)$$

  on the training set!

- Most famous example: ridge regression $\mathrm{EPE}(f) = \|Y - f(X)\|_{l_2}^2 + \lambda \|\beta\|_{l_2}^2$
- Model validation is provided by, e.g., *k-cross-validation*
- Penalty terms can be interpreted in a Bayesian framework

The lasso (Tibshirani, 1996)

- l_1-penalty on the size of regression coefficients

$$\mathrm{E}_n^\lambda(\beta) = \|Y - f_\beta(X)\|_{l_2}^2 + \lambda \|\beta\|_{l_1}$$
$$\|\beta\|_{l_1} = |\beta_1| + |\beta_2| + \ldots + |\beta_n|$$

- Performs automatic variable selection!
- Breaks intractability of subset selection
- Efficient path algorithms are available
- Can be over-zealous in eliminating correlated features
- Corresponds to a Laplace distribution prior
- `http://videolectures.net/kdd08_hastie_rpcd/`

The Actuarial Profession
making financial sense of the future

Interpretation of the lasso

- How does the lasso achieve variable selection?
- Compare lasso and ridge regularisation

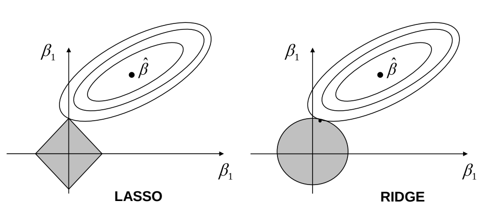

The Actuarial Profession
making financial sense of the future

Lasso – Results on our example

- Results obtained with the R package "LARS" by Hastie (2007)

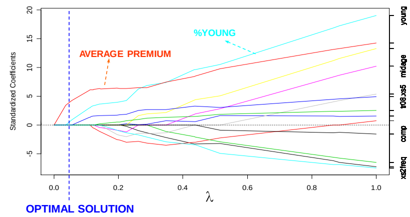

$$E_n^{\lambda}(\beta) = \|Y - f_\beta(X)\|_{l_2}^2 + \lambda\|\beta\|_{l_1}$$

Lasso – Model validation

- Optimal solution is for the regularisation parameter ~ 0.05
- Corresponds to Y ~ AvgDirPremium + %Young, as for GLM

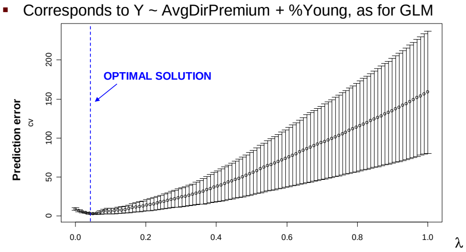

Other types of regularisation

- Elastic net (Zou & Hastie, 2005)

$$E_n^{\lambda}(\beta) = \|Y - f_\beta(X)\|_{l_2}^2 + \lambda\|\beta\|_{l_1} + \mu\|\beta\|_{l_2}^2$$

- Enforces sparsity while avoiding the excesses of lasso
- Can address situations where

  # of parameters » # of observations !!!

  - E.g. microarray data analysis, with groups of correlated genes

What about neural networks?

- Nothing but non-linear statistical models

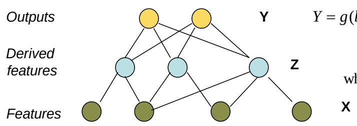

$$Y = g(b_0 + b^T Z) \overset{\text{REG}}{=} b_0 + b^T Z$$

$$Z = \sigma(a_0 + a^T X)$$
$$\text{where } \sigma(v) = (1 + \exp(-v))^{-1}$$

- Can approximate any function
- No need for detailed specification of the model
- Provide “prediction without interpretation” (Hastie et al., 2001)

Comparison of GLM and regularisation

- GLM
  - Limited by linearity (but a large dictionary of functions is possible)
  - “log P” loss function more general than squared loss
  - Greedy algorithms may get stuck in local minima
- Regularised regression
  - Breaks intractability and can be extremely efficient
  - Can address cases where there # variables » # data points
  - Use of quadratic loss function is a limit – or is it?
- Hybrid approaches
  - Regularised GLM

