MARKET-BASED SYSTEMS

University of St Andrews
Department of Computer Science
Distinguished Lecture Series, March 2008

Lecture 3

Dave Cliff

Department of Computer Science
University of Bristol

dc@cs.bris.ac.uk

University of BRISTOL

The Contents

NB This is a 3-hour partial tutorial overview of Market-Based Systems …in three 60min chunks

- partial as in incomplete: we can't cover everything in three hours
- partial as in biased: this is my version of the story…

- Lecture 1: Rationale and Background

  Here we'll find out why computer scientists should care about market-based systems, review some notable applications, and also cover some of the background economics. They call economics "the dismal science" for a reason, so that background economics stuff won't delay us too long...

- Lecture 2: Artificial Trading Agents for Fun and Profit

  This lecture tells the story of some of the best-known algorithms used for autonomous "trader-robots", and how they were found to consistently beat human traders.

- Lecture 3: What's hot, what's not, and where next: Tales from the City
  ---
  Looks at work on automatic optimization and design of trader-agents, and online market mechanisms, with particular reference to the current hot topics in the automated trading technology in the financial markets.

Recap of Lectures 1& 2…

- Transition to utility-scale/"cloud" computing & service-oriented business models gives need for automated trading strategies/software
- Human traders do a lot, quickly
- More than zero intelligence is necessary
- Non-economists at HP and at IBM motivated to develop trader "robots"
- The Autonomous Agents academic research community start playing games
- 1997: ZIP open-sourced
- 2001: IBM do ZIP a very nice favour
- The financial markets start to take an interest
- But there was something that IBM didn't know…

GA-optimization of ZIP traders

- Meanwhile at MIT & HP Labs…
  - Genetic Algorithms (GAs) to optimize ZIP-trader parameters to particular markets
    - First presented at CIFEr'98, New York, Mar'98; & ASCMA98, Minneapolis, May '98
    - Tech report in 2001: http://www.hpl.hp.com/techreports/2001/HPL-2001-99.html
  - if GA-tuned ZIP traders had been used by IBM, maybe ZIP would have dominated
- ZIP-trader marketplaces have 8 control parameters
  - initially set by educated guesswork
- Trivial to use a genetic algorithm (GA) to optimize the 8 parameters
  - genome: a single point in the eight-dimensional real hyperspace $\mathbf{R}^8$
    - actually, points constrained to lie within the unit hypercube in 8-space
  - vanilla GA: with annealing mutation; population size 30; 500 generations

ZIP: quantitative margin adjustments

- ZIP algorithm is adaptive: adjusts margins up or down using simple machine learning rules
- Quote-price $p_i(t)$ set by limit price λ_i and margin $\mu_i(t)$: $p_i(t) = \lambda_i \cdot (1+\mu_i(t)$
  - Seller A $\mu_i(t)$ in $[0,\infty]$ forall t; $\mu_i(t)$+= raises margin; $\mu_i(t)$-= lowers margin
  - Buyer $\mu_i(t)$ in [illegible] raises margin; $\mu_i(t)$+= lowers margin
- ZIP uses Widrow[illegible]g rule to adjust actual output A wrt desired D using rate β:
  - $A(t+1) = A(t) + \Delta(t)$; where $\Delta(t) = \beta \cdot (D(t) - A(t))$
  - With momentum (damping) factor γ in $[0,1]$: $A(t+1) = \gamma \cdot A(t) + ((1$ [illegible] $0 \leq \gamma$
- So for ZIP we have:
  - $\mu_i(t+1) = (p_i(t) + \Delta_i(t)) / \lambda_i - 1$
  - $\Delta_i(t) = \beta_i \cdot (\tau_i(t) - p_i(t))$; where target price $\tau_i(t) = (A_i(t) + R_i(t) \cdot q(t))$; $A()$ & $R()$ stochastic
- Giving:
  - $\mu_i(t+1) = (p_i(t) + \Gamma_i(t)) / \lambda_i - 1$; where $\Gamma_i(0)=0$ and $\Gamma_i(t+1) = \gamma_i \cdot \Gamma_i(t) + ((1- \gamma_i) \cdot \Delta_i(t))$

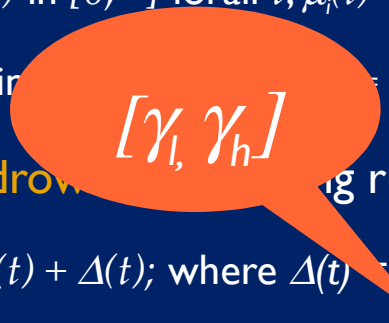

The Genetic Algorithm: not a lot of code

See e.g.:

D. Goldberg
Genetic Algorithms
Addison-Wesley, 1986

M. Mitchell
An Introduction to Genetic Algorithms
MIT Press, 1998

```
/*population size*/
POP_SIZE=30;
/*number of generations to evolve for*/
MAX_GENS=200;
/*create initial population at random*/
generate_random_population(Pop1[]);
/*evolve…*/
for(g=0;g<MAX_GENS;g++)
{  /*evaluate fitness of each member of population*/
   for(i=0;i<POP_SIZE;i++)
   { evaluate_fitness(Pop1[i]);}
   /*identify the best member*/
   elite=find_fittest(Pop1[]);
   /*breed new population using tournament selection*/
   for(i=0;i<POP_SIZE;i++)
   { /*randomly pick 3 distinct possible parents*/
     parents[1]=irandpick1st(POP_SIZE);
     parents[2]=irandpick2nd(POP_SIZE,parents[1]);
     parents[3]=irandpick3rd(POP_SIZE,parents[1],parents[2]);
     /*sort them into order of fitness*/
     sort_fitness(parents[]);
     /*the two fittest parents "breed" to make a kid*/
     Pop2[i]=breed_new_kid(parents[1],parents[2]);
   }
   /*preserve the elite*/
   Pop2[0]=Pop1[elite];
   /*new population replaces old population*/
   Pop1[]=Pop2[];
}
```

GA optimizing from "easy" genomes

- Initial population genomes seeded with original parameter values as used in the initial ZIP trader studies. 200 generations; populationsize=30.

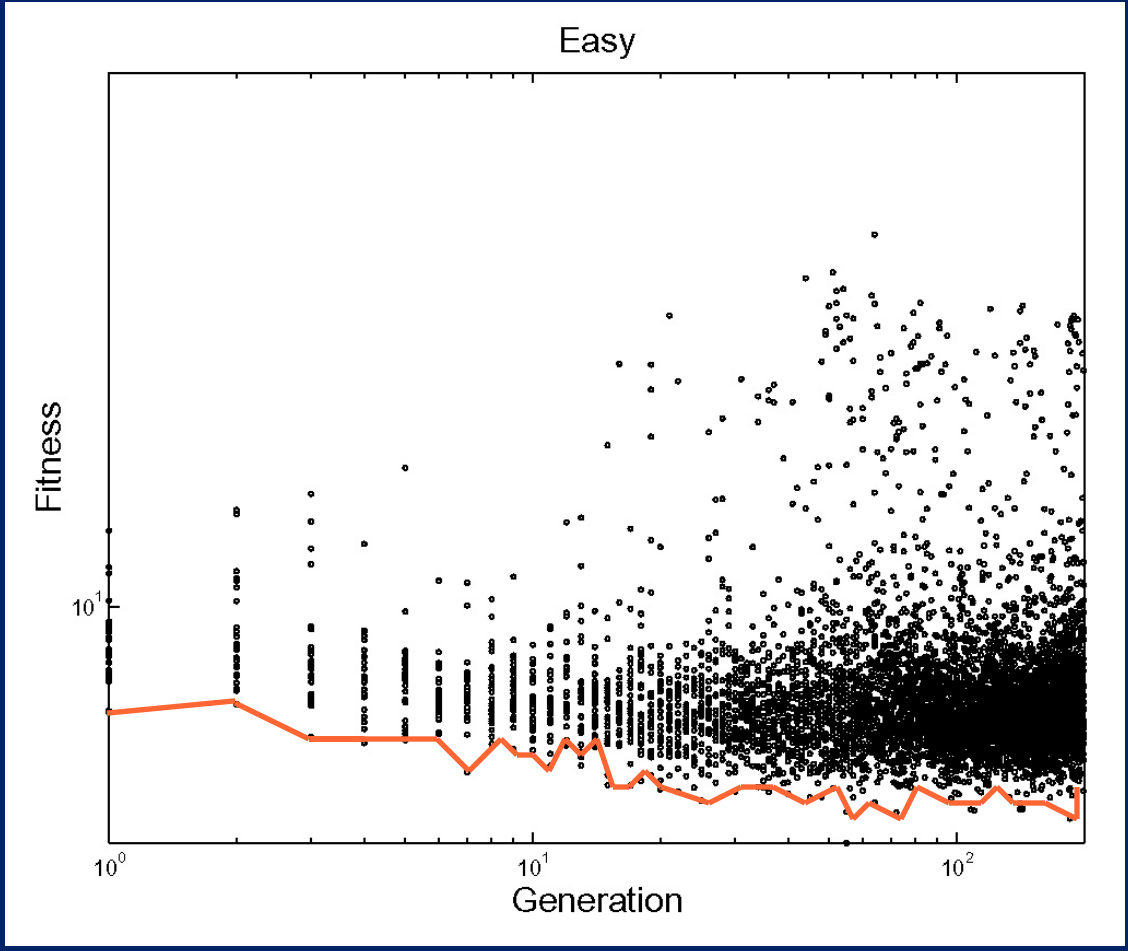

Slight improvement: shows that the GA can improve on the parameter values used by the inventor of the ZIP algorithm

GA optimizing from "zero" genomes

- Initial population genomes all seeded with positively unhelpful (0,0,0,0,0,0,0,0) values
- 200 generations; populationsize=30

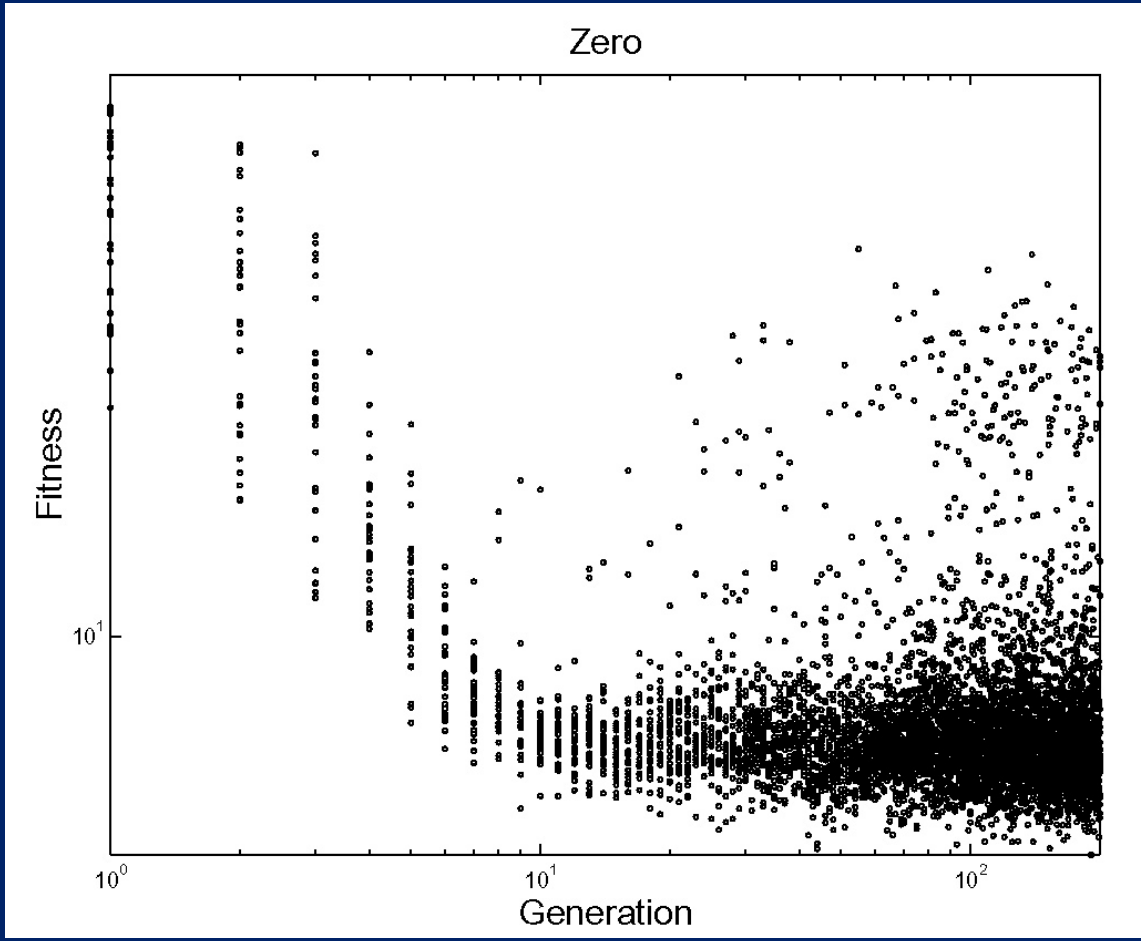

Definite improvement: shows that the GA can improve even when commenced inconveniently with traders that (initially) have no adaptation or memory

GA optimizing from "hard" genomes

- Initial population genomes all seeded with absolutely ridiculous parameter values: each value deliberately way too high or way too low. 200 gens; populationsize=30.

