Analysis and Synthesis: Multi-Agent Systems in the Social Sciences

at: Market-Based Control of Complex Computational Systems

University of Liverpool

1 September 2008

Robert E. Marks
School of Economics/AGSM/ALL
UNSW

bobm@agsm.edu.au

OUTLINE

1. Introduction
2. Sufficiency and Necessity
3. Analysis in the Social Sciences
   — Simulation and Analysis
   — Validation
4. Economists’ Uses of Multi-Agent Systems
5. Synthesis in the Social Sciences
   — Designing Markets
6. Similarities and Differences

1. Introduction

I assert that computer scientists are concerned with finding solutions to the issues of market design, whereas —

social scientists in general and economists in particular have been concerned with explaining and predicting social phenomena.

These two approaches have both demanded sufficiency, but social scientists (or at any rate economists) have also demanded necessity:

Not just: "This is a solution"

but also: "This set contains all possible solutions."

2. Sufficiency and Necessity

Simulation demonstrates existence, sufficiency,

but not necessity.

Simulation can demonstrate the untruth of a proposition,

but not provide proofs or theorems,

simulations cannot provide generality.

What, never?

Does this matter?

Formal Simulation

Mathematical "model A" comprises the conjunction $(a_1 \wedge a_2 \wedge a_3 \cdots \wedge a_n)$, where $\wedge$ means "AND", and the a_i denote the elements (equations, parameters, initial conditions, etc) that constitute the model.

Sufficiency: If model A exhibits the desired target behaviour B, then model A is sufficient to obtain exhibited behaviour B: $A \Rightarrow B$

Thus, any model that exhibits the desired behaviour is sufficient, and demonstrates one conjunction of conditions (or model) under which the behaviour can be simulated.

But if there are several such models, how can we choose among them? And what is the set of all such conjunctions (models)?

Necessity

Necessity: Only those models A belonging to the set of necessary models $\mathcal{N}$ exhibit target behaviour B.

That is, $(A \in \mathcal{N}) \Rightarrow B$, and $(D \notin \mathcal{N}) \not\Rightarrow B$.

A difficult challenge: determine the set of necessary models, $\mathcal{N}$.

Since each model $A = (a_1 \wedge a_2 \wedge a_3 \cdots \wedge a_n)$, searching for the set $\mathcal{N}$ of necessary models means searching in a high-dimensional space, with no guarantee of continuity, and a possible large number of non-linear interactions among elements.

Lack of Necessity Means ...

For instance, if $D \not\Rightarrow B$, it does not mean that all elements a_i of model D are invalid or wrong, only their conjunction, that is, model D.

It might be only a single element that precludes model D exhibiting behaviour B.

But determining whether this is so and which is the offending element is a costly exercise, in general, for the simulator.

Without clear knowledge of the boundaries of the set of necessary models, it is difficult to generalise from simulations.

Simulation Can Demonstrate Necessity . . .

Only when the set $\mathcal{N}$ of necessary models is known to be small (such as in the case of DNA structure by the time Watson & Crick were searching for it) is it relatively easy to use simulation to derive necessity.

They had much information about the properties of DNA (from others):

when they hit on the simulation we know as the “double helix”, they knew it was right.

But still “*A* structure ...”, not “*The* structure” in the title of their 1953 *Nature* paper.

3. Analysis in the Social Sciences

In the social sciences:

- "positive" analysis
- explanation of existing phenomena, understanding
- prediction.

In engineering, crudely:

- "normative" analysis
- solving problems
- synthesis
- design.

3.1 Simulation and Analysis

The anecdote about the economist looking for his lost car keys:

"An accurate answer to the wrong question"? (using closed-form methods)

or: simulation (numerical methods)

"Approximate answers to the right questions"

Helped by the developments in computer hardware and software.

Meanwhile: C.S. has borrowed simulation tools from the natural world:
artificial neural nets, simulated annealing, genetic algorithms/programming

Want: dynamics, out-of-equilibrium characterisations.

3.2 Validation

For whom?

With regard to what?

A good simulation is one that achieves its goals:

- to explore
- to predict
- to explore

Or

- what is?
- what could be?
- what should be?

Validation

To the extent that the social sciences are concerned with real-world, historical phenomena,

any simulations must be verified (no bugs) and validated (does the model provide behaviour which matches the stylised facts of the historical phenomenon?)

Midgley et al: verification + validation = *assurance*

Back-predictions.

Docking.

Consider the historical market data:

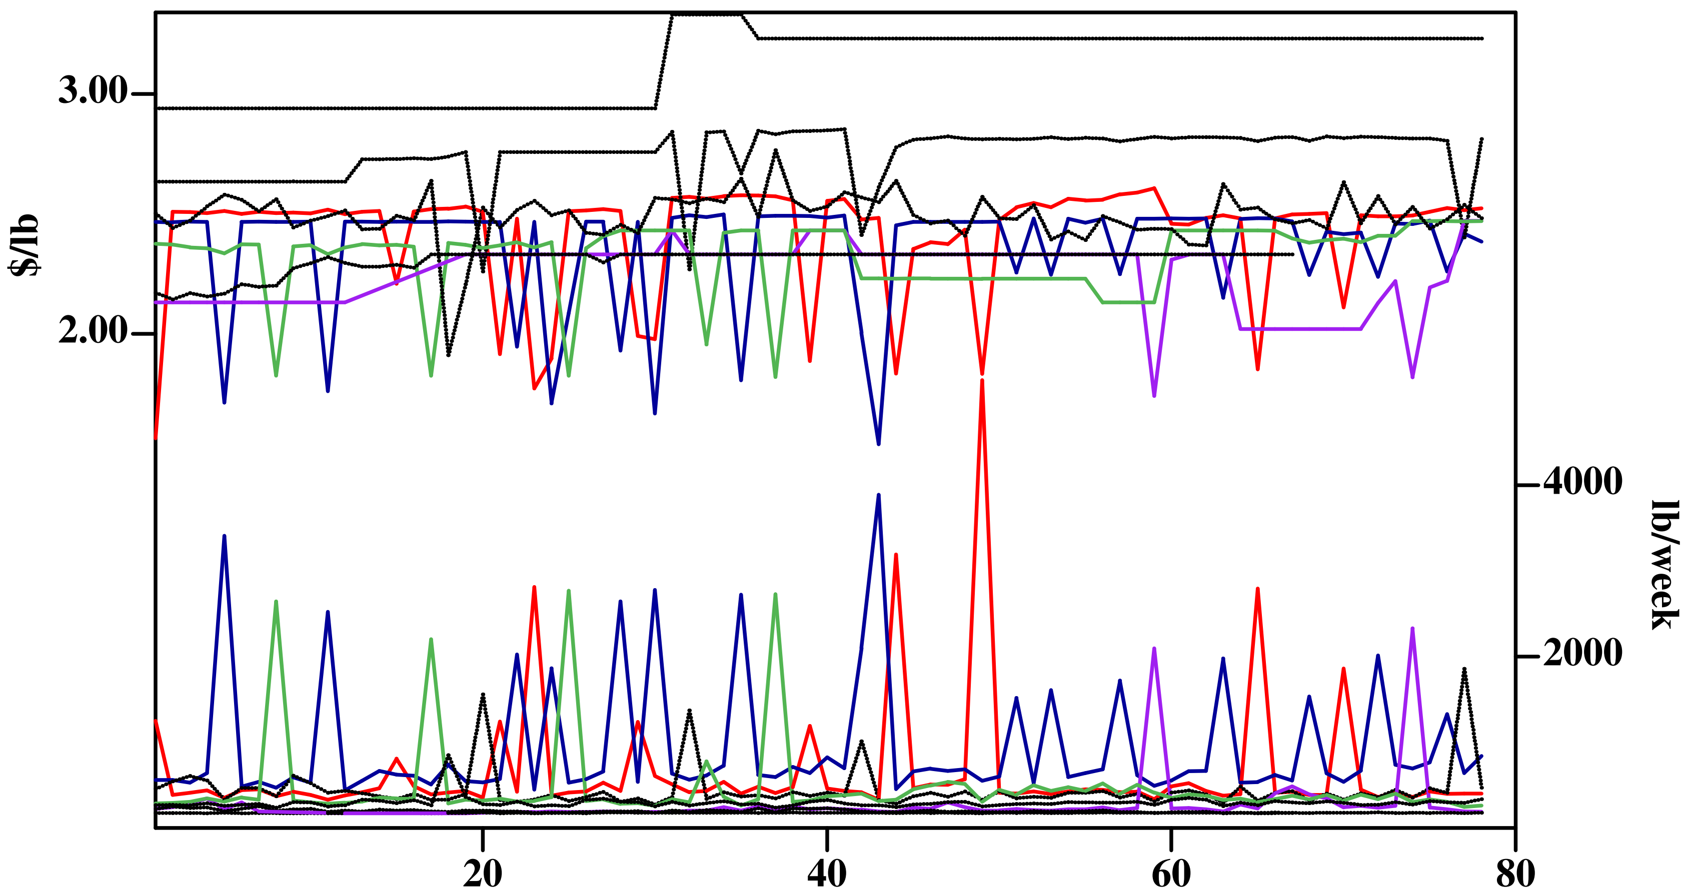

Figure 1: Weekly Sales and Prices (Source: Midgley et al. 1997)

Stylised Facts of the Market Behaviour

- Much movement in prices and quantities of four brands — a rivalrous dance.
- Pattern: high price (and low quantity) punctuated by low price (and high quantity).
- Another four brands: stable prices and quantities

Questions:
What is the cause of these patterns?

— shifts in brand demand?

— reactions by brands?

— actions by the supermarket chain?

— unobserved marketing actions?

Explanations?

Interactions of profit-maximising agents, plus external or internal factors $\longrightarrow$ via a model $\longrightarrow$ behaviour

Similar (qualitatively or quantitatively) to the brands' behaviours of pricing and sales.

Note: assuming profit-maximising (or purposeful) agents means that we are not simply curve-fitting or description using D.E.s. Going beyond the rivalrous dance.

Further ...

With a calibrated model, we can:

perform sensitivity analysis of endogenous with respect to exogenous variables.

Prediction only requires sufficiency, not necessity ("These are the *only* conditions under which the model can work.")

Examine:

- **limits of behaviour (Miller's Automated Non-linear Testing System)**
- **regime-switching**
- **range of behaviour generated**
- **sensitivity of the aggregate (or emergent behaviour) to a single agent's behaviour.**

Validation

Moss & Edmonds (2005): for AB models at least two stages of empirical validation.

1. the micro-validation of the behaviour of the individual agents in the model, by reference to data on individual behaviour.
2. macrovalidation of the model's aggregate or emergent behaviour when individual agents interact, by reference to aggregate time series.

   with the emergence of novel behaviour, possible surprise and possible highly non-standard behaviour, difficult to verify using standard statistical methods.

   $\therefore$ only qualitative validation judgments might be possible.

Formalisation of Validation

Let set P be the possible range of observed outputs of the real-world system.

Let set M be the exhibited outputs of the model in any week.

Let set S be the specific, historical output of the real-world system in any week.