II. Dealing with uncertain and soft knowledge

Overview

- A significant portion of the things we know about risk is
  - Uncertain (model, parameter, data uncertainty)
  - Soft or qualitative
  - Fuzzy
  - Anecdotic
- Techniques to deal with uncertain/soft knowledge
  - Rule-based systems, e.g. expert systems
  - Fuzzy set theory
  - Bayesian analysis
  - Dempster-Shafer belief/possibility theory
  - Non-monotonic reasoning

An example: severity distribution with data uncertainty and prior knowledge

- The problem: find the parameters of the loss severity distribution
- A very simple example:
  - Single-parameter Pareto distribution (large losses)
  - Data uncertainty depends on amount already paid, size of loss, date of loss…
  - Underwriting guidelines: α between 2 and 5, α = 3.5 default recommendation
- Crisp data, no prior knowledge
  - Use MLE for point estimates and Fisher information matrix or bootstrap for standard error

Fuzzy set theory (Zadeh, 1965)

- Captures the notion of an object whose value is not sharply defined: e.g. “large loss”, “risky policyholder”
- Membership $\mu_A(x)$ to a set can be any real number between 0 and 1
- Fuzzy numbers: fuzzy subsets of $\Re$
- Fuzzy arithmetic can be defined quite naturally, e.g. A (+) B:

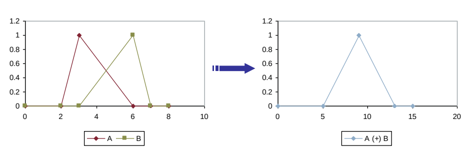

x

Example – Using fuzzy set theory (I)

- Non-settled loss:
  triangular fuzzy number with width larger for losses that have a large outstanding percentage and are recent

- Use fuzzy arithmetic to produce an MLE-like estimate of the parameters, bootstrap for standard errors. E.g., for a Pareto distribution:

$$\alpha = \frac{n}{\sum \ln x_i - n \ln \theta}$$

- *Requires: crisp functions of fuzzy numbers, sum of fuzzy numbers, adding crisp and fuzzy numbers…*

Example – Using fuzzy set theory (II)

- The result is a fuzzy number with membership function $\mu_\alpha(x)$
- Prior knowledge on α: another fuzzy number $\mu_{\alpha'}(x)$
- Final result: the fuzzy intersection $\mu_{A \cap A'}(x)$ of the two estimates

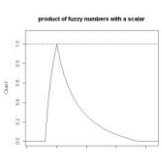

$\mu_\alpha(x)$

(data-driven)

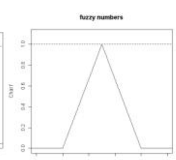

$\mu_{A'}(x)$

(underwriting guidelines)

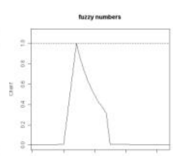

$\mu_{A \cap A'}(x)$

(combined)

Example – Using rule-based systems

- Rule-based approach – Example 1
  - Exclude from the data set all losses whose uncertainty is greater than 30%
  - Calculate α with MLE based on the remaining data points
  - If α (MLE) is between 2 and 5, keep it
    - Else if it is < 2 choose α=2
      - Else choose α=5

- Rule-based approach – Example 2
  - Use parametric bootstrap to get α and se(α)
  - Use a credibility approach to combine the above with underwriter's opinion: e.g.,
    - α(cred) = Z α(bootstrap) + (1 – Z) α(underwriter)
    - Z = se(α)^2 / (Var(α)+ se(α)^2), where Var(α) is the variance of the underwriter's estimate
    - (Subject to α being between 2 and 5)

Example – Using a Bayesian approach

- Non-settled losses: $P(\hat{X} \mid X) \sim \text{Gamma}(E(\hat{X}) = X, CV = a)$
- Prior knowledge: $\alpha \sim \text{Beta}(\min = 2, \max = 5, \text{mode} = 3.5, \ldots)$

- **No data uncertainty, no prior knowledge**: maximise likelihood

$$P(X_1, \ldots X_n \mid \alpha)$$

- **Prior knowledge, no data uncertainty**: maximise *posterior* likelihood

$$P(\alpha \mid X_1, \ldots X_n) \propto P(X_1, \ldots X_n \mid \alpha) P(\alpha)$$