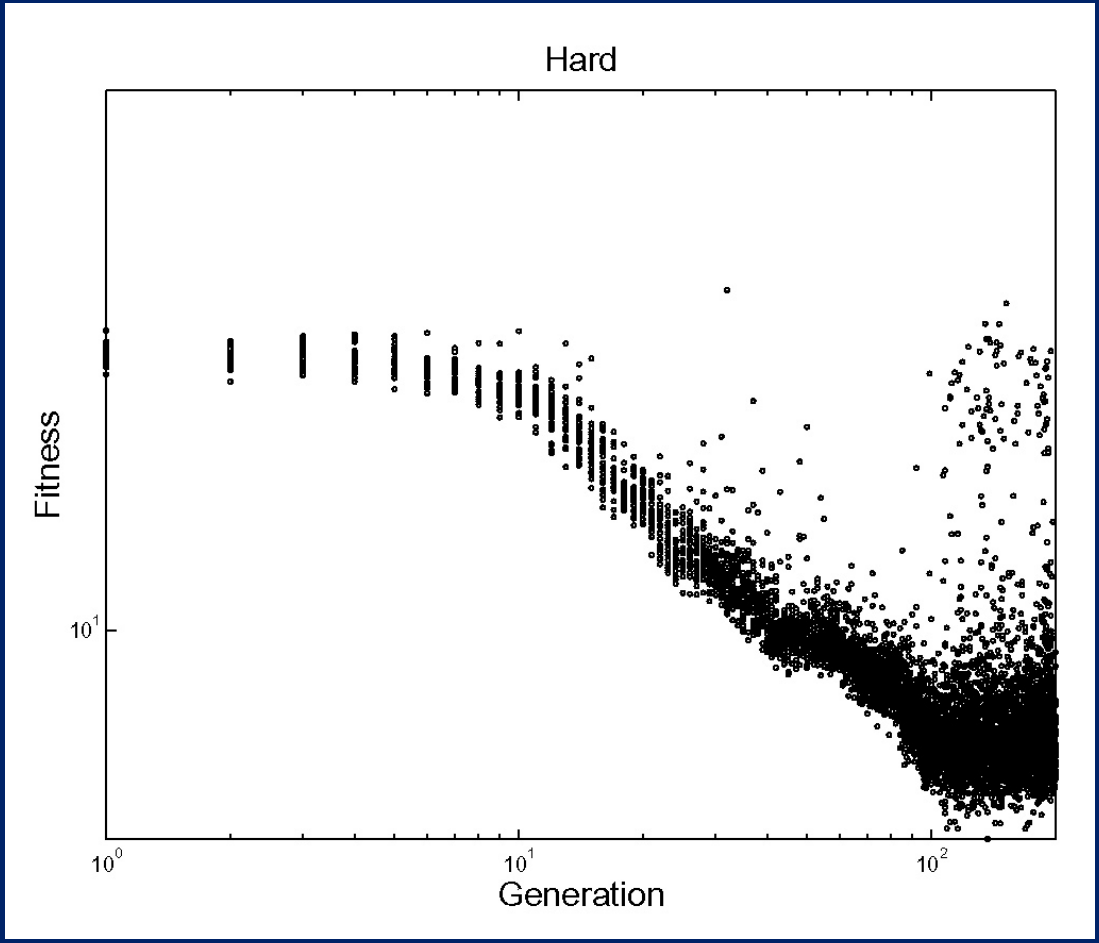

Definite improvement: shows that the GA can improve even when initiated from trader genomes with maliciously/idiotically poorly-chosen parameter values

If the computer twiddles the knobs…

…why limit the number of knobs?

- ZIP60, and ZIP100, both have too many parameters for a human to set by hand without dying of boredom.
- But a small cluster of PCs can happily spend a couple of days/weeks finding the right settings for the 60 or 100 knobs.
- And, on a decent-sized HPC facility, you get the twiddling done in a few minutes.
- "Real-world" trading algorithms are now routinely being developed and deployed that assume the availability of adaptation, learning, and/or optimization.
- But if all the traders (or just the majority of traders) in the marketplace are robots, why use a human-compatible marketplace?

Now evolve the auction mechanism too

- If you're going to use trading agents instead of humans, then why use market mechanisms designed by humans for humans?
- Use GA to search a space of possible auction types
- GA simultaneously co-adapts ZIP trader parameters, as before
- Fitness measure: *minimize* root mean square deviation of transaction prices from equilibrium price
  - (front-weighted average of Smith's α, as before)
- Problem: how to encode for a range of auction styles?

- D. Cliff. Evolution of market mechanism through a continuous space of auction-types. *Computational Intelligence in Financial Engineering (CIFEr)*, Hawaii, May 2002. *http://www.hpl.hp.com/techreports/2001/HPL-2001-326*

A continuum of auction mechanisms

- Let Q_s denote the probability that the next quote comes from a seller,
  - i.e. Q_s denotes the probability that the next quote is an offer/ask
  - NB probability of a quote coming from a buyer Q_b=1.0-Q_s
- In the English Auction, Q_s=0.0
- In the Dutch Auction, Q_s=1.0
- In a CDA a quote is equi-probable from either the seller-side or the buyer-side, so Q_s=0.5
- What if we interpret these 3 human-designed mechanisms (0.0, 0.5, & 1.0) as points on a Q_s continuum?
  - e.g. how about auction based on Q_s=0.1?
- Nonstandard values of Q_s are easily implementable in online e-marketplaces

Doing it for real: really not a lot of code.

```
Real q_s=0.1;


r=uniform_random_real(0.0,1.0);
if(r<q_s)
{ get_next_quote(sellers); }
else
{ get_next_quote(buyers); }
```

Doing it for real: manually, with humans

- Spin-the-wheel…

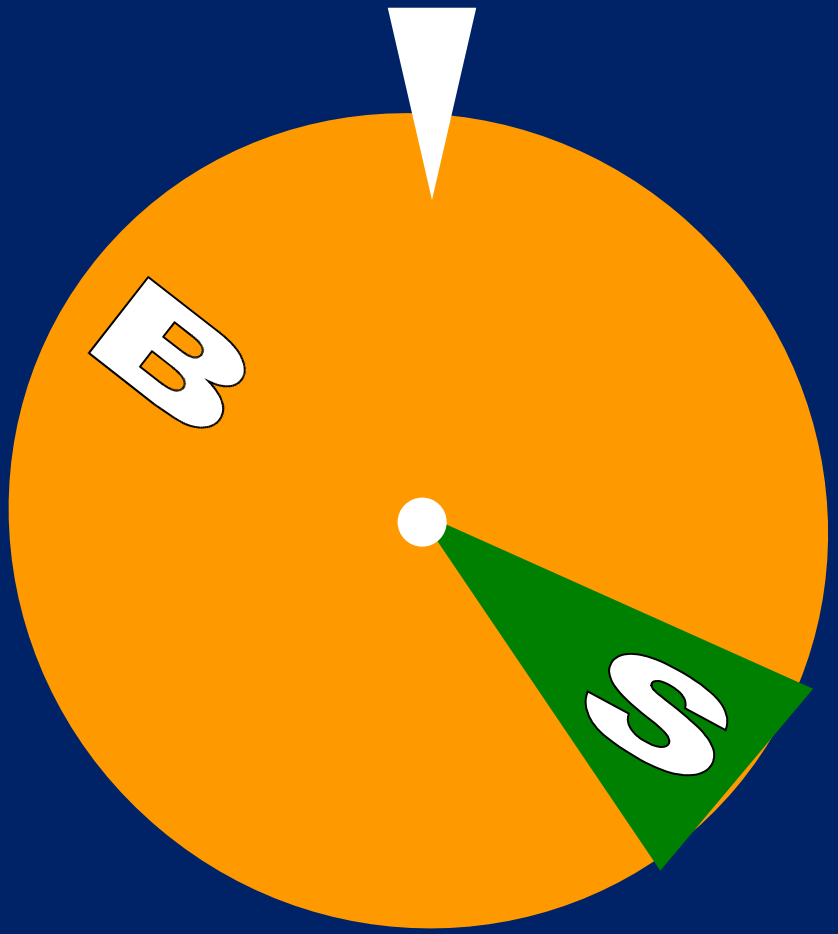

Effects of shock-changes to schedules

- Maybe the fact that a single unchanging schedule is used for the duration of each experiment is too simplistic?
  - The ZIP traders don't have to deal with any changes in supply/demand
  - Over-fitting?
- Using same experiment set-up, see what Q_s evolves when, partway through each evaluation, there is a sudden "shock" change in the supply and demand schedules

- http://www.hpl.hp.com/techreports/2002/HPL-2002-128

Lots of results…

- data in orange is significantly different from CDA Q_s=0.5 (Wilcoxon-Mann-Whitney@1%)
- The non-CDA auctions are all better than CDA
- D. Cliff. "Explorations in evolutionary design of online auction market mechanisms", *Journal of Electronic Commerce Research and Applications.* 2(2):162-175, 2003.
- Hypothesis: there are exploitable regularities in most markets, most of the time.

M1 only: Q_s=~0.0

M2 only: Q_s=~0.0

M3 only: Q_s=~0.16

M1M2 "shock": Q_s=~0.25

M2M1 "shock": Q_s=~0.50

M2M3 "shock": Q_s=~0.56

M3M2 "shock": Q_s=~0.14

M1M2M1 "shock": Q_s=~0.51

M2M1M2 "shock": Q_s=~0.50

M2M3M2 "shock": Q_s=~0.58

M3M2M3 "shock": Q_s=~0.52

M1M2M3 "shock": Q_s=~0.56

M3M2M1 "shock": Q_s=~0.47

That was just the start

- Replications and extensions by postgraduate students…
  - Zengchang Qin (Bristol, 2002) replicated & extended to pure English/Dutch auctions
  - Vibhu Walia (Birmingham, 2002) demonstrated evolution of hybrid auctions in ZI-C traders
  - Neil Robinson (Sussex, 2002) evolved hybrid auction mechanisms for MBC of UDC
  - David Shipp (Leeds, 2004) explored longer sequences of supply/demand shocks
  - Dan Wichett (Birmingham, 2004) explored co-adaptive dynamics of heterogeneous gene-pools

- Andrew Byde (HP) demonstrated GA-evolution of optimally hybrid n^{th}-price sealed-bid auctions, regardless of intelligence of traders

- Dave Cliff (HP) developed new 60-parameter super-variant of ZIP traders
  - 60-dimensional hyperspace better than the original 8-parameter version
  - evolutionary control of dimensionality demonstrated to be beneficial
- See www.ziptrader.org and www.ziptrader.org/zip60; *IEEE TransEC:* 2008

New Scientist May 25th, 2002

TRADING PLACES: who needs people to do business?

Market maker

Earn megabucks from natural selection

A STOCK market that's not only run by computers but designed by them too would generate hundreds of billions of dollars' worth of extra profit for their human controllers. So says artificial intelligence expert David Cliff, who applied the principles of evolution to allow his software to develop the ideal marketplace.

Cliff's software created a super-auctioneer that can orchestrate buying and selling in a new type of stock market: one where the traders are all software-based agents. The idea could even introduce an ethical dimension to trading. Cliff found that in his marketplace, fewer people paid inflated prices and everyone was able to maximise profits.

Last year, *New Scientist* reported on research from IBM which showed that software agents trading commodities such as precious metals or coffee beans outperformed their human counterparts, earning on average 7 per cent more (11 August 2001, p 21). One of the software agents used in that experiment was developed by Cliff, who works at Hewlett-Packard's research lab in Bristol.

But Cliff found it odd that these software agents were being shoehorned into computer models of marketplaces designed by people for people. So he set about designing a marketplace that is better suited to artificial buying and selling agents.

To do this, he designed a range of possible markets using software incorporating a genetic algorithm, a type of program inspired by evolution. The GA uses a process similar to natural selection to "evolve" continually until it reaches the best solution to a problem. The idea is that for every set of trading conditions, there is an ideal marketplace, and it's this the GA seeks. It evolves by putting a range of parameters through a series of mutations until it maximises profit. For example, Cliff's GA found that one of the most critical factors is the number of "traders" allowed to bid at any one time.

Today's markets use a system known as a continuous double auction, where buyers and sellers simultaneously yell offers at each other, with the sellers slowly lowering their asking prices and the buyers gradually increasing their offers.

The GA found that the fittest market of this sort was one that allowed three buyers to make bids for every sell offer, generating 2 per cent more profit than a typical market. In the real world, with trillions of dollars percolating through the markets, this small gain might translate into hundreds of billions of dollars, says Cliff.

One crucial difference between the evolved marketplace and a human one is that traders have to reveal their maximum and minimum prices to the auctioneer program, though not to each other, before the market can be set up for trading. This gives it key data on which to base its evolution.

But this strategy would only work if the information were kept confidential and was known only to the marketplace itself, says Lawrence Barwick, head of proprietary trading at the Bank of America in London. "If it's closed like that and everybody has to play by the same rules then it could work."

Duncan Graham-Rowe

25 May 2002 ◆ New Scientist ◆ www.newscientist.com

11

Quantum foot in the door

ALL around us are tiny doors that lead to the rest of the Universe. Predicted by Einstein's equations, these quantum wormholes offer a faster-than-light shortcut to the rest of the cosmos—at least in principle. Now physicists believe they could open these doors wide enough to allow someone to travel through.

Quantum wormholes are thought to be much smaller than even protons and electrons, and until now no one has modelled what happens when something passes through one. So Sean Hayward at Ewha Womans University in Korea and Hisa-aki Shinkai at the Riken Institute of Physical and Chemical Research in Japan decided to do the sums.

They have found that any matter travelling through adds positive energy to the wormhole. That unexpectedly collapses it into a black hole, a supermassive region with a gravitational pull so strong not even light can escape.

But there's a way to stop any would-be traveller being crushed into oblivion. And it lies with a strange energy field nicknamed "ghost radiation". Predicted by quantum theory, ghost radiation is a negative energy field that dampens normal positive energy. Similar effects have been shown experimentally to exist.

Ghost radiation could therefore be used to offset the positive energy of the travelling matter, the researchers have found. Add just the right amount and it should be possible to prevent the wormhole collapsing—a lot more and the wormhole could be widened just enough for someone to pass through.

It would be a delicate operation, however. Add too much negative energy, the scientists discovered, and the wormhole will briefly explode into a new universe that expands at the speed of light, much as astrophysicists say ours did immediately after the big bang.

For now, such space travel remains in the realm of thought experiments. The CERN Large Hadron Collider in Switzerland is expected to generate one mini-black hole per second, a potential source of wormholes through which physicists could try to send quantum-sized particles. But sending a person would be another thing. To keep the wormhole open wide enough would take a negative field equivalent to the energy that would be liberated by converting the mass of Jupiter.