Let set Q be the intersection, if any, between the set M and the set S, $Q \equiv M \cap S$.

We can characterise the model output in several cases. (Mankin et al. 1977).

Five Cases for Validation

a. no intersection between M and S ($Q = \emptyset$), then the model is *useless.*

b. intersection Q is not null, then the model is *useful,* to some degree: will correctly exhibit some real-world system behaviours, will not exhibit other behaviours, and will exhibit some behaviours that do not historically occur. Both incomplete and inaccurate.

c. If M is a proper subset of S ($M \subset S$) then all the model's behaviours are correct (match historical behaviours), but the model doesn't exhibit all behaviour that historically occurs: accurate but *incomplete.*

d. If S is a proper subset of M ($S \subset M$) then all historical behaviour is exhibited, but will exhibit some behaviours that do not historically occur: complete but *inaccurate.*

e. If the set M is equivalent to the set S ($M \Leftrightarrow S$), then (in your dreams!) the model is complete and accurate.

Or Graphically ...

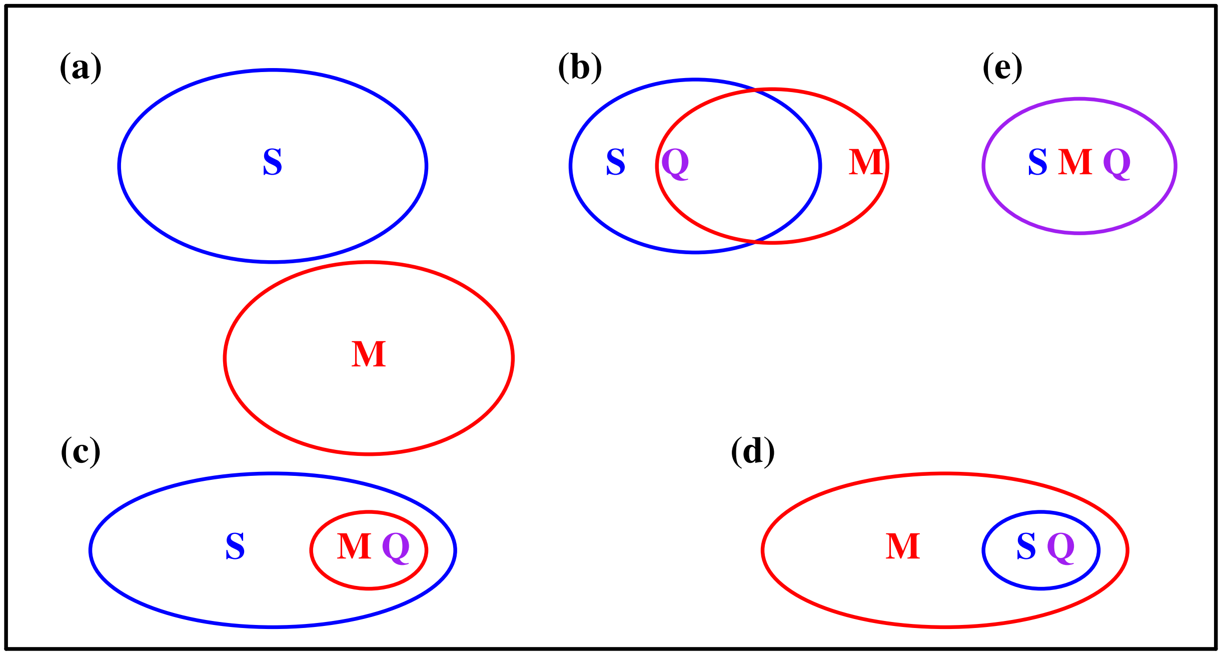

Figure 2: Validity relationships (after Haefner (2005)).

a. useless

b. useful, but incomplete and inaccurate

c. accurate but incomplete

d. complete but inaccurate ← possibly the best to aim for

e. complete and accurate

Modelling Goals

One goal: to construct and calibrate the model so that

$M \approx Q \approx S$: there are very few historically observed behaviours that the model does not exhibit,

and there are very few exhibited behaviours that do not occur historically.

The model is close to being both complete and accurate.

In practice, a modeller might be happier to achieve case d., where the model is complete (and hence provides sufficiency for all observed historical phenomena), but not accurate.

4. Economists´ Uses of Multi-Agent Models

Economic theory moved beyond the reduced forms of perfect competition and monopolies

to oligopolies — markets with small numbers of sellers:

oligopolists interact strategically (not amenable to reduced forms — such as "market clearing")

strategic interactions best modelled by game theory.

But Nash equilibrium, although not reduced — all competitors' behaviours are inputs — but focussed on the *equilibrium*,

while most transactions take place off-equilibrium

Further: real-world firms are heterogeneous ("asymmetric")

First Multi-Agent Systems in Economics

At first, to deal with the non-convexities and discontinuities of the strategy space, agent-based simulation techniques were used to search for equilibrium results, Marks and others used Genetic Algorithms (GAs)

to search a strategy space that was rugged, and non-stationary, when the problem was co-evolutionary, as the other firms also searched for "better" mappings from market state to action.

Coevolution of asymmetric firms requires separate populations in the GA, lest extra-market communication occur: collusion is generally illegal.

"Social" learning v. "individual" learning — Vriend

Coevolution was different from most engineering applications of the GA:
which looked for numerical optima, not the characteristics of the final generation's population.

Agents and the Market Interact

Other AB simulations in economics have used agents that "learnt",
as opposed to the GA's "social" learning of successive generations of agents.

How good do agents have to be?

Gode & Sunder (1993) found — not very, at least for the "double action" market, in which "zero-intelligence players" (ZIPS, who tossed coins or dice) do as well (or even better) than highly rational players (as game theory assumes).