- **Prior knowledge, data uncertainty**: maximise posterior with hidden variables

$$P(\alpha \mid \hat{X}_1, \ldots \hat{X}_n) \propto \Pr(\alpha) \int P(\hat{X}_1, \ldots \hat{X}_n \mid X_1 = x_1, \ldots X_n = x_n) P(X_1 = x_1, \ldots X_n = x_n \mid \alpha)\, dx_1 \ldots dx_n$$

(solve by numerical methods, e.g. Markov Chain Monte Carlo)

- Simplifications possible by including conditional independence constraints, e.g. by using Bayesian networks: crucial with many variables

The Actuarial Profession
making financial sense of the future

Example – Using a Bayesian network

- Bayesian networks are compact representations of the joint probability distribution through directed acyclic graphs.

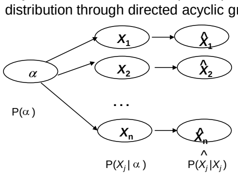

$$\Rightarrow P(\alpha, \hat{x}_1, \ldots \hat{x}_n, x_1, \ldots x_n) = P(\alpha) P(x_1 \mid \alpha) \ldots P(x_n \mid \alpha) P(\hat{x}_1 \mid x_1) \ldots P(\hat{x}_n \mid x_n)$$

- The posterior probability $P(\alpha \mid \hat{X}_1, \ldots \hat{X}_n)$ can be calculated by using inference by enumeration (see Appendix):

$$P(\alpha \mid \hat{x}_1, \ldots \hat{x}_n) = c \Pr(\alpha, \hat{x}_1, \ldots \hat{x}_n) = c \int \Pr(\alpha, \hat{x}_1, \ldots \hat{x}_n, x_1, \ldots x_n)\, dx_1 \ldots dx_n = c \Pr(\alpha) \prod_{k=1}^{n} \int \Pr(\hat{x}_k \mid x) \Pr(x \mid \alpha) dx$$

- Information hidden (for simplicity): the distribution below θ

The Actuarial Profession
making financial sense of the future

Fuzzy set theory *v* Bayesian approach

- Common criticism: "FST just an unwieldy version of probability theory". Is that fair? Both deal with uncertainty, but…
  - Conceptual difference: *a* loss may be exactly £130,000 – but whether this loss should be called "large" is vague
  - However, the uncertainties we care about are *quantitative* and not *linguistic/logical*
  - Effective toolbox available for those who embrace Bayes: MCMC, Gibbs sampling…
  - … whereas anything beyond basic arithmetic is tricky with fuzzy set theory

- FST poor at addressing *parameter* and *model* uncertainty
  - Fuzzy numbers always behave as *perfectly correlated variables*
  - Unlike FST, Bayesian methods allow to address model uncertainty

- *FST: a much-needed, rigorous extension of set theory, but… is it for us?*

The Actuarial Profession
making financial sense of the future

III. Dealing with a changing environment

The temporal element

- Dynamic vs static environments:
  - In most actuarial problems, the environment is not static – e.g., the frequency of losses may be changing due to improved risk control mechanisms…
  - … and experience reveals itself gradually
- One needs knowledge-update mechanisms
- Why does one need a model which evolves in time?
  - Agents have limited time/space resources
  - Imagine a situation where very large collections of data are updated…
- Techniques
  - Kalman filtering
  - Hidden Markov models
  - Dynamic Bayesian networks

Example: Reserving

- Source: Claims reserving, state-space models and the Kalman filter by P. De Jong and B. Zenhnwirth (1983)

Development year

	0	1	2	3	4	5	6	7	8	9
2000	19,272	56,333	84,499	111,183	131,937	131,937	137,867	131,937	131,937	131,937
2001	30,269	78,700	100,494	117,445	134,396	125,921	127,131	129,553	129,707	
2002	32,447	72,268	110,615	109,140	115,039	125,363	126,838	127,247		
2003	59,030	134,293	166,759	190,371	178,565	193,322	194,371			
2004	81,347	131,298	156,986	165,549	165,549	165,549				
2005	63,433	117,805	131,398	132,531	138,000					
2006	47,545	86,095	95,090	98,945						
2007	42,141	67,836	81,198							
2008	44,613	85,880								
2009	34,822									