Charles Choi

More at: www.arxiv.org/abs/gr-qc/0205041

The Economist Nov 30th, 2002

The future's still not clear

French national champion, by pushing Crédit Lyonnais into the arms of Crédit Agricole. But when Crédit Agricole refused to pay €44 per share ($44), the lowest price he was prepared to accept, he lost patience. BNP Paribas, the auction's winner, bid €2.2 billion, or €58 a share, for the stake, a 49% premium over the market value and a quarter more than the next-highest bid. Advisers to Mr Mer joked that BNP Paribas's bid document must have contained a typing error.

The losers from the sudden sale are Crédit Agricole and Crédit Lyonnais's chief executive, Jean Peyrelevade. Having refused to pay €44 a share in a private deal, Crédit Agricole reportedly offered that amount in the public auction, only to lose. Part of the problem is that Crédit Agricole's decentralised, mutual structure makes it difficult for its "barons" in the countryside to agree on anything quickly.

Mr Peyrelevade, who wants to keep his bank independent, has so far maintained what one French banker calls a Soviet balance of power between the bank's biggest shareholders. As soon as one looms too large, he courts another. This time though, he miscalculated, stringing out his negotiations with Crédit Agricole to a point where Mr Mer stopped believing that the two banks had any real desire to merge.

Mr Peyrelevade now faces the likelihood that BNP Paribas or, just possibly, Société Générale, will bid for the whole of Crédit Lyonnais once a shareholder pact designed to protect the bank from takeover ends at the end of June 2003. Either bank would fold Crédit Lyonnais's wholesale banking business into their own operations. To avoid that, Mr Peyrelevade may have to run back to Crédit Agricole, which already owns 10.5%, with more conciliatory words than those of recent months.

For now, BNP Paribas denies that it intends to buy the whole of Crédit Lyonnais. But Michel Pebereau, its chairman, is unlikely to have paid such a high price for the stake unless he has designs on the whole. With Crédit Lyonnais in its grip, BNP Paribas would increase its share of the French retail banking market from 8% to 14%.

One likely opponent of any such move would be AGF, the French insurance subsidiary of Germany's Allianz, which owns about 10% of Crédit Lyonnais. French bankers are preparing themselves for a battle to echo the one that Banque Nationale de Paris (BNP), Paribas and Société Générale waged in 1999, which ended with BNP wresting control of Paribas. As one banker says, "BNP Paribas is usually the winner." ■

Financial markets

Robo-traders

Computerised trading agents may help humans build better markets

THANKS to slumping markets, investment banks are shedding many of their highly-paid traders. When markets recover, the banks might be tempted to replace them with rather cheaper talent. One alternative has been around for a while but has yet to catch on: autonomous trading agents—computers programmed to act like the human version without such pesky costs as holidays, lunch breaks or bonuses. Program trading has, of course, been done before; some blamed the 1987 stockmarket crash on computers instructed with simple decision-making rules. But robots can be smarter than that.

Dave Cliff, a researcher at Hewlett-Packard Laboratories in Bristol, England, has been creating trading robots for seven years. In computer simulations he lets them evolve "genetically", and so allows them to adapt and fit models of real-world financial markets. His experiments have suggested that a redesign of some markets could lead to greater efficiency.

Last year, a research group at IBM showed that Mr Cliff's artificial traders could consistently beat the human variety, in various kinds of market. Nearly all take the shape of an auction. One well-known type is the English auction, familiar to patrons of the salerooms of Christie's and Sotheby's, where sellers keep mum on their offer price, and buyers increase their bids by stages until only one remains.

At the other extreme is the Dutch auction, familiar to 17th-century tulip-traders in the Netherlands as well as to bidders for American Treasury bonds. Here, buyers remain silent, and a seller reduces his price until it is accepted. Most markets for shares, commodities, foreign exchange and derivatives are a hybrid of these two types: buyers and sellers can announce their bid or offer prices at any time, and deals are constantly being closed, a so-called "continuous double auction".

Mr Cliff's novel idea was to apply his evolutionary computer programs to marketplaces themselves. Why not, he thought, try and see what types of auction would let traders converge most quickly towards an equilibrium price? The results were surprising. In his models, auctions that let buyers and sellers bid at any time—like most of today's financial exchanges—were less efficient than ones that required relatively more bids from either buyers or sellers.

These "evolved auctions" also withstood big market shocks, such as crashes and panics, better than today's real-world versions. Mr Cliff's most recent results, which will be presented in Sydney, Australia, on December 10th, show that the best type of auction for any market depends crucially on even slight differences in the number of buyers and sellers.

Bank of America has been investigating these new auctions, along with robotic traders, for possible use in electronic exchanges. The hope is that today's financial auctions and online marketplaces might work better by becoming more like their English and Dutch forebears. But what to call such multi-ethnic hybrids? Here's introducing the "Cliffhanger". ■

Financial News, 24 March 2005

"Algorithmic Trading has the potential to replace human traders completely, according to one of the scientists who pioneered the use of advanced trading systems in the financial markets. Speaking at last week's Financial News Trading for Investors conference in London, David Cliff, head of complex adaptive systems research at Hewlett-Packard, said: "I don't see why everything a trader does can't be done by a machine. I don't want to say when it will happen as the technology is not there yet."

http://www.efinancialnews.com/

The Trade, Jan-Mar 2005, p.30

"'Traders under threat from rise of machines', screamed a recent headline in the Financial News. It seems a professor of complex adaptive systems who dresses like the bad guy in a Bond movie has been talking about traders being an unnecessary cost who can be replaced by algorithms."

Richard Balarkas,
Global Head of AES Sales,
Credit Suisse

http://www.thetrade.ltd.uk/cissue_realitycheck.htm

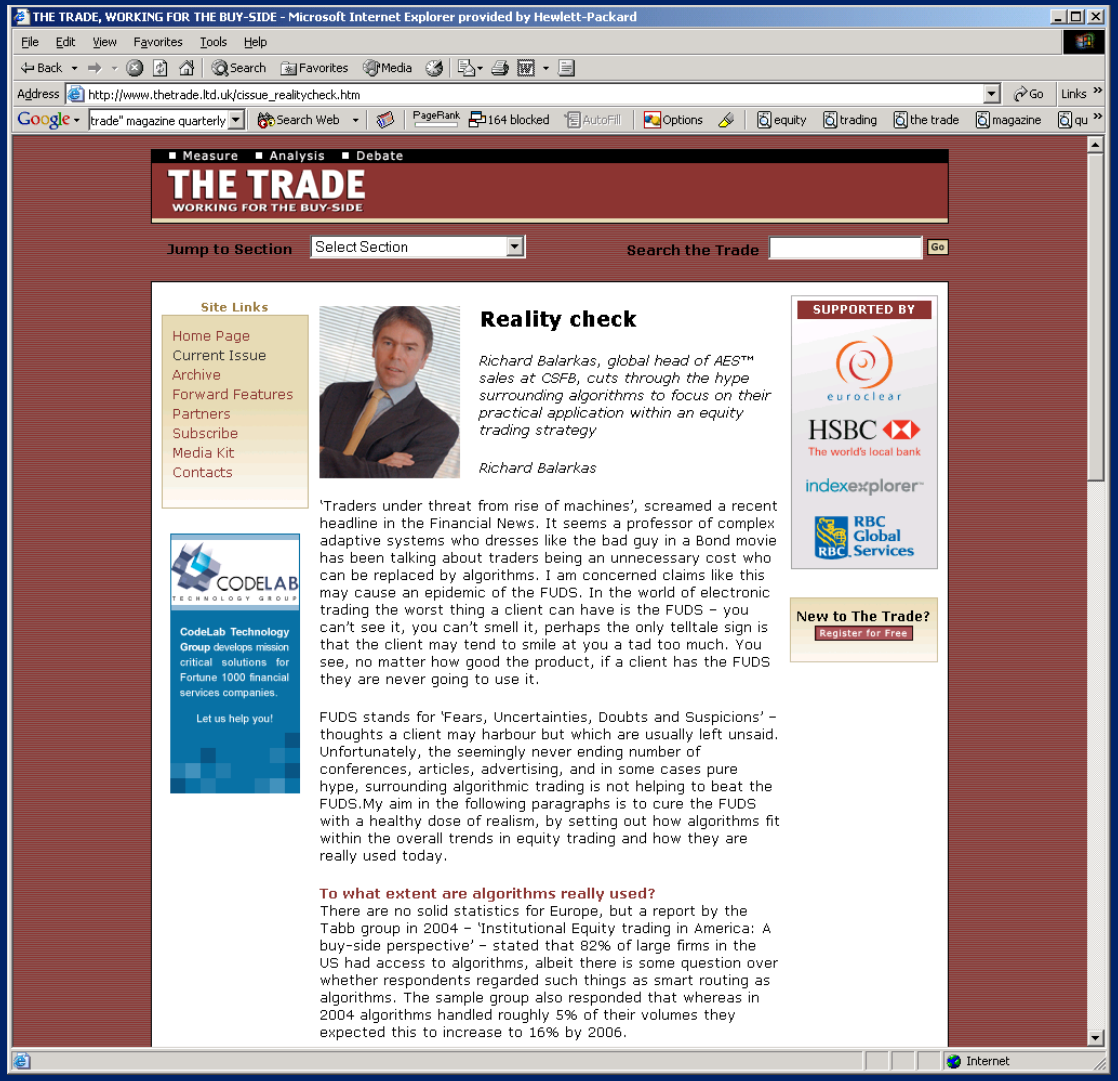

HP Labs & London Stock Exchange

- Alas, NDAs mean I can say very little about this project
- I can tell you that I led the HP side of it from the outset

- Five years ago it was abundantly clear that:
  - The current wave of M&A activity in the exchange-operator space was coming soon
  - Algorithmic Trading & Automated Execution increasing very rapidly
  - New technology (e.g. Net/Web/DMA/STP) lowering barriers to entry for disruptive technologies, and for disruptive new companies

- Technology innovations could give LSE an edge

- HP Labs team worked with LSE's nascent R&D unit
- HP: global revenues of $80bn p.a.; $4bn R&D spend; $200m of that on HPLabs; HP established in 1939 – centralised innovation established at HP Labs in 1966
- LSE: a UK-based SME; significant revenue slice from data provisioning; enjoys accidental privilege of national monopoly status; R&D unit a recent addition in its 200yr history.

And then I became a banker…

- Joined Deutsche Bank London FX Complex Risk Group
  - Wondered when I would be shown the secret technology
- Started out on a tour around the trading floor on "rotation", learning what everyone did and how it fitted together: Spot & Derivs; Prop & Client; Trading & Sales
- Coming in as a techie, with expertise in Computer Science & AI, thought I would be able to help out a lot of folk by coming up with whizzy innovative solutions to their problems, but…
  - Dissatisfaction/distrust of technology was much higher than I expected
  - A lot of the solutions needed did not need a PhD to solve
  - A lot of the solutions needed did not actually need *any* technology innovation
  - My impression is that the above three points are true:
    - At many investment banks, not just Deutsche
    - And in many asset classes, not just FX

Algorithmic Trading in Foreign Exchange

- FX is interestingly different from equities
  - No single central exchange pumping out price/vol data; still a lot of OTC
  - Massively liquid: $1,900bn/day, 24*7
  - Regulatory structure is pleasantly relaxed (verging on the non-existent)
  - Approx 2-3 years or more behind equities wrt uptake of algorithmics
  - Oh, and there's not much in the way of historical data. Really there isn't.

- Deutsche Bank London FX Trading Floor
  - biggest & best in Europe
  - 70% of DB London FX spot trades go through a single robot (>>£1bn/day)
  - aggressively developing more advanced technologies

  - My main job was to help improve the robot for high-frequency trading

What I did at DB (1): Price Discovery

- EBS screen (pre-ICAP)
- The whole curve

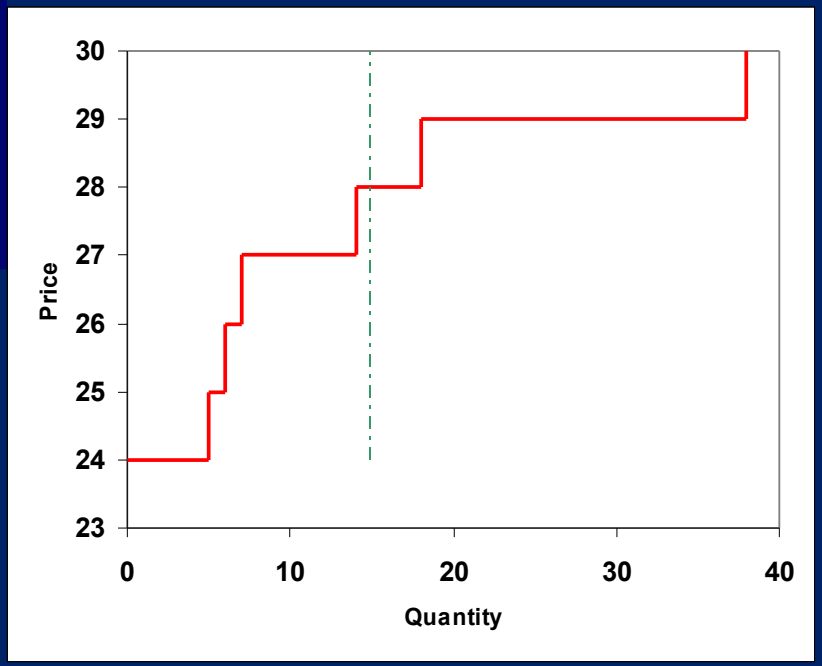

- "Laddered" prices

24	25	26	27	28	29
5	1	1	7	4	20

Really Reality: Reuters FX dealing

cf Price "Ladders" showing a list of ~180 prices

- Reuters D2000-2 spot USD/DEM supply and demand cu
  internal order-book data (hidden by the GUI) – usuall
- EBS must have this sort of data too, but their GUI
- At DB, I worked on ways to recover the hidden curves from what EBS does show