We're still not sure what it is about the structure and rules of the DA market that help the ZIPs.

5. Synthesis in the Social Sciences

Designer markets are exemplified by

1. derivatives markets, in general
2. markets for pollution permits
3. auctions for electro-magnetic spectrum
4. markets for the trade of electricity, and
5. on-line or automated markets

6. Also: contract design.

Methods to Help Designing Markets

Roth (1991) spoke of:

- traditional closed-form game-theoretic analysis
- human-subject experiments
- computer simulations

But Mirowski argues for a typology of market types, perhaps begun in McMillan (2002)

5.1 Designing Markets

Designing markets is complex:

searching for a mapping —

- **— from genotype to phenotype**
- **— from design (structure & rules) space to performance space**
- **— from genotype to phenotype (in evolution)**

In general this is complex (but see Byde 2006): "syntactic complexity" (Edmonds & Bryson 2003). which requires the performance behaviour to "emerge", as Simon (1996) agreed.

A Framework ...

MacKie-Mason & Wellman's Marketplace Design Framework:

Three fundamental steps of a market transaction:

1. the connection (search, discovery)
2. the deal (negotiating, agreeing)
3. the exchange (execution of the transaction).

$\therefore$ two design decisions:

1. the market mechanism (mechanisms for connection, deal, and exchange)
2. the agents

Mirowski argues that economists have focussed on agents and ignored market mechanisms, except in reduced form ("market-clearing")

Design Tradeoffs

For simulation, tradeoffs among possible goals of modelling and simulation must be explicit.

LeBaron's seven basic design questions:

1. the economic environment
2. modelling agents' preferences $\leftarrow$ agents
3. price formation and market clearing
4. the model's fitness
5. information processing and communication $\leftarrow$ agents
6. learning: individual, or social $\leftarrow$ agents
7. benchmarking

6. Similarities and Differences

In 1978 I wrote a Ph.D. thesis, the title of which included the phrase: “disequilibrium dynamics”.

It allowed exchange to take place out of equilibrium — of course, you might say, how else is market-clearing to be attained?

But then it was seen as strange: although there was a short-lived literature along similar lines — Malinvaud, Barro & Grossman, and others.

But this focus on out-of-equilibrium dynamics is precisely what we now need as we simulate using market-based controls.

Simon's Bounded Rationality

Agent-based models, following Simon (1982), also assume Bounded Rationality. Indeed, in the absence of Turing machine (universal calculator), it is difficult not to.

But Epstein (2006) reflects:
"One wonders how the core concerns and history of economics would have developed if, instead of being inspired by continuum physics ... blissfully unconcerned as it is with effective computability — it had been founded on Turing. Finitistic issues of computability, learnability, attainment of equilibrium (rather than mere existence), problem complexity, and undecidability, would then have been central from the start. Their foundational importance is only now being recognized.

Epstein on the virtues of boundedly rational agents …

“As Duncan Foley summarizes:

> `The theory of computability and computational complexity suggest that there are two inherent limitations to the rational choice paradigm.
>
> One limitation stems from the possibility that the agent’s problem is in fact undecidable, so that no computational procedure exists which for all inputs will give her the needed answer in finite time.
>
> A second limitation is posed by computational complexity in that even if her problem is decidable, the computational cost of solving it may in many situations be so large as to overwhelm any possible gains from the optimal choice of action.’ (See Albin 1998, 46).”

ABM $\rightarrow$ Generative Explanation:

Generative explanation (Epstein 2006):

"If you haven't grown it, you haven't explained its emergence."

To answer: how could the autonomous, local interactions of heterogeneous boundedly rational agents generate the observed regularity (that emerges)?

— Generative sufficiency is a necessary but not sufficient condition for explanation. Each realisation is a strict deduction.

Grüne-Yanoff (2006) argues to distinguish *functional explanations* (easier for simulators) from *causal explanations* (much less achievable for social scientists).

Truth and Beauty

Epstein (2006): does AB simulation lack beauty?

Russell: Mathematics as cold, austere, supreme beauty.

Russell: Beauty when "the premises achieve more than would have been thought possible, by means which appear natural and inevitable."

The first damns computer simulation, but the second can occur with emergence from AB models.

Epstein compares different schools of classical music: German v. French.

Truth (from agent-based modelling) can be beautiful too.

Formalisation of Agent-Based Models

Epstein (2006): every agent model is a computer program.

$\therefore$ Turing computable

But for every Turing machine, $\exists$ a unique corresponding and equivalent

partial recursive function.

They might be extremely complex and difficult to interpret, but they exist.

Hence: ”recursive” or “effectively computable” or “constructive” or “generative” (after Chomsky) social science.

7. Cross-Fertilisation

8. Conclusion

9. Simulation

Social Science, not Physical Science

At the aggregate level, similar.

But at the micro level, the agents in social science models are people, with self-conscious motivations and actions.

Aggregate behaviour may be well described by differential equations, with little difference from models of inanimate agents at the micro level.

A Third Way of Doing Science

(from Axelrod & Tesfatsion 2006)

Deduction + Induction + Simulation.

- Deduction: deriving theorems from assumptions
- Induction: finding patterns in empirical data
- Simulation: assumptions $\rightarrow$ data for inductive analaysis

S differs from D & I in its implementation & goals.

S permits increased understanding of systems through controlled computer experiments

Emergence of self-organisation

Examples: ice, magnetism, money, markets, civil society, prices, segregation.

Defn: emergent properties are properties of a system that exist at a higher level of aggregation than the original description of the system.
Not from superposition, but from interaction at the micro level.