Z_{t+1} (NEW INFORMATION)

- Ingredients of a probabilistic temporal model
  - Hidden variables $\mathbf{x}_t$ → true parameters
  - Evidence variables $\mathbf{z}_t$ → latest diagonal
  - Transition model $\mathbf{P}(\mathbf{x}_{t+1}|\mathbf{x}_t)$ → e.g. random walk model, x(t+1) = x(t) + v(t+1)
  - Observation model $\mathbf{P}(\mathbf{z}_t \mid \mathbf{x}_t)$ → e.g. $z_d(t) = x(t-d)\phi(d)+u_d(t)$, $\phi(d) = (d+1)\exp(-d)$
  - Prior probability $\mathbf{P}(\mathbf{x}_0)$ → based on triangle available at time 0, e.g. 2003

A possible approach: Kalman filtering

- Kalman filtering: regression analysis with a mechanism for updating parameters. Classical application: *radar tracking*
- Key assumption: the current state follows multivariate Gaussian

UPDATED PARAMETER ESTIMATE · PARAMETER ESTIMATE @ t · INFO UPDATE

$$\mu_{t+1} = F\mu_t + K_{t+1}(z_{t+1} - HF\mu_t)$$

TRANSITION MODEL · KALMAN GAIN (CREDIBILITY) · OBSERVATION MODEL

- Complete set of equations in the Appendix

A more general approach: Dynamic Bayesian Networks

- DBN's: Bayesian networks representing temporal models
- Assumptions
  - Stationarity (laws governing change don't change!)
  - (First order) Markov process (current state depends only on previous one)
- Complete joint distribution:

PRIOR PROBABILITY · TRANSITION MODEL · OBSERVATION MODEL

$$P(X_0, X_1, \dots X_t, E_1, \dots E_t) = P(X_0) \prod_j P(X_j \mid X_{j-1}) P(E_j \mid X_j)$$

- This is all we need to solve the prediction problem:
  - $P(X_{t+1} \mid e_{1:t+1}) = \alpha P(e_{t+1} \mid X_{t+1}) \sum_{xt} P(X_{t+1} \mid x_t) P(x_t \mid e_{1:t})$
- Representation as a (unrolled/rolled-up) dynamic Bayesian networks

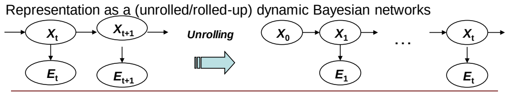

Comparison of temporal models

- All Kalman filters can be represented as a DBN (but not vice versa)
  - Multivariate Gaussian hypothesis, linearity for Kalman are critical
  - Serious non-linearities (e.g. changes of reserving guidelines, judicial decisions…) require DBN's with both discrete and continuous variables
  - Extended Kalman filters attempt to deal with non-linearities
- Markov Chain Monte Carlo methods can be used for approximate inference in DBNs
- Hidden Markov Models (HMM) and DBN are equivalent formulations – however, DBN's are more compact and allow gains from sparsity
- <u>The biggest limitation for all temporal models is stationarity: in all cases, a prior model of the possible future changes is needed</u>

IV. Making decisions in an uncertain environment

Overview

- Ultimately, we want to understand risk because we have to make informed decisions
- Examples:
  - Buying an insurance policy
  - Choosing an investment
  - Making a business plan
  - Buying reinsurance
  - Dynamic financial analysis
- In all cases, what is the likely outcome of the decisions we make?
- This is a well-known problem in computational intelligence:
  - **Designing an intelligent agent which can move in an environment making the best decisions – i.e., the decisions which maximise utility**
- Main recommended reading: Russell and Norvig, *Artificial Intelligence: A modern approach*, Prentice Hall, 2003

Model players as *intelligent agents*

- Each agent (e.g. an insurance company) is autonomous and incorporates strategies to interact with the environment
- Ignore the other players in the market –all blurred into the "environment"
- Characteristics of the environment
  - Fully v partially observable
  - Deterministic *v* stochastic
  - Stationary *v* non-stationary
  - Discrete v continuous
- Intelligent agents in a partially observable, stochastic environment can be modelled as Dynamic Decision Networks (DDN's)