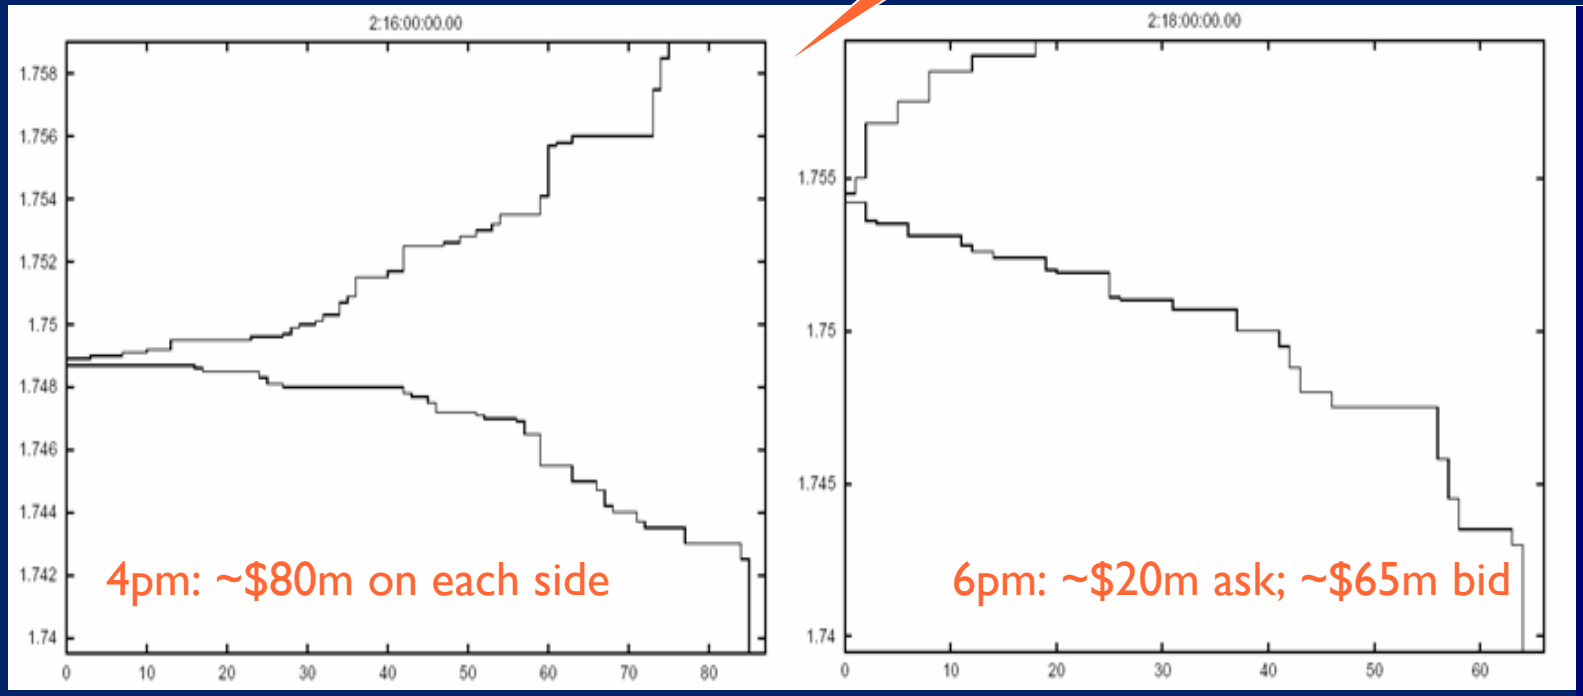

- J. Danielsson & R. Payne "Measuring & explaining liquidity on an electronic limit order book: evidence from Reuters D2000-2". Manuscript, London School of Economics (2003). www.riskresearch.org

What I did at DB (2): Execution Logic

- Pre-existing hand-coded "points" in strategy space
- Smooth interpolation between the points
  - generated stochastic "hybrid" strategies
- Not a million miles from what I had previously done on ZIP etc
- No pretty pictures, and no details here either ☺

What I did at DB (3): Visualisation

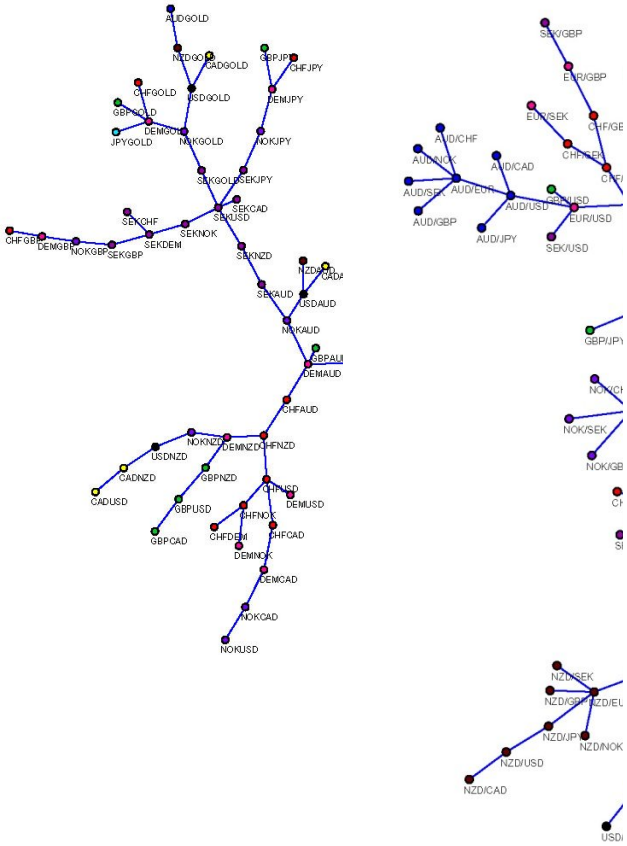

FIG. 4: (Color) The Min the correlations between the years 1993 and 1994.

FIG. 11: (Color) Currenc June 2004.

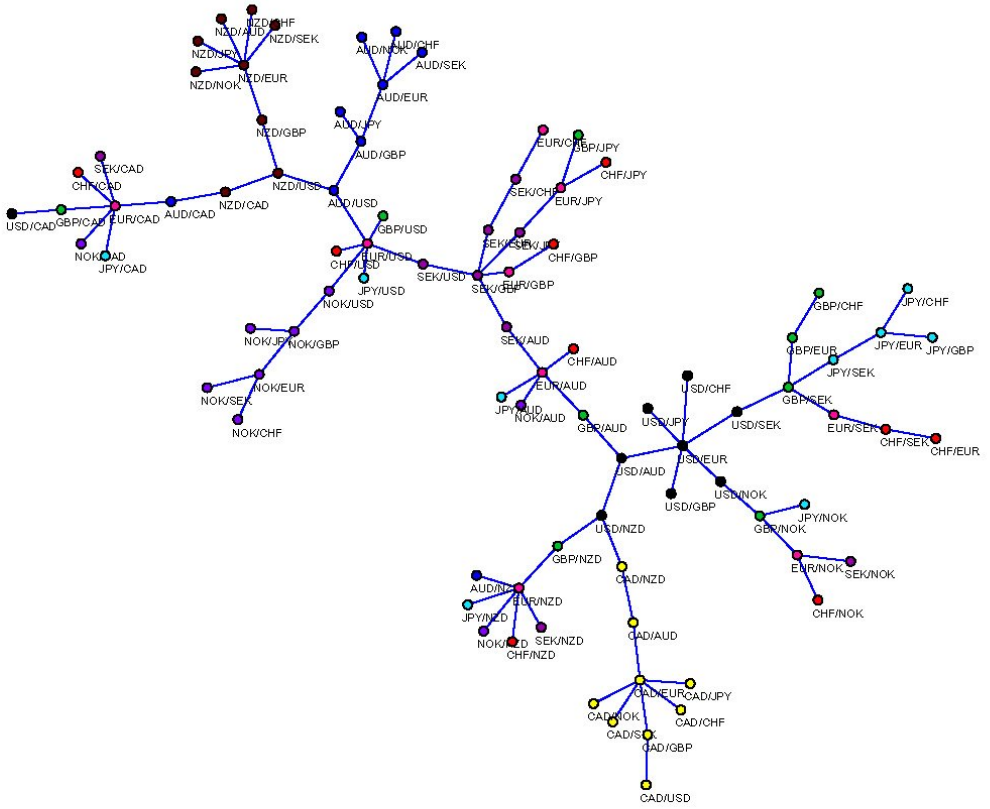

FIG. 12: (Color) Currency tree (MST) for a 2 week period in July 2004.

Where next?

- All three of these areas still hold some interesting challenges
- Will come to some future research challenges in a bit
- But first: some observations on technology innovation…

Innovation at HP & DB: spot the differences

- Research & Development Process
  - HP: Well established processes & checkpoints; woven into the company's DNA
    - Recruit skilled scientists and engineers; separate paths for managers & techies
    - Innovation workshops & training in creative/innovative thinking
    - Mapping of Intellectual Property "landscape" and of business sector "ecosystems"
    - Patent Strategy Review Panels staffed by senior technologists
    - Execution is an issue: mañana, mañana…
  - DB: Often ad-hoc, accidental, introspection-based, suck-it-and-see
    - Recruit immensely bright and quick-witted talent, from diverse educational backgrounds
    - Lack of formal innovation process(es) gives great speed and agility
    - Execution is almost immediate: Is it done yet? Is it working? How much has it made?

- Performance/Effectiveness Timescales & Measures
  - HP: 20-year patent lifetime; delay on feedback means intangibility is an issue
    - "10% success bankrolls the 90% of failures"; plus a specific success, squeezing ink onto paper
  - DB: 3-month P&L driven; conservatism is favoured: better out than down

Innovation at HP & DB: **why** the differences?

- HP has, from its outset, been a technology company
  - Makes its money from excelling in particular technologies
  - The specific technology HP excels in has changed from decade to decade over 60yrs
  - Hewlett invented HP's first patent; the patent made the money; that set the trend
  - HP employment contract lays claim to all IP of its employees, now and forevermore
- Corporate/Investment Banks have traditionally not been technology companies
  - Successful CIBs excel in customer relationship management
  - Successful CIBs excel in innovation too, but not necessarily innovating technology
    - An innovation might be a new financial product; a new business model
  - My DB contract non-compete clauses, stated that on departure:
    - No stealing DB's customers
    - No poaching DB's staff
    - Er, that's it
- But technology is now inter-woven and on the critical path for all CIBs
  - CIBs are fast becoming technology companies, partially at least
  - even if they don't realise it yet

Financial Times January 9th, 2007

INTELLECTUAL PROPERTY

Banks lay traps for copycats

The starting pistol has been fired on the race to register financial patents, write Saskia Scholtes and Gillian Tett

Finding the formula: the increasing complexity of financial instruments and trading systems has spurred growth of patents AFP/Getty

When bankers at UBS, the Swiss investment bank, recently invented a commodity investment scheme for devout Muslims, they were eager to start marketing their brainchild.

But first they took a step that once would have been almost unimaginable in the banking world: they asked American lawyers to file a patent to stop their idea being pinched.

"People haven't traditionally recognised the value in patenting these types of products," says Peter Ghavami, global head of commodities at UBS. "But these days the conventional wisdom is being challenged. We believe this patent should give us some protection."

Bankers have historically been highly aggressive in protecting their turf – and profits – but surprisingly bad at stopping rivals from copying ideas. While patents are a vital tool for, say, pharmaceutical companies, their use has not been widespread in finance.

As a result, whenever a bank produces a startling innovation, it is quickly copied elsewhere. When ABN Amro, for example, came up with a product called a "constant proportion debt obligation" last summer, other banks replicated the scheme within days.

In recent years, however, attitudes to intellectual property have undergone a subtle shift, as banks and other financial companies have become much more active in filing patents, particularly in the US.

In Europe, financial services or so-called "business method" patents are rare because they are granted only where they can be shown to have solved a specific technical problem, usually defined as non-financial or non-economic. And, say lawyers, the few financial services patents that do exist are seldom enforced in the courts.

But in the US, a 1998 ruling in a case involving State Street Bank made clear that business processes could be patented. A flood of applications followed and those are now feeding through into issued patents.

"This was the waving of the chequered flag. At that point, financial services companies started looking at every innovation to see if it was patentable and started hiring lawyers to build patent portfolios," says one intellectual property lawyer at a big investment bank.

The number of US patents issued for financial services soared in 2006, rising three-fold to 226. Many of these patents were issued to large US financial companies such as Citigroup, JPMorgan and Goldman Sachs. General Electric's finance arms headed the field with 11.

Edward Filardi, an intellectual property lawyer at Skadden, Arps, Wickersham & Taft, says the growth in patent applications for financial services is far from over.

A reason for this is that patents not only protect intellectual property, they can also be turned against companies (see below).

"Patent applications are snowballing and it's going to continue," Mr Filardi says. "Large financial institutions don't want to be in a position where a small company or individual can exert undue influence on their business."

Citigroup is one of more than 30 US banks being sued over patents that cover remote cheque image processing. Some banks, including JPMorgan, have settled with the company that was founded by the patent owner, Claudio Ballard, but most are fighting.

This is one of very few court legal battles in this area, but others could soon unfold over...

Eurex, the European exchange, plans to launch a credit futures contract later this spring. But the intellectual property for a technique that calculates a key component of this type of contract – known as a credit fixing – has been claimed by Credex. A US-based company, Credex is determined not to let rivals use it. This has forced Eurex to turn to alternative technology, complicating its plans.

Banks now view intellectual property as an important operational risk, say lawyers. Investment banks and other financial services companies are actively building portfolios of business process patents.

Some in the financial ...finance too?" asks a lawyer at a US bank.

Other trends are also raising the pressure to protect intellectual property rights.

One is the fact that the spread of technology has made it easier than ever for ideas to travel across the globe. In the 1970s it typically took 20 to 25 years for products to become ubiquitous in the market place. Now innovations diffuse "in less than five years", according to Mercer Oliver Wyman, the consultant.

At the same time, financial innovation is becoming more complex and thus more amenable to protection from patents. One way that banks are trying to gain a competitive edge on each other is in the use of mathematical algorithms. Another is in trading systems or software.

The Chicago Mercantile Exchange is now so determined to protect its intellectual property it has a full-time patent attorney on staff. "We have lodged several patents," says James Krause, chief information officer.