Adam Smith's Invisible Hand $\longrightarrow$ prices

Schelling's residential tipping (segregation) model:
People move because of a weak preference for a neighbourhood that has at least 33% of those adjoining the same (colour, race, whatever) $\longrightarrow$ segregation.

Need models with more than one level to explore emergent phenomena.

Families of Simulation Models

1. System Dynamics SD
   (from differential equations)
2. Cellular Automata CA
   (from von Neumann & Ulam, related to Game Theory)
3. Multi-agent Models MAM
   (from Artificial Intelligence)
4. Learning Models LM
   (from Simulated Evolution and from Psychology)

Comparison of Simulation Techniques

Gilbert & Troitzsch compare these (and others):

Technique	Number of Levels	Communication between agents	Complexity of agents	Number of agents
SD	1	No	Low	1
CA	2+	Maybe	Low	Many
MAM	2+	Yes	High	Few
LM	2+	Maybe	High	Many

Number of Levels: “2+” means the technique can model more than a single level (the individual, or the society) and the interaction between levels.

This is necessary for investigating emergent phenomena.

So “agent-based models” excludes Systems Dynamics models, but can include the others.

Verification & Validation

Verification (or internal validity): is the simulation working as you want it to:

— is it "doing the thing right?"

Validation: is the model used in the simulation correct?

— is it "doing the right thing?"

To Verify: use a suite of tests, and run them every time you change the simulation code — to verify the changes have not introduced extra bugs.

Judd's ideas (2006)

"Far better an approximate answer to the right question ... than an exact answer to the wrong question."

— John Tukey, 1962.

That is, economists face a tradeoff between:

the numerical errors of computational work
and
the specification errors of analytically tractable models.

Axelrod on Model Replication and "Docking"

Docking: a simulation model written for one purpose is aligned or "docked" with a general purpose simulation system written for a different purpose.

Four lessons:

1. Not necessarily so hard.
2. Three kinds of replication:
    a. numerical identity
    b. distributional equivalence
    c. relational equivalence
3. Which null hypothesis? And sample size.
4. Minor procedural differences (e.g. sampling with or without replacement) can block replication, even at (b).

Reasons for Errors in Docking

1. Ambiguity in published model descriptions.
2. Gaps in published model descriptions.
3. Errors in published model descriptions.
4. Software and/or hardware subtleties.

   e.g. different floating-point number representation.

(See Axelrod 2006.)

AGENT-BASED MODELS

AB Models are used where the interactions are decentralised, and the autonomous agents make their own decisions (perhaps constrained).

- ∴ **AB models are suitable for interactions which are *bottom-up*, not *top-down*.**

- ∴ **social and market interactions, rather than engineering or internal organisational interactions.**

Using AB models

In ABM/ACE models, a population of software objects is:

- instantiated, and each agent is given:
  - certain internal states (e.g., preferences, endowments) and
  - rules of behaviour (e.g., seek utility improvements).

The agents are then permitted to interact directly with one another and a macrostructure emerges from these interactions.

Patterns Emerge

Patterns in this macrostructure may then be (Axtell, 2005):

— compared with empirical data,

— to revise agent internal states and rules, and

— the process repeated until an empirically plausible model obtains.

e.g. ACE stock markets have been used to model heterogeneous agents: will the stylised features of such markets emerge? *Yes.*

What is an Agent?

An agent: a self-centred program that controls its own actions based on its perceptions of its operating environment.

Derived from the Distributed AI notion of a network of calculating nodes.

Example: the automata in Conway's Game of Life or Schelling's Segregation game or the couples in March & Lave's Sons and Daughters game..

Another example of an agent that won $2,000,000 in a challenge by the U.S. Department of Defense in October 2005 ...

Stanley here.

Agents and agency

Wooldridge & Jennings (1995) would give computer agents these properties:

- autonomy: no others control their actions and internal state,
- social ability: can interact and communicate with other agents
- reactive: they perceive their environment and respond
- pro-active: they initiate goal-directed actions
- (intentionality: metaphors of beliefs, decisions, motives, and even emotions)

Further agent features:

plus (Epstein 1999):

- **heterogeneity: not "representative" but may differ**
- **local interactions: in a defined space**
- **boundedly rational (Simon): information, memory, computational capacity**
- **non-equilibrium dynamics: large-scale transitions, tipping phenomena**

Eight Desired Attributes of Modelled Agents (G&T)

1. Knowledge & beliefs.
   Agents act based on their knowledge of the environment (including other agents), which may be faulty — their beliefs, not true knowledge.
2. Inference.
   Given a set of beliefs, an agent might infer more information.
3. Social models.
   Agents, knowing about interrelationships between other agents, can develop a “social model”, or a topology of their environment: who’s who. etc.

Eight Desired Attributes ...

4. **Knowledge representation.**
Agents need a representation of beliefs: e.g. predicate logic, semantic (hierarchical) networks, Bayesian (probabilistic) networks.

> [Sebastian] Thrun [leader of the winning team in the 2005 DARPA Grand Challenge] had a Zen-like revelation: "A key prerequisite of true intelligence is knowledge of one's own ignorance," he thought. Given the inherent unpredictability of the world, robots, like humans, will always make mistakes. So Thrun pioneered what's known as probabilistic robotics. He programs his machines to adjust their responses to incoming data based on the probability that the data are correct. — Pacella (2005).

Eight Desired Attributes ...