Example: insurers as dynamic decision networks (business planning, DFA)

- DDN's are DBN's extended with *decision nodes* and *utility nodes*

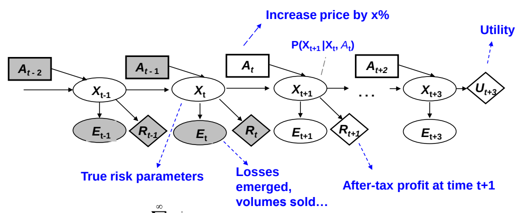

- Utility: $U(\{s_j\}_{j\in\mathbb{N}}) = \sum_{j=0}^{\infty} \gamma^j R(s_j)$

IV.Modelling collective behaviour

Overview

- The environment is not really a "blur"… and the other players cannot be ignored
- Two main problems:
  - Agent design: *e.g.*, maximise utility in the face of competition
  - Mechanism design: *e.g.* you're the regulator
- Two main ingredients
  - Multi-agent systems
  - Game theory (with tournaments, *à la* Axelrod, but without one-to-one encounters)

Example – Personal insurance market

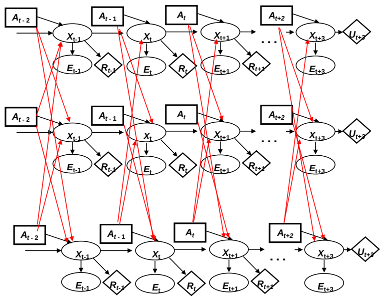

Example (cont'd)

- Formally, the market can be modelled as a network of DDN's (except for individual customers)
- The effectiveness of different rules can be tested via stochastic simulation, e.g. through an Axelrod-like tournament (game theory)
- The exercise is severely limited by the patchy and fuzzy knowledge that each player has of the other players
- The regulator might be in a better position than individual players to run such an exercise
  - Unlikely to provide exhaustive answers on mechanism design but might lead to discover, e.g., unforeseen side-effects of regulation

The emergence of truly collective behaviour

- Running these simulations might allow to deepen our understanding of certain genuinely collective behaviours, e.g. cycles, bubbles…
- E.g., at what level of complexity the typical features of the insurance market (including the undewriting cycle) start to emerge? Can its length be predicted?

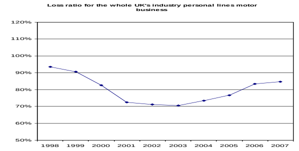

Conclusions

So what is the appropriate framework?

- The "intelligent agents" paradigm provides an adequate framework for describing and understanding risks
  - Risk agents need to learn from data… (machine learning)
  - … deal with uncertain/soft knowledge… (Bayesian networks)
  - … deal with changes in the environment (dynamic Bayesian networks)
  - … make decisions in that environment and modify it… (dynamic decision networks)
  - … and interact/compete with other risk agents for resources (multi-agent systems, game theory)
- Computational intelligence is now more than a collection of heuristics, thanks among the others to the "Stanford school" of statisticians (Efron, Tibshirani, Hastie, Zou…)

Practical findings

- Regularisation *an efficient alternative to GLM* for predictive modelling
- Bayesian networks better than fuzzy set theory for dealing with uncertain and expert knowledge
- Dynamic bayesian networks (DBN) are a more general method than Kalman filtering to capture the changing nature of risk
- Dynamic decision networks – an extension of DBN's – a good model for agents making decisions in a risky environment
- The main ingredients to understand the collective behaviour of markets are multi-agent systems and game theory (stochastic tournaments)
- **IT'S A BAYESIAN JUNGLE OUT THERE!**

Limitations

- Current computational intelligence techniques capture the "ecological" aspect of risk only up to a point:
  - "Prediction" always means prediction in a somehow stationary environment
  - The laws themselves change here and people "work the system"… only Asimov-style artificial intelligence could address this!
  - Soft knowledge on non-stationarity can be introduced in a Bayesian fashion, but…
- Complexity of some of the techniques (e.g. multi-agent systems)
- A parochial view of what "risk" means?
  - E.g. where do methods such as derivatives fit in all this?