Glenn Manchester, head of Thunderhead, a small British company that develops software for credit deriva-

> 'If you look at chemical compounds, it is clear what they are – but financial ideas tend to be amorphous'

...tives trades, echoes the sentiment. "We definitely want to protect our intellectual property ... with patents," he says.

However, the efficacy of such techniques is relatively untested, not least because defining what is unique in financial innovation can be extremely hard.

"If you look at chemical compounds, it is clear what they are – but if you look at financial ideas, they tend to be amorphous," says Mr Das.

In a world where financial innovation can be spectacularly profitable – and where banks are desperate to retain control of the output of their scientific-sized staff – nobody expects the patent trend to end soon.

As one head of credit derivative product development at a major investment bank explains: "Filing a patent is a very low-cost option to protect something that you have spent time and

Beware the 'patent trolls' of finance

The growth of patenting in the financial services industry is reminiscent of the nuclear arms race: patents are very seldom used but provide protection from the patents of others.

companies to ransom for business methods that are sometimes widely accepted. In the pharmaceutical industry, those filing patents would typically support their application by

been extensively studied in the finance literature since the 1960s," says Prof Lerner.

In June, the company to which Mr Daughtery licensed the patents filed a

partner at Skadden, Arps, Wickersham & Taft, a US law firm. "The patent office has very limited resources in terms of searching publications and websites for existing business

Don't all rush at once

- Data from USPTO database of granted patents, June 2007, results of keyword searches:

- "algorithmic trading": 0
- "trading algorithm": 2
  - 6892186(HP), 5101353(Lattice); + 2 irrelevant
- "implementation shortfall": 1
  - 7110974(Lehman).
- "volume weighted average price": 6
  - 7228289(Trading Tech.), 7110974(Lehman), 6462758 (Reuters); 6098051+6012046+5845266(Optimark).
- "automated execution"+"financial": 6
  - 7209896+7181424(Nasdaq), 7085739(Accenture), 6112189+6098051+6012046(Optimark) + 19 irrelevant + 8 designs.

So why don't CIBs have Laboratories?

- Just because big companies like HP (and IBM, and MSFT, and Xerox, etc) have centralised research & development labs, doesn't mean that CIBs should too
- In fact, the idea of a central internal R&D Lab/function is starting to look distinctly last-century
- The new buzzwords on the Tech Innovation street:
  - Open Innovation
  - R&D replaced by A&D

- Attractively Darwinian
- R&D risk is borne by the investors in the start-ups
  - Not by the customers of the products
- Start-up exit strategy is acquisition by a gorilla
  - Not bubble-style IPO