5. **Goals.**
   Agents driven by some internal goal, e.g. survival, and its subsidiary goals (food, shelter). Usually definition and management of goals imposed on the agent.
6. **Planning.**
   Agent must (somehow) determine what actions will attain its goal(s). Some agents modelled without teleology (simple trial-and-error), others with inference (forward-looking), or planning.
7. **Language.**
   For communication (of information, negotiation, threats). Modelling language is difficult. (Want to avoid inadvertent communication, e.g. through the genome of a population in the GA.)
8. **Emotions.**
   Emergent features? Significant in modelling agents? Or epiphenomenal?

How to Model Agent Architecture?

Early approach to modelling cognitive abilities (symbolic paradigm) was fragile, complex, and lacked common sense.

Since then, five approaches:

1. Production Systems
2. Object Orientation
3. Language Parsing & Generation
4. Machine-Learning Techniques, and (most recently)
5. Probabilistic Robotics — Stanley (Thrun et al. 2005).

Ignore 3., 4. last lecture, 5. too new.

Economic Journal June 2005 Feature —

- focussed on Complex Adaptive Systems CAS in economics
- appeared just after Leombruni & Richiardi asked, “Why are economists sceptical about agent-based simulations?” (*Physica A* 355: 103−109, 2005.)
- included 4 papers: introduced by Markose, with papers by Axtell, Robson, and Durlauf
- addressing, respectively,
    - markets as complex adaptive systems,
    - formal complexity issues,
    - the co-evolutionary “Red Queen” effect and novelty, and
    - the empirical and testable manifestations of CAS in economic phenomena.

Markose and the *EJ* Feature on CAS:

- many “anomalies” not understood or modelled using conventional optimisation economics:
  - innovation,
  - competitive co-evolution,
  - persistent heterogeneity,
  - increasing returns,
  - “the error-driven processes behind market equilibrium,”
  - herding,
  - crashes and extreme events such as October 1987.
- need the “adaptive or emergent methods” of ACE simulation

Moreover ...

Axtell (2005) argues that:

- **the decentralised market as a whole can be seen as a collective computing device**
- **the parallel distributed agent-based models of *k*-lateral exchange $\rightarrow$ the specific level of complexity (polynomial) in calculations of equilibrium prices and allocations.**

Validation Relationships

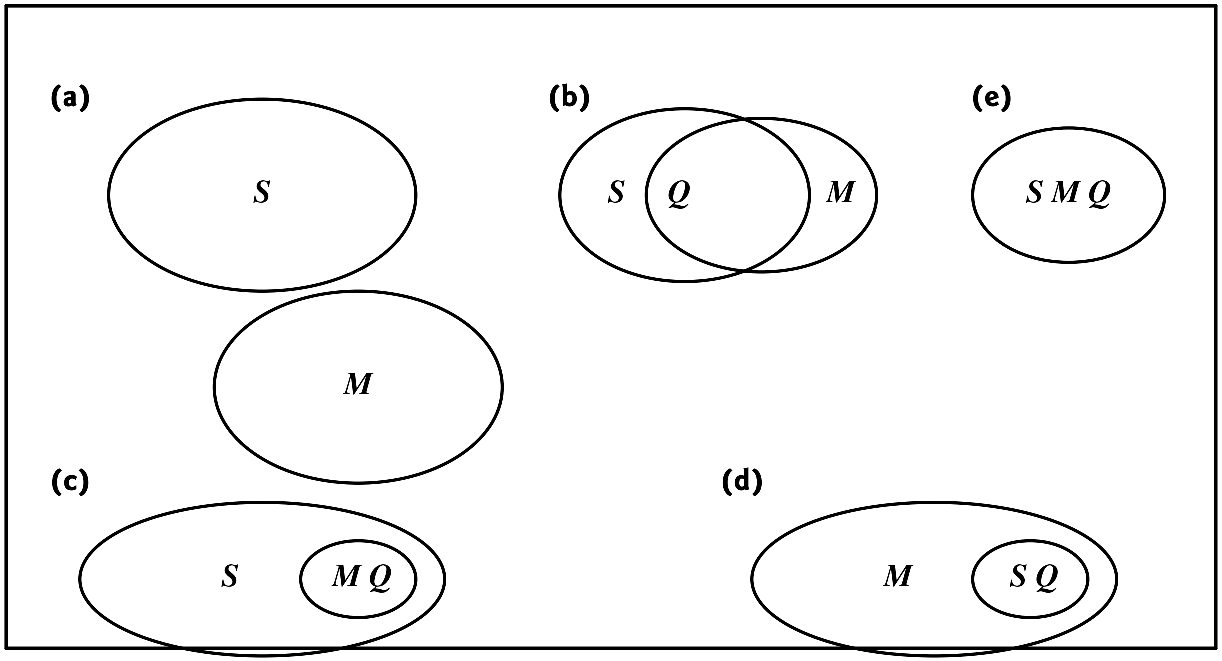

Figure 2: Validity relationships (after Haefner (2005)).

Modelling Goals

One goal: to construct and calibrate the model so that:

$M \approx Q \approx S$: there are very few historically observed behaviours that the model does not exhibit,

and there are very few exhibited behaviours that do not occur historically.

The model is close to being both complete and accurate.

In practice, a modeller might be happier to achieve case d., where the model is complete (and hence provides sufficiency for all observed historical phenomena), but not accurate.

Measures of Validity

A measure of validity which balances the Type I error of inaccuracy with the Type II error of incompleteness.

Define a metric $m()$ (a ratio scale) on the sets.

Define *inaccuracy* α as

$$\alpha \equiv 1 - \frac{m(Q)}{m(M)}, \tag{1}$$

and *incompleteness* γ as

$$\gamma \equiv 1 - \frac{m(Q)}{m(S)}. \tag{2}$$

Continued ...