References

- Many sources… but two must-haves:
  - Hastie, Tibshirani and Friedman, "The elements of Statistical Learning: Data Mining, Inference and Prediction", Springer, 2001
    - The book that has given a solid statistical foundation to machine learning, by those who invented the bootstrap, the lasso, and much else
  - Russel and Norvig, "Artificial Intelligence: A Modern Approach", 2nd Ed, Prentice Hall, 2003
    - The main reference for AI, also known as **"The intelligent agents book"**: responsible for changing the way we look at the discipline

Questions?

Appendix I. Fuzzy set theory

Fuzzy set theory (Zadeh, 1965)

- Fuzzy membership: a fuzzy set A in Ω is a set of ordered pairs
  $A = \{x, \mu_A(x)\}$, x in Ω, $\mu_A : X \to [0,1]$ (degree of membership)
- Fuzzy set operations can be defined naturally as:

$$\mu_{A \cup B}(x) = \max\{\mu_A(x), \mu_B(x)\}, \quad \mu_{A \cap B}(x) = \min\{\mu_A(x), \mu_B(x)\}$$

- Fuzzy number: (informal definition) a fuzzy subset of R whose membership function is centred around a given real number. *It's a fancy range!*

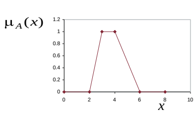

or

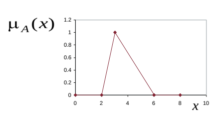

Fuzzy arithmetic

- Fuzzy arithmetic is based on Zadeh's extension principle: if * is a binary operation, and A, B are two fuzzy numbers,

$$\mu_{A(*)B}(z) = \sup_{x,y}\{\min(\mu_A(x), \mu_B(y)) \mid z = x * y\}$$

- Crisp functions can be defined similarly:

$$\mu_{f(A)}(z) = \sup_{x \in \Re}\{\mu_A(x) \mid z = f(x)\}$$

- Quick reference: `http://videolectures.net/acai05_berthold_fl/`
- A brand-new R package for fuzzy arithmetic: fuzzyOP (March 2009)

Appendix II. Bayesian networks

Bayesian networks

- Formally, a Bayesian network is a directed acyclic graph (DAG) where
  - Each node represents a random variable
  - There is an arc from X to Y if X affects directly Y ("X is a parent of Y")
  - Each node has a conditional probability distribution Pr(X | Parents(Y))
- The topology + conditional probability tables of Bayesian networks are a *compact* representation of the joint probability distribution $\Pr(E_1,\ldots E_n)$
  - The compactness derives from the sparsity of the connections
  - Alternatively, they can be viewed as a collection of independence statements (a node is independent of its non-descendants, given its parents)
- The chain rule can be written more compactly as the formula below, which defines the full joint distribution as the product of the local conditional distributions:

$$\Pr(E_1,\ldots E_n) = \prod_{i=1}^{n} \Pr(E_i)\Pr(E_i \mid \text{Parents}(E_i))$$

The Actuarial Profession
making financial sense of the future

Constructing Bayesian networks

- 1. Choose an ordering of variables $X_1, \ldots, X_n$
- 2. For i = 1 to n
  - add X_i to the network
  - select parents from $X_1, \ldots, X_{i-1}$ such that

$$\mathbf{P}(X_i \mid Parents(X_i)) = \mathbf{P}(X_i \mid X_1, \ldots X_{i-1})$$

This choice of parents guarantees:

$\mathbf{Pr}(X_1, \ldots, X_n) = \pi_{i=1}^{n} \mathbf{Pr}(X_i \mid X_1, \ldots, X_{i-1})$ (chain rule)

$= \pi_{i=1} \mathbf{Pr}(X_i \mid Parents(X_i))$ (by construction)

Source: http://aima.eecs.berkeley.edu/slides-ppt/m14-bayesian.ppt (Resources for Russell & Norvig's book)