A few examples from my time at Deutsche…

Flextrade FX Algo Platform

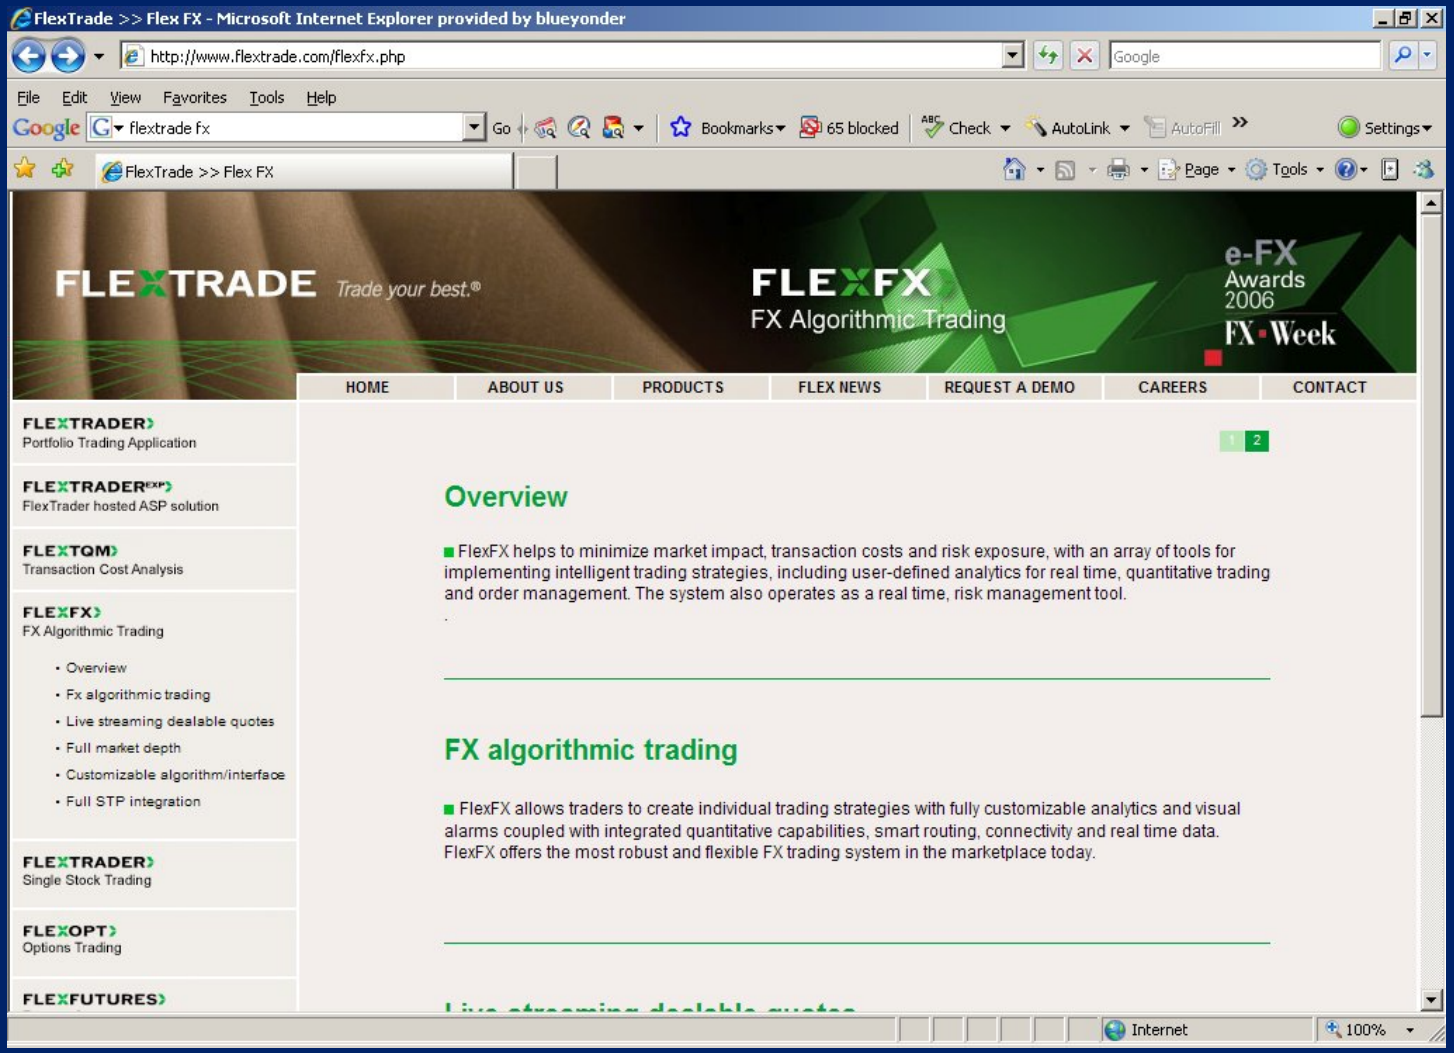

Codefarm Automated Credit/CDO Structuring

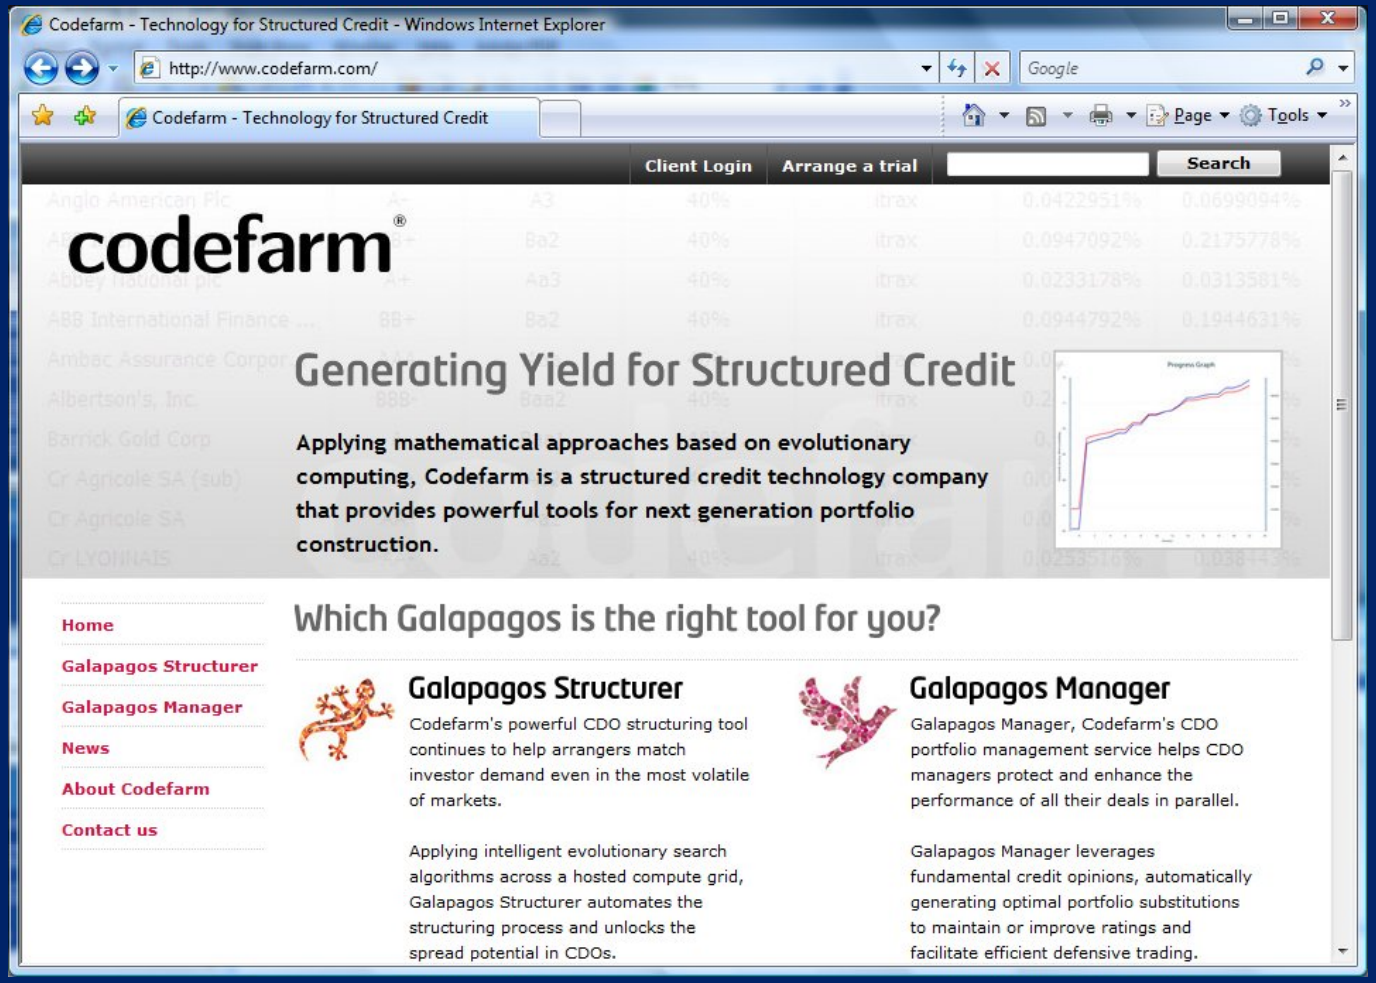

Complexity Science for real: www.eurobios.com

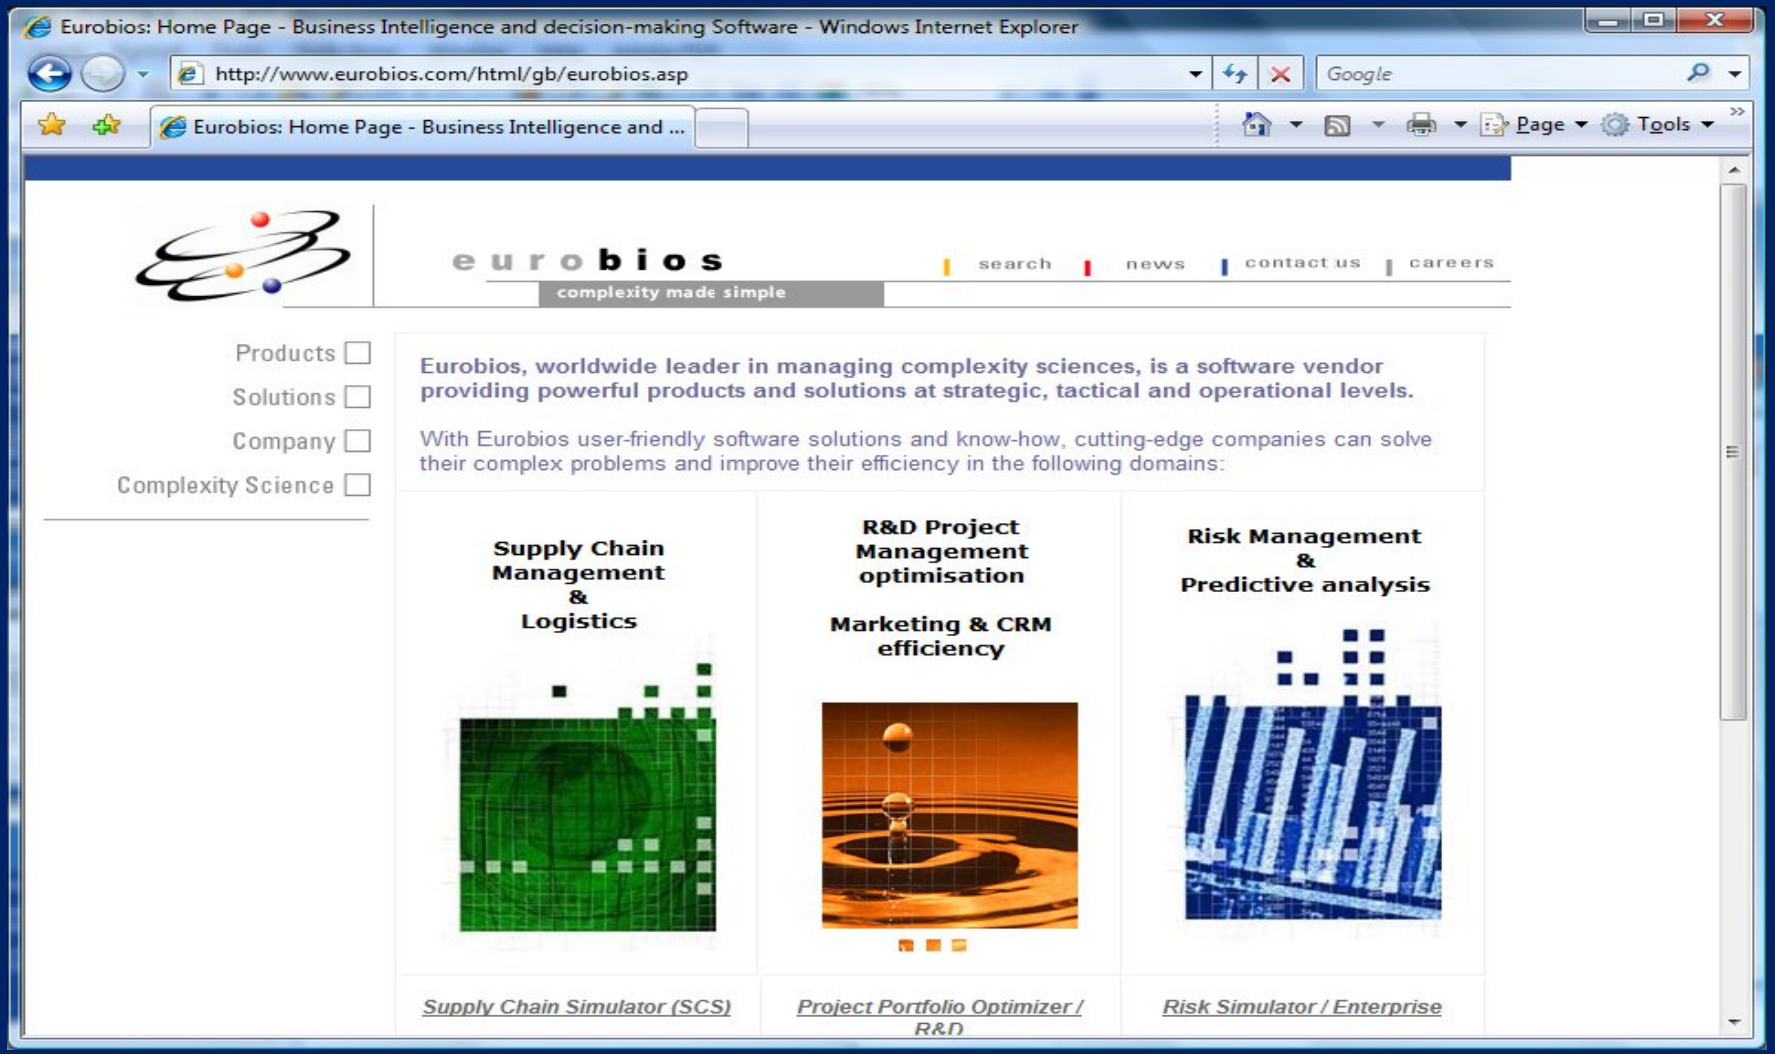

Oh, and Syritta Algorithmics too, one day…

The truth lies in the middle

- RoI for an industrial R&D Lab has time horizons way longer than are commonly acceptable in CIBs – that culture will not change in a hurry
- But the basic practice of an R&D Lab *can* be incorporated into a CIB community, whether cash or dreivs, trading or sales, or any cross-product too
- What I call the "P-I-P-E-R" loop: Predict – Invent – Protect – Exploit - Repeat
  - Anticipate/guess at likely future scenarios
  - Imagine the opportunities and problems in those scenarios
  - Plan to exploit the opportunities, and to avoid/ameliorate the problems
  - Protect the intellectual property/capital thereby generated
  - Exploit – get your RoI
- This takes practice/training, but it is way cheaper than setting up Central Labs
  - If you are practiced/trained, you can better evaluate the offerings of A&D innovators
- It's this practice of technology innovation that students should get educated in

Some predictions for the next few years

- "Semantic web" machine-readable semantic tags on ticker news-feeds and real-time Natural Language Processing on time-based media (voice channels, video feeds) combine to allow algo trades to be triggered/modified by text/voice/video data
- Algorithms structuring & executing cross-asset trades on basis of searching vast databases for high-order nonlinear correlations will become the de facto norm
- "Arms-race" co-adaptive dynamics in algorithm sophistication continues to drive up complexity of algorithms; so hand-design gives way to machine-design & optimization
- Barriers to entry continue to lower: plug-and-play piping & interfaces & algos & services
- Speed & Latency:
  - Co-location "proximity servers" give way to direct hosting on exchange servers
  - Algorithms move into silicon (FPGA, ASIC, etc) for nanosecond execution
  - Exchange API via a motherboard/backplane (…leading to stochastic polling?)
- Trader Interfaces: "ungameability" (impact-reducing obfuscation) is a hot topic
  - More sophistication needed to for X-asset "battlefield command" "head-up display"

When you have HPC/Grid/Utility capacity…

You can do this…

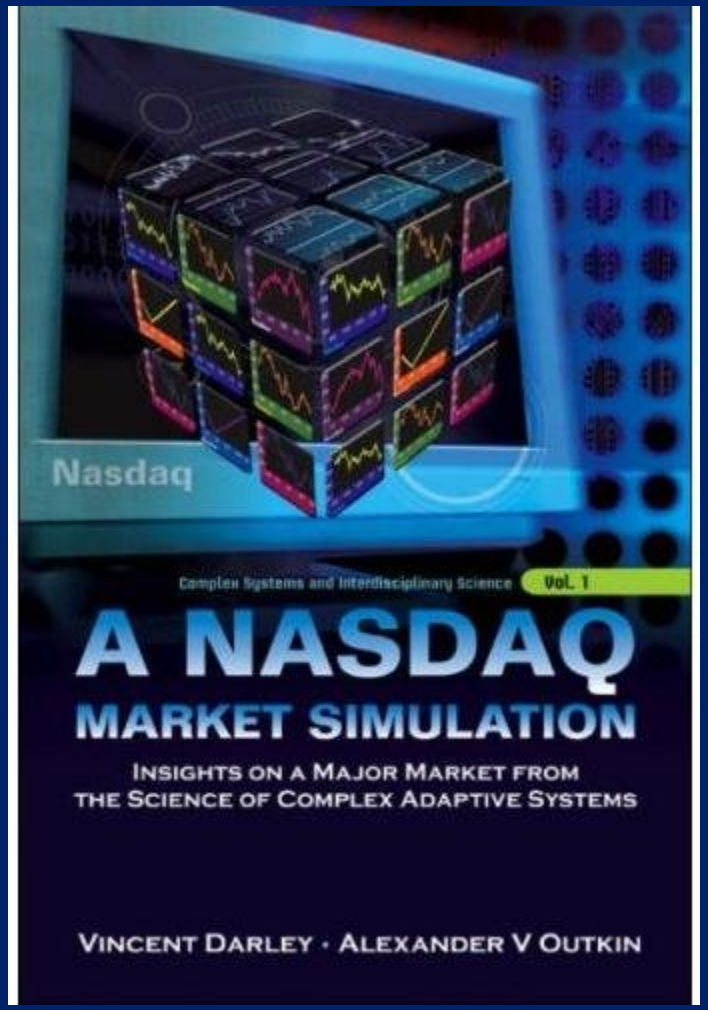

…and then wrap a GA around it?

Trader HUD/UUI

- Human traders will be around for several more years
- But the job could get much more complex
- The trader's human-computer interface is an area ripe for productive innovation
  - numeric tickers & spreadsheets not enough
- Lots of numeric data can instead be combined into dynamic graphic displays
  - with audio warnings
  - maybe with haptic input/feedback
- It's easy to get this sort of stuff very wrong
- Combat aircraft could hold some clues…

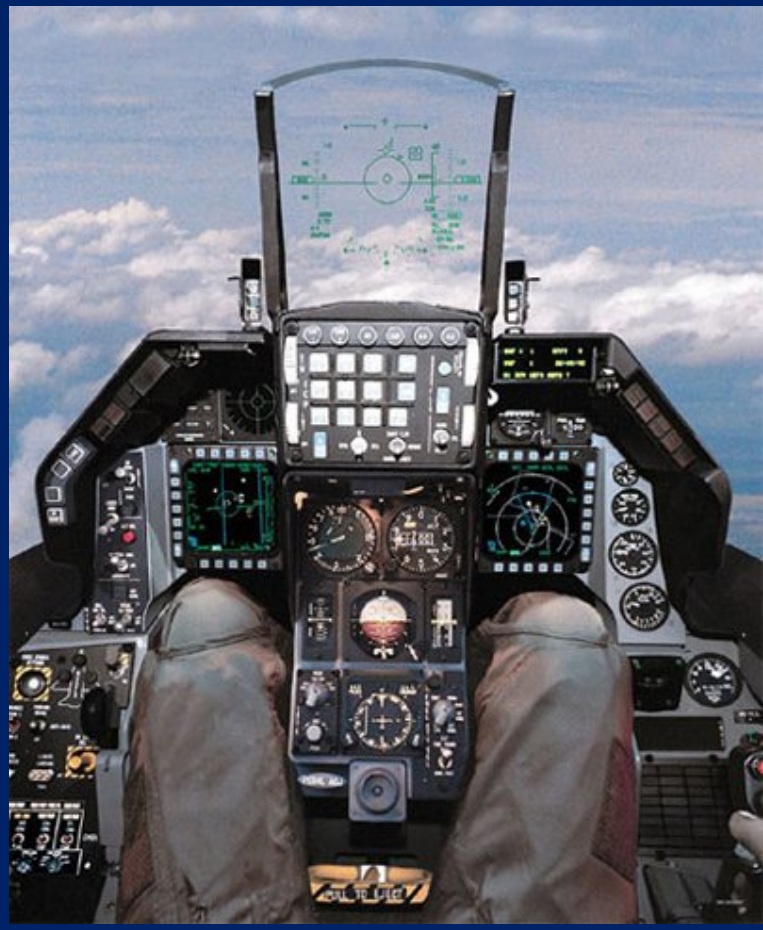

Trader HUD/UUI

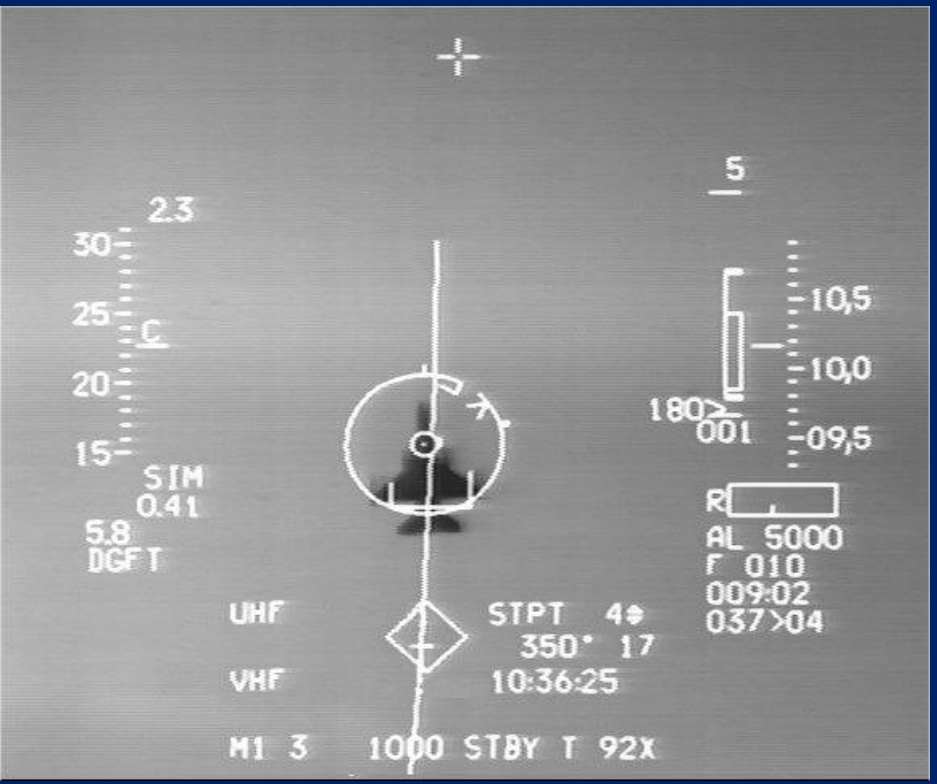

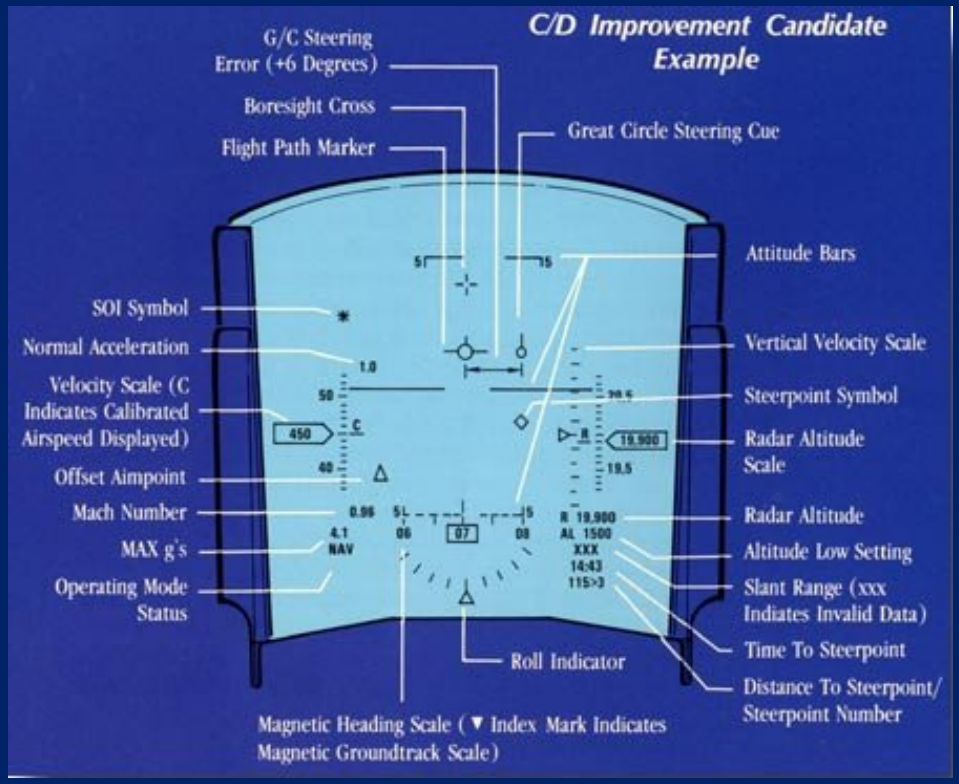

Good interfaces support good people…

…they don't replace them

(video from www.zawodny.com/blog/archives/006426.html)

New Scientist June 2nd, 2007

Where have all the traders gone?