A measure of degree of validation V: a weighted average of inaccuracy α and incompleteness γ:

$$V \equiv v(1-\alpha) + (1-v)(1-\gamma) \tag{3}$$

$$\therefore V = v\,\frac{m(Q)}{m(M)} + (1-v)\,\frac{m(Q)}{m(S)}$$

$$\therefore V = m(Q)\left(\frac{v}{m(M)} + \frac{1-v}{m(S)}\right) \tag{4}$$

The value of the weight v, $0 \leq v \leq 1$, reflects the tradeoff between accuracy and completeness.

Trade-offs

Possible to reduce incompleteness by generalising the model and so expanding the domain of set *M* until *S* is a proper subset of *M*, as in case d.

Or by narrowing the scope of the historical behaviour to be modelled, so reducing the domain of *S*.

Also be possible to reduce inaccuracy by restricting the model through use of narrower assumptions and so contracting the domain of *M*.

If *M* is sufficiently small to be a proper subset of *S*, as in case c., then the model will never exhibit anhistorical behaviour.

But not guaranteed to maintain a non-null intersection *Q*, and it is possible that the process results in case a., with no intersection.

Look in the Right Place

Reminiscent of the economist looking for his lost car keys under the street light (*M*), instead of near the car where he dropped them in the dark (*S*).

Advocates of simulated solutions, such as Judd (2006), have argued that it is better to “have an approximate answer to the right question, than an exact answer to the wrong question,” to quote Tukey (1962).

References

[1] Arthur, W.B. (1991), Designing economic agents that act like human agents: a behavioral approach to bounded rationality, *American Economic Review* 81: 353–359.

[2] Arthur, W.B. (1993), On designing economic agents that behave like human agents, *Journal of Evolutionary Economics* 3: 1–22.

[3] Arthur, W. B., J. Holland, B. LeBaron, R. Palmer and P. Tayler (1997), Asset pricing under endogenous expectations in an artificial stock market, in: W. B. Arthur, S. Durlauf and D. Lane, eds., *The Economy as an Evolving Complex System II* (Addison-Wesley, Reading, Mass.) 15–44.

[4] Axelrod R., The evolution of strategies in the iterated Prisoner's Dilemma, in L. Davis (ed.), *Genetic Algorithms and Simulated Annealing*, London, Pittman, 1987.

[5] A. Beltratti, S. Margarita and P. Terna, *Neural Networks for Economic and Financial Modelling*, London: International Thomson Computer Press, 1996,

[6] Brenner T., 2006, Agent Learning Representation: Advice on Modelling Economic Learning, Ch 3 in the *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science).

[7] Chattoe, E. (1998), Just how (un)realistic are evolutionary algorithms as representations of social processes? *Journal of Artificial Societies and Social Simulation* 1(3). http://www.soc.surrey.ac.uk/JASSS/1/3/2.html

[8] Curzon Price, T. (1997), Using co-evolutionary programming to simulate strategic behavior in markets, *Journal of Evolutionary Economics* 7(3): 219–254.

[9] Dawid, H. (1999), *Adaptive Learning By Genetic Algorithms: Analytical Results and Applications to Economic Models* (Springer, Berlin), 2nd ed.

[10] Duffy J., 2006, Agent-Based Models and Human Subject Experiments, Ch 4 in the *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, ed. by L. Tesfatsion and K.L. Judd, (Amsterdam: Elsevier Science).

[11] Erev, I. & A.E. Roth (1998), Predicting how people play games: reinforcement learning in experimental games with unique mixed strategy equilibria, *American Economic Review* 88(4): 848–881.

[12] Gjerstad S. & J. Dickhaut (1998), Price formation in double auctions, *Games and Economic Behavior* 22: 1–29.

[13] A. Hailu & S. Schilizzi (2004), Are Auctions More Efficient Than Fixed Price Schemes When Bidders Learn? *Australian Journal of Management*, 29(2): 147–168.

[14] Holland J.H. (1975), *Adaptation in Natural and Artifical Systems,* Univ. of Michigan Press.

[15] Holland J.H. et al., *Induction: Processes of Inference, Learning and Discovery,* Bradford, Cambridge, Mass., 1986.

[16] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, *Neural Networks*, 2(5), 360–365, 1989.

[17] Koza, J.R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection.* MIT Press, Cambridge.

[18] Marks, R.E. (1988) Niche strategies: the Prisoner's Dilemma computer tournaments revisited, http://www.agsm.edu.au/~bobm/papers/niche.pdf

[19] Marks, R.E. (1989) Breeding optimal strategies: optimal behavior for oligopolists, *Proc. of the Third International Conference on Genetic Algorithms,* George Mason University, June 4–7, 1989, ed. by J.D. Schaffer, San Mateo, Calif.: Morgan Kaufmann Publishers, pp.198–207.

[20] Marks R.E., 2006, Market Design Using Agent-Based Models, Ch 12 in the *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, (Amsterdam: Elsevier Science).

[21] Marks R.E., and Schnabl H. (1999) Genetic Algorithms and Neural Networks: a comparison based on the Repeated Prisoner's Dilemma, Thomas Brenner (ed.), *Computational Techniques for Modelling Learning in Economics*, (Dordrecht: Kluwer Academic Publishers), pp. 197–219.

[22] T. Masters, *Practical Neural Network Recipes in C++*, San Diego: Academic Press, 1992.

[23] Roth, A.E. and I. Erev (1995), Learning in extensive form games: experimental data and simple dynamic models in the intermediate term, *Games and Economic Behavior* 8: 848−881.

[24] N.N. Schraudolf & J.J. Grefenstette, A User's Guide to GAucsd 1.4, Technical Report CS92-249, CSE Department, UC San Diego, 1992.

[25] Selten, R. (1998), Features of experimentally observed bounded rationality, *European Economic Review* 42: 413−436.