The Actuarial Profession
making financial sense of the future

Inference in Bayesian networks

- Bayesian networks can be used to calculate the *posterior* distribution of the parameters/random variates we are interested in (the *query variables*)
- A typical query requires to calculate P(X|e) where X is the query variable and e is an instance of the evidence variable E. There are also hidden variables Y with values y.
- Exact inference by enumeration:
  - $P(X \mid e) = \alpha\, P(X,e) = \alpha \sum_y P(X,e,y)$
- Other inference methods are
  - By variable elimination (exact)
  - Direct sampling (approximate)
  - Markov Chain Monte Carlo methods (approximate)

The Actuarial Profession
making financial sense of the future

Appendix III. Markov decision processes and DDN's

The Actuarial Profession
making financial sense of the future

Fully observable environments

- **Source: Russel & Norvig, 2003**
- **In the case of fully observable environment, an easy, complete solution to optimal decision making by an agent is provided by Markov decision processes (MDPs)**
- **Markov decision processes**
  - **Assumptions:** fully observable environment, stochastic, stationary
  - **Markovian transition model**: At each time *t*, an agent will be in state *s* and will be able to perform an action *a*. As a consequence, it will move to state *s'* with probability *T*(*s,a,s'*)
  - **Utility function**: In each state *s*, the agent receives a reward *R*(*s*). The utility of a state sequence can either be additive or additive-discounted (no other possibilities!!):

$$U(\{s_j\}_{j\in\aleph}) = \sum_{j=0}^{\infty} R(s_j) \qquad U(\{s_j\}_{j\in\aleph}) = \sum_{j=0}^{\infty} \gamma^j R(s_j)$$

  - A **rule** ρ specifies what each agent should do in any state that it might reach
  - An optimal rule is one which maximises expected utility
  - The solution can be found with the so-called value iteration algorithm, which is guaranteed to converge to a unique solution (see Russell & Norvig, Section 17.2)

The Actuarial Profession
making financial sense of the future

Partially observable environments

- **Partially observable MDPs (pom-dee-pees)**
  - The agent does not necessarily know which state it is in
  - The utility depends on *s* **and** on how much the agent knows about *s*
  - A **belief state** *b* is defined as the probability distribution over all possible states
  - It can be shown that the optimal action depends on the agent's current belief state *b*
  - The problem of solving a POMDP on a physical state space can be reduced to that of solving an MDP on the corresponding belief state space
  - A comprehensive approach to POMDPs is provided by **dynamic decision networks**

Dynamic Decision Networks (DDN's)

- **DDN's provide a comprehensive approach to agent design for partially observable stochastic environments**
- **DDN's are DBN's extended with *decision nodes* and *utility nodes***
- **In the network below (which looks ahead three steps), X_t are the state variables, E_t are the evidence variables, A_t is the action at time t, R_t is the reward @ t and U_t is the utility of the state @ t**
- **Note that: the transition model is now $P(X_{t+1} | X_t, A_t)$, the observation model is as before and U_t is assumed to be available only in approximate form!**

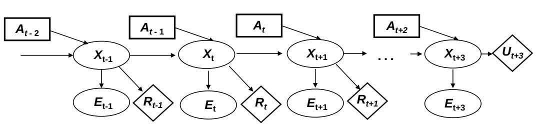

Appendix IV. Kalman filter

Kalman filtering: model description

- Kalman equations “*in their full, hairy horribleness*” (Russell and Norvig, 2003):

  Transition model: $P(x_{t+1} \mid x_t) = N(Fx_t, \Sigma_x)(x_{t+1})$

  Sensor model: $P(z_t \mid x_t) = N(Hx_t, \Sigma_z)(z_t)$

  F, Σ_x : linear transformation, noise covariance for transition model

  H, Σ_z : linear transformation, noise covariance for sensor model

  Update equation for the mean: $\mu_{t+1} = F\mu_t + K_{t+1}(z_{t+1} - HF\mu_t)$

  Update equation for the covariance: $\Sigma_{t+1} = (1 - K_{t+1})(F\,\Sigma_t F^t + \Sigma_t)$

  Kalman gain: $K_{t+1} = (F\,\Sigma_t F^T + \Sigma_x)H^T(H(F\,\Sigma_t F^T + \Sigma_x)H^T + \Sigma_x)^{-1}$

- Note that the Kalman gain gives the credibility of the new observation