Time is running out for Wall Street's high rollers. A new breed of traders is muscling in, says **Robert Matthews**

BRAD BAILEY was visiting the trading floor of an investment bank in New York City when he first noticed it. As a former Wall Street trader, he should have felt at home amid all the screens, phones and bustle of billions of dollars in trades. But that was just it: there wasn't any bustle. In fact, there were hardly any traders. "You could hear a pin drop," he recalls. Then it sank in: machines had taken over the role of people and computer servers don't make any noise.

There's a quiet revolution happening all over the financial world. Gone are the days of Gordon Gekko lookalikes screaming obscenities and dumping a loss-making stock onto an unsuspecting market. Investors have realised that the processing speed and sheer volume of trades a computer can make can help them to outwit the sharpest of dealers. As a result, they are investing heavily in what has fast become an arms race between investors. Their goal is to develop the best "algorithmic trading" systems – software that helps decide which trades are the most profitable, and then does the deals. Ten years ago, algo-trading was almost non-existent, but according to a recent report by Bailey, now at the Boston-based consulting firm Aite Group, one-third of all trading decisions in US markets are now made by machines. He predicts that by 2010 more than half will be done this way. At Deutsche Bank in London, over 70 per cent of a category of foreign currency trades, called "spot trades", are now carried out without human intervention every day. All this will have an

"70 per cent of foreign currency trades are carried out without human intervention"

impact on more than just high-rolling investors. Even if you don't own any shares you can bet that millions of those owned by your pension fund are already being bought and sold using "algo" trading techniques.

It's not hard to see why algorithmic trading is so attractive. Machines can make multiple trades, monitor thousands of stocks and do it all at breakneck speed. Crucially, they can do it without anyone noticing. There are big profits to be made in buying and selling shares that other traders haven't yet realised are being lucratively traded. The more discreetly you can do this – by spreading the deal over lots of small trades, for example – the less likely other traders are to wake up to the opportunity and dilute your profit potential. Such discretion is near impossible for a human, as it requires constant monitoring of the market to make sure your trades don't alter stock prices in an unfavourable direction.

As a result, investment houses are becoming increasingly tech-savvy. Eavesdrop on traders today and you are more likely to hear talk of "low latency access" (which we'll get to later) than of what they'd like to do to a rival's neck.

"Anyone who's been on Wall Street since the early 1990s will have had to reinvent themselves," says Bailey. Back then, success as a trader hinged on an instinct for what the market was "thinking", plus reactions fast enough to make the most of the opportunities the market presented. In the early years of computerised trading, when machines were simply communication tools, hitting a key a few tenths of a second faster than a competitor could make a real difference to your profit margin. Nowadays human traders struggle to keep up. "Silicon is taking over from carbon on Wall Street," says Bailey, as beige boxes proliferate across trading floors.

Computers have the edge over humans in many ways. Take something as simple as reaction time. When a human trader sees a stock change price, he might react in a few hundred milliseconds. A computerised trader is at least 10 times faster: depending on how much you are willing to fork out to speed things up. A few hundred milliseconds might seem insignificant, but if the price changes by a fraction of a per cent in the split-second before a trade worth many millions, it could mean a swing of tens of thousands of dollars. The key for any trader is "low latency" market access – that is, minimal delay between placing an order and seeing it fulfilled.

To achieve this, traders naturally have ultra-fast software, running on top-of-the-line computers with the very best processors and memory capacity. But there is a more direct, and perhaps less obvious, way to speed things up: moving closer to the source. Trading companies pay top dollar to smuggle their servers as close as possible to those of the stock exchange. With access to such "proximity servers", their electrons can beat those of their rivals to the punch. Last year one of the biggest hitters in the algo-based trading world, Deutsche Bank, paid an undisclosed sum to proximity-server supplier BT Radianz to shave milliseconds off its trading times.

Stealth-trading is another area in which machines have the upper hand. For example, many of the leading brokerage firms now have computers running so-called volume-weighted average price (VWAP) algorithms. These algos slice up big transactions, then ▸

Bloomberg Markets, June 2007

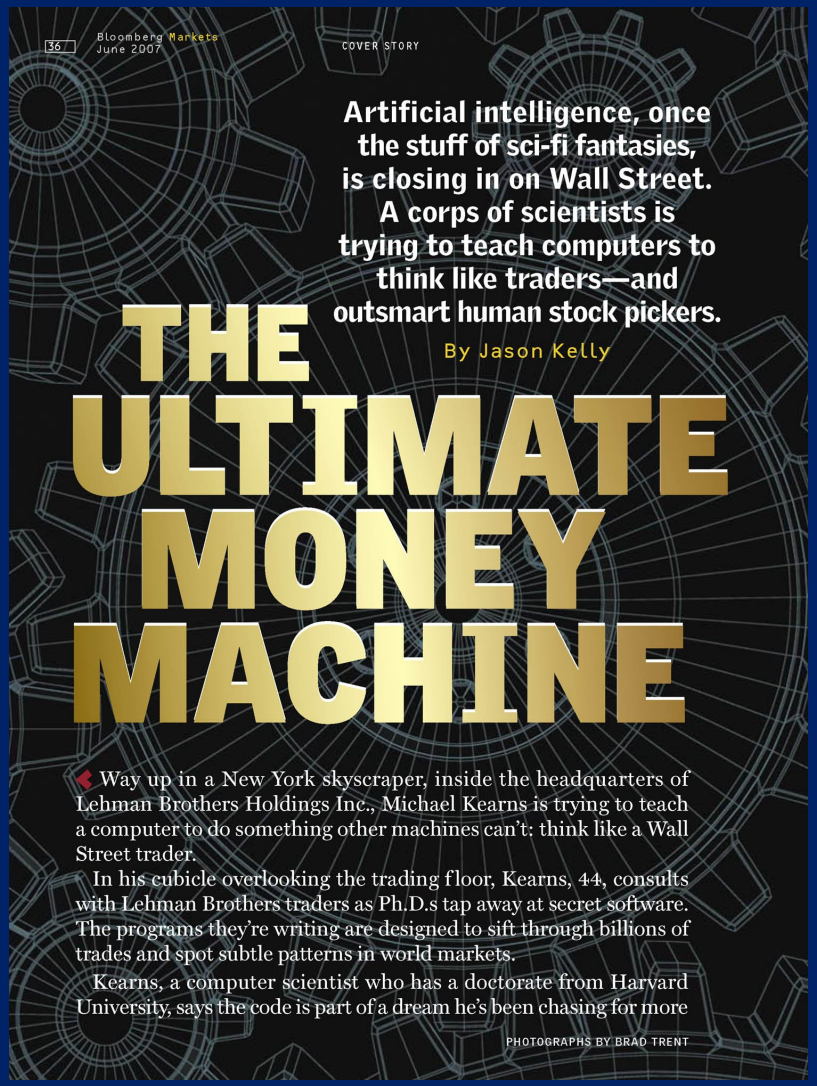

36 Bloomberg Markets June 2007 COVER STORY

Artificial intelligence, once the stuff of sci-fi fantasies, is closing in on Wall Street. A corps of scientists is trying to teach computers to think like traders—and outsmart human stock pickers.

By Jason Kelly

THE ULTIMATE MONEY MACHINE

Way up in a New York skyscraper, inside the headquarters of Lehman Brothers Holdings Inc., Michael Kearns is trying to teach a computer to do something other machines can't: think like a Wall Street trader.

In his cubicle overlooking the trading floor, Kearns, 44, consults with Lehman Brothers traders as Ph.D.s tap away at secret software. The programs they're writing are designed to sift through billions of trades and spot subtle patterns in world markets.

Kearns, a computer scientist who has a doctorate from Harvard University, says the code is part of a dream he's been chasing for more

PHOTOGRAPHS BY BRAD TRENT

Euromoney Algorithmic Trading, July 2007

My family and other algorithms

by Professor Dave Cliff, University of Bristol

IN THE SUMMER OF 1996, I INVENTED A TRADING ALGORITHM THAT OUTPERFORMS HUMAN TRADERS. THE ALGORITHM, KNOWN BY THE ACRONYM 'ZIP' (FOR ZERO-INTELLIGENCE-PLUS) IS EMBARRASSINGLY SIMPLE. BEING VERY SIMPLE MAKES ZIP VERY FAST, WHICH, NOWADAYS, WITH MAJOR PLAYERS RUNNING THEIR ALGORITHMIC TRADING OPERATIONS ON PROXIMITY-SERVERS HUDDLED AS CLOSE AS PHYSICALLY POSSIBLE TO THE MAIN SERVERS ON MAJOR EXCHANGES, IS A NOTABLE POINT IN ZIP'S FAVOUR. ANOTHER THING THAT MAKES ZIP ATTRACTIVE TO MANY PEOPLE IS THAT YOU CAN USE IT FOR FREE (AND I'VE BEEN TOLD THAT SOME MAJOR ALGO HOUSES DO EXACTLY THAT) FULL DETAILS OF THE ALGORITHM, INCLUDING THE 'C' SOURCE-CODE, WERE PUBLISHED ON THE WEB IN 1997 AND HAVE BEEN AVAILABLE THERE EVER SINCE.

5

In the 10 years since it was invented, other algorithms have come along to challenge ZIP; but over the same period I've been involved in the development of various extended versions of the original ZIP - there is now a whole family of ZIP algorithms, and the original has been re-christened 'ZIP8' (for reasons that will become clear later). This family of extended ZIPs demonstrate significant improvements in performance, with little or no increase in latency. The most recent of the ZIPs to be given away for free goes by the name of 'ZIP60'. I will come back to ZIP60 later.

ZIP – the basics

In 1996 I was working as a lecturer in computer science and artificial intelligence at the University of Sussex, and supplementing my meagre academic salary by day-trading equity options on LIFFE, the main London futures exchange. The good folk at Hewlett-Packard labs in Bristol offered me a seven month stint working with them as a visiting academic, and told me that I could come up with my own research project for my stay there. In those days, trading at LIFFE was still based heavily on face-to-face open-outcry trading

Professor Dave Cliff, Department of Computer Science, University of Bristol

tel: +44 (117) 954 5264 or +44 7977 552250

e-mail: dc@cs.bris.ac.uk or davec@syritta.com

New Scientist June 2nd, 2007

limit the amount bought and sold in each trade to a specified proportion of the total value of trades in those stocks on that day. The idea is to avoid drawing other traders' attention to what you are doing, and so prevent rivals getting in on the act.

Or at least that's how it used to be, until the Darwinian imperative kicked in. "At the start of the decade, if you used a simple slice-and-dice program, you got a big improvement in trading performance," says Richard Balarkas, an algorithmic-trading expert at Credit Suisse in London. Now those big improvements have evaporated. People started using monitoring algorithms to spot trades where, for example, someone was selling 5000 shares every 15 minutes, he says. "It's what we call 'signalling risk', and it has become a real issue. Other people can make money from working out what you're trying to do."

Arms race

As a result, the arms race has accelerated. Pattern-spotting software now looks for signs of algos trying to sneak their transactions onto the market, and studies the size of the transactions to make a guess at what's going on. That in turn has led to the development of algos which create smokescreens by combining their transaction strategies with a small amount of randomness in, say, the timing of the trades.

This particular arms race now seems to be heading towards some kind of stalemate, driven by the fact that there's a limit to how small and apparently random a transaction can be. "It's getting to the point where there is so much noise that it's going to be hard to detect market opportunities," says Balarkas. And that is pushing algo traders to find new ways of reaching their ultimate goal: "adding alpha".

Everyone has heard the stories of kids sticking pins in stock market listings and outperforming Wall Street's finest. There is a solemn truth behind such anecdotes: investment managers who try to beat the market with their cunning usually fail. According to David McCraw of Aberdeen Asset Management, based in Edinburgh, UK, two out of three investment managers secured worse returns than those achieved by tracker funds that automatically follow the overall market. In the jargon, these managers fail to add "alpha" – roughly speaking, the measure of extra return on investment achieved, compared to a tracker.