[26] Simon, H.A. (1955), A Behavioral Model of Rational Choice, *Quarterly Journal of Economics*, 69, 99−118.

[27] H. A. Simon. *Models of bounded rationality,* MIT Press, 1982.

[28] H. A. Simon (1983) Why should machines learn?, in *Machine Learning: An Artificial Intelligence Approach,* R. S. Michalski, J. G. Carbonell and T. M. Mitchell, editors 1983, Springer-Verlag. Berlin. pp. 25−37.

[29] H. A. Simon , *The Sciences of the Artificial,* The MIT Press, 2nd ed., 1996.

[30] Szpiro G., The emergence of risk aversion, *Complexity*, 2(4), 31−39, 1997.

[31] N. J. Vriend, An illustration of the essential difference between individual and social learning, and its consequences for computational analyses, *Journal of Economic Dynamics & Control,* 24: 1−19, 2000.

References:

- R. Axelrod, Advancing the Art of Simulation in the Social Sciences, in J.-P. Rennard (ed.), *Handbook of Research on Nature-Inspired Computing for Economy and Management,* (Hershey, PA: Idea Group Inc., 2006)
- R. Axelrod & L. Tesfatsion, On-Line Guide for Newcomers to Agent-Based Modeling in the Social Sciences, in L. Tesfatsion & K.L. Judd (eds.), *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, North-Holland, Amsterdam, 2006. www.econ.iastate.edu/tesfatsi/abmread.htm
- S. Durlauf, Complexity and empirical economics, *The Economic Journal*, 115 (June), F225–F243, 2005.
- N. Gilbert & K.G. Troitzsch, *Simulation for the Social Scientist,* Open Uni Press, 2nd ed. 2005.
- A. Hailu & S. Schilizzi, Are Auctions More Efficient Than Fixed Price Schemes When Bidders Learn? *Australian Journal of Management*, 29(2): 147–168, December 2004. www.agsm.edu.au/eajm/0412/hailu_etal.html
- S. Hartmann, The world as a process: Simulations in the natural and social sciences. In R. Hegselmann, U. Mueller, & K.G. Troitzsch, eds., *Modelling and simulation in the social sciences: From the philosophy of science point of view,* vo. 23 of *Series A: Philosophy and methodology of the social sciences,* pp. 77–100. Kluwer Academic Publishers, 1996.
- K. L. Judd, Computationally Intensive Analyses in Economics, *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, ed. by Leigh Tesfatsion & Kenneth L. Judd, Amsterdam: Elsevier Science, 2006, Ch. 2.
- B. Latané, Dynamic social impact: Robust predictions from simple theory. In R. Hegselmann, U. Mueller, & K.G. Troitzsch, eds., *Modelling and simulation in the social sciences: From the philosophy of science point of view,* vo. 23 of *Series A: Philosophy and methodology of the social sciences,* pp. 287–310, Kluwer Academic Publishers, 1996.
- J. March & C. Lave, *Introduction to Models in the Social Sciences*, New York: HarperCollins, 1975.
- M. Resnick. *Turtles, termites, and traffic jams: Explorations in massively parallel microworlds*. MIT Press, 1997.

References:

- P.S. Albin, ed. 1998. *Barriers and Bounds to Rationality: Essays on Economic Complexity and Dynamics in Interactive Systems*. PUP, Princeton. Introduction by D.K. Foley.
- R. Axtell, The complexity of exchange, *The Economic Journal,* 115 (June), F193−F210, 2005.
- S. Durlauf, Complexity and empirical economics, *The Economic Journal*, 115 (June), F225−F243, 2005.
- J.M. Epstein, Agent-based computational models and generative social science, *Complexity*, 1999.
- J.M. Epstein, Remarks on the Foundations of Agent-Based Generative Social Science, *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, edited by Leigh Tesfatsion and Kenneth L. Judd, Amsterdam: Elsevier Science, 2006.
- T. Grüne-Yanoff, The Explanatory Potential of Artificial Societies, *Models and Simulations* (Paris, 2006).
- A. Hailu & S. Schilizzi, Are Auctions More Efficient Than Fixed Price Schemes When Bidders Learn? *Australian Journal of Management*, 29(2): 147−168, December 2004.
- R. Leombruni & M. Richiardi, Why are economists sceptical about agent-based simulations? *Physica A* 355: 103−109, 2005.
- S.M. Markose, Computability and evolutionary complexity: markets as complex adaptive systems (CAS), *The Economic Journal*, 115 (June), F159−F192, 2005.
- R.E. Marks (2007), Coffee, segregation, energy, and the law: validating simulation models, *Computational Economics*, forthcoming.
- D.F. Midgley, Marks R.E., and Kunchamwar D. (2007) The Building and Assurance of Agent-Based Models: An Example and Challenge to the Field, *Journal of Business Research*, Special Issue: Complexities in Markets, forthcoming.
- R.M. Pacella (2005), DARPA Grand Challenge, *Popular Science*, October.
- A.J. Robson, Complex evolutionary systems and the Red Queen, *The Economic Journal,* 115 (June), F211−F224, 2005.

- Simon, H.A. 1982. *Models of Bounded Rationality*. Cambridge, MA: MIT Press.
- Simon, H.A. 1987. Rationality in Psychology and Economics. *Rational Choice: The Contrast Between Economics and Psychology*, Hogarth, R.M. and Reder, M.W. eds. Chicago: University of Chicago Press.
- S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics*, MIT Press, September 2005.
- M. Wooldridge & N. R. Jennings, Intelligent Agents: Theory and Practice, *Knowledge Engineering Review,* 1995.