Being able to predict where the market is heading would be one way to beat the trackers – but that is easier said than done. "Any pension fund trader will have hundreds of trades to make, and there's an awful lot of information to assimilate," says Balarkas. "It's

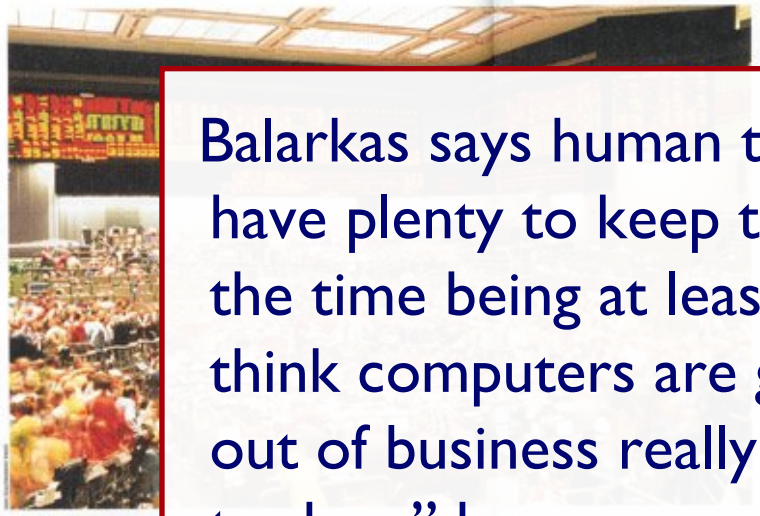

"Traders now ve… from the beate… of 'dark pools o…

really impossible to expect on… analyse it all equipped with ju… receptors, two aural receptors… Human traders are increasingly relying on computerised sidekicks to boil all the data down to a "buy" or "sell" punchline.

Another strategy for stealing a march on trackers is to look for signs of "liquidity" – the presence of buyers and sellers for particular stocks, which opens up a chance to trade. As before, there's scope for cunning. Trade too eagerly, and you could ruin the price, wait too long and the window of opportunity will slam shut as others beat you to it. This has led to the emergence of "sniper" algos, which wait for a suitable buyer or seller to emerge and then conduct the trade as fast as possible, before the price can be affected by other traders.

The obvious place to find such liquidity is in the big stock markets of London and New York. Looking in the obvious places, however, pools of liquidity". These are groups of buyers and sellers trading specific stocks in markets outside the mainstream exchanges. "If the big stock markets are the equivalent of regular stores, these dark pools are like eBay," says Bailey. "Around 15 per cent of US market liquidity is hidden away in them." Algos are being designed that go "fishing" in the dark pools, dangling small numbers of shares in the markets like bait. The speed with which they get snapped up indicates the liquidity level and the likelihood of profitable trading.

As with much else in algo-trading, implementing this simple idea is getting harder all the time as the competition between financial institutions for custom increases with Darwinian relentlessness. "As soon as one bank has a system significantly

THE RISE OF THE MACHINES

More and more investors on the US financial exchanges are relying on algorithmic trading software to tell them when to buy and sell their stocks and shares

University and founder of Syritta, a UK-based consultancy firm that develops algo-trading software. Banks are finally realising it's the only way to preserve their profit margins.

Some banks are trying to protect their algos using patent and copyright laws, but they can do little to hold back the real driving force for change: the continued battle for supremacy in which only the fittest algos survive. So instead of trying to deny the inevitable, some programmers are turning to evolutionary concepts for inspiration in designing new algos. Cliff is one of them. In the mid-1990s, he designed one of the first commercially successful algos. It didn't make the decision about when to buy or sell, but in experiments where human traders were pitted against his system it turned a bigger

"The pro… of the parameter… is that you're also throw… information," he says. Instead, … different parameter values and examin… them, looking for clues as to why some work better than others. On some days it turns out that only a handful of the 60 values are really needed to get the algo to work well. All this requires truckloads of computing power and investment banks just cannot get enough of it. Much of it is hired from companies offering metered access to giant computer clusters, but in many banks it's not unusual to see desktop PCs pressed into service, running trading analysis programs whenever they fall idle. For the company's most treasured algorithm processing, the number-crunching is done on

…efore unleashing the algo upon … it's rather like the fly-by-wire … by the latest fighter aircraft," … "The pilot can only instruct … form manoeuvres that the … will tolerate."

…man traders will still have plenty to keep … occupied, for the time being at least. "People who think computers are going to put them out of business really don't understand traders," he says. "We don't have algos that predict where the market is going – we have people, and they are still much better at it." The question is, for how much longer? ●

Robert Matthews is visiting reader in science at Aston University, Birmingham, UK

Balarkas says human traders will still have plenty to keep them occupied, for the time being at least. "People who think computers are going to put them out of business really don't understand traders," he says.

End of the road for human traders?

- Human traders are just humans; humans are just animals
- Animals are the result of 4 billion years of trial-and-error evolution;
- For 99.9999% of those 4bn yrs
  - humans and their ancestors were not evolving to become traders
  - Rather, they were evolving to satisfy the Four Fs
- For last 2000/200/20 years humans have tried to adapt to trading in markets
- But we humans are not purpose-built to trade…

  …we're limited in speed and in bandwidth; and trading just keeps getting harder/faster/higher-bandwidth
- ZIPs (and other "robot" traders) are purpose-built to trade; this gives them a distinct advantage
- No convincing argument in principle for why human traders could never be replaced by machine traders, but the technology isn't quite there yet

If you won't believe me, will you believe IBM?

"The trader is dead, long live the trader!"

- 20-page IBM Consulting (+EIU) report from www.ibm.com/services/fm2015

- "Power will shift from the traders who have benefited from merely facilitating transactions, to the buyers & sellers who take positions on either end of the trade…"
- "Transparency and speed are driving firms to develop a true client orientation and optimize risk/return efficiency, and are pushing them to become specialist enterprises – a task that will require a conscientious approach to innovation and significant modification of their operating models."

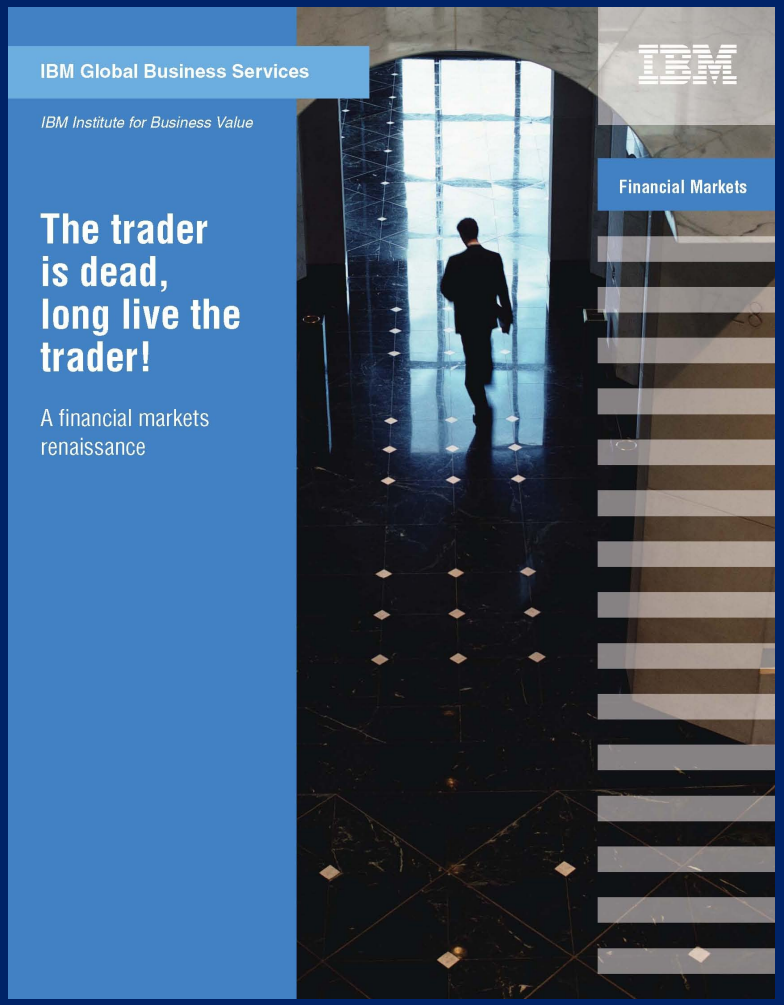

The Joker's Prerogative: question authority

"Traders who think computers ARENT going to hugely change their business really don't understand technology"

My day-job: www.lscits.org

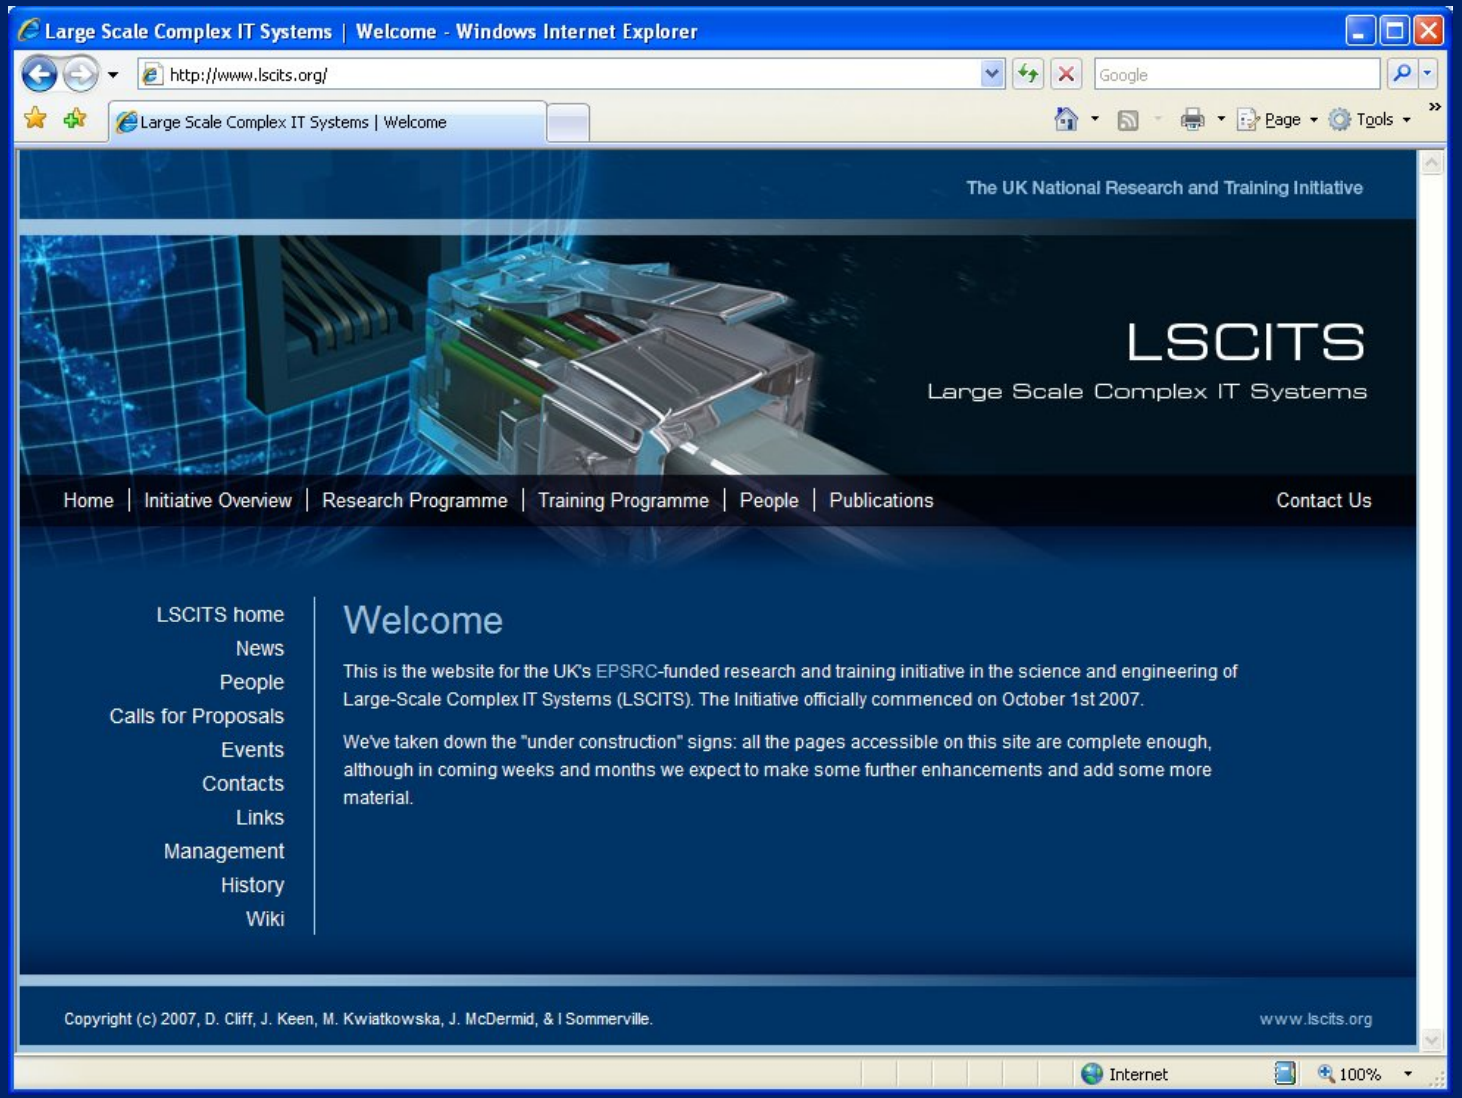

Summary

- Market-based systems are worth knowing about.
- Know you know everything that I know.

Thanks

- Collaborators, Colleagues, Competitors
  - At HP: Andrew Byde, Janet Bruten, Steve Gjerstad, Chris Preist, John Cartlidge
  - At MIT: Rod Brooks, Jake Beale, Won-Suk Chun, and the FABLab UROPs
  - At London Stock Exch: David Birch, Howard Miller, Laura Pandit, Tom Stenhouse
  - At Deutsche Bank: Tony Hall, Gio Pilliteri, Rhomaios Ram, Phil Wood.
  - At Southampton: Nick Jennings, Krishnen Vytelingum
  - The IBM guys

- At St Andrews:
  - Ian Sommerville for the invite & hosting
  - You lot, for listening. ☺