MARKET-BASED SYSTEMS

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

Lecture 1

Dave Cliff

Department of Computer Science
University of Bristol

dc@cs.bris.ac.uk

University of BRISTOL

The Advertising

- Market-Based Systems

  Over the last ten years, computer giants IBM and Hewlett-Packard have each invested significant research effort in developing algorithms that embody strategies for trading in "electronic marketplaces", and in algorithms that offer radical new types of electronic marketplace too. This industrial research has been paralleled internationally by a number of academic research groups with similar ambitions. Some of this research is motivated by the desire to create autonomous agents for e-commerce applications, some of it is aimed at doing better resource allocation and control in large-scale distributed data-centers and grid systems, and some of it is aimed at creating predictive models of real financial systems. As it happens, in the last few years there has been an explosion of interest in using such techniques in the global financial markets.

  These three lectures take a selective walk through the motivation, the background, the key results, the state of the art, and end with some wild hand-wavy speculations on where things will go next.

  Absolutely no previous knowledge of economics is required.

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.

The Bloke Giving The Lectures

- BSc in Computer Science; MA & PhD in Cognitive Science
  … I am not a trained economist

- Sussex Uni
- MIT Artificial Intelligence Lab
- Hewlett-Packard Labs
- Deutsche Bank FX Complex Risk Group
- Southampton Uni
- Bristol Uni

- Also some trading technology companies:
  - Syritta Algorithmics Ltd (www.syritta.com)
  - Duvacy Ltd

The Bottom Line for you lot at St Andrews

- Market economics looks set to become an important metaphor for new methods in computer science/engineering dealing with allocation of scarce resources,
  - so-called Market-Based Control (MBC) or M-B Resource Allocation (MBRA)
  - gives fast, robust, distributed control for dynamically varying problems
- Requires software versions of:
  - traders (quoting prices, haggling)
  - marketplaces (where buyers and sellers meet and interact)
- Designing traders and marketplaces is hard, but can be done automatically
- Mucking around with computer models/implementations of traders and market mechanisms can lead to innovations relevant to real-world economic/financial systems ($$$)

- Typical computer science undergrad/masters courses ignore all this
- There is enough material for a one-term course, but I've only written 3hr worth so far
- Follow the links/references in this presentation for more content

Natural systems can be useful inspiration

- The last twenty years has seen a number of computing techniques inspired by biology
  - Artificial Neural Networks (ANNs)
  - Genetic Algorithms (GAs)
  - Artificial Immune Systems (AIS)
  - etc (…Swarms/Colonies, Boolean Nets, Morphogenesis…)

- But ANNs, GAs, AISs, etc are all just examples of Complex Adaptive Systems
  - Systems which, at one level of analysis, are seen to be composed from many relatively simple components, interacting with each other in relatively simple ways, and yet for which at a higher level of analysis the behavior of the overall system is difficult or impossible to predict even given perfect knowledge of the individual components and their interactions… often due to the result of nonlinearities in the component interactions compounding over the entire system, giving rise to apparently emergent behaviors: path-dependencies and adaptation to environment in particular, often coupled with resilient self-organization/self-regulation/self-healing & graceful degradation wrt component failures

- But although most CAS are biological, biology is not the only field to study CAS…

Economies as CAS

… Economics studies CASs too: for example, markets possess:

- Many components (e.g. people, firms, buyers, sellers, speculators, arbitrageurs)
  - non-linear and/or unpredictable (can appear random)
- High connectivity & many nonlinear interactions between components
- Emergent system-level properties:
  - e.g. supply & demand automatically balance at equilibrium price (later…)

- This is not a new/recent observation…
  - P.Anderson & K.Arrow (eds.) *The Economy as an Evolving Complex System*, Addison-Wesley, 1988.
  - W.B.Arthur, S.Durlaf, & D. Lane (eds.) *The Economy as an Evolving Complex System II*, Addison-Wesley, 1997.

Economics: truly a dismal science

- Useful to distinguish between macroeconomics and microeconomics

- Macroeconomics is high-level impenetrable/mystical stuff, not needed here
  - e.g. the relationship between a nation's inflation rate and unemployment levels

- Microeconomics is just maths really, so much more useful for problem solving:
  - "Microeconomics is a set of theories with one aim: to help us gain an understanding of the process by which scarce resources are allocated among alternative uses, and of the role of prices and markets in this process." (Gravelle & Rees, *Economics*, 1992, p. 1)

  - Why should computer scientists care about allocation of scarce resources?

Dismal, but necessary

- There are lots of obviously important obviously resource-limited systems
  - e.g. telecoms, computer OS, logistics, staffing/timetabling, etc etc
  - automating their control/regulation is attractive for all the usual reasons…

  … and is difficult for all the usual reasons too
- So ideas from microeconomics could help us do automated dynamic resource allocation and control in engineered systems:
  - so-called Market-Based Control (MBC)
  - AKA Market-Based Resource Allocation (MBRA)
- MBC is not a radical new idea; it was initially explored in:
  - B. Huberman (ed): *The Ecology of Computation*. North-Holland, 1988;
  - S. Clearwater (ed): *Market-Based Control*. World Scientific, 1995.
  - cf also "contract nets"

Giving something back to microeconomics?

- In addition to borrowing ideas from microeconomics for work in PPSN, tools and techniques common in PPSN can be applied to problems in microeconomics
  - Modelling existing/potential economic systems: agent-based computational economics
  - Developing new styles of agent-enabled e-commerce and online trading systems
- Designing traders & markets for market-based systems is often v. hard: more of art than science
- Automated design and optimization techniques, such as genetic algorithms and genetic programming, can be applied to help design/optimize trading algorithms and the associated market mechanisms (…covered later, in Lecture 3)
- Furthermore the trading algorithms being optimized may use machine learning techniques to be adaptive in some way on shorter timescales (Lecture 2)
- It's possible that interaction/intersection of Comp.Sci. optimization and adaptation/learning techniques and microeconomics will lead to innovations applicable beyond market-based control, i.e. in real-world economic/financial systems
  - cf. interplay between computational neuroscience <=> artificial neural networks

Not just of academic interest…

- Major industrial research labs have pioneered and continue to explore the use of market-based systems for distributed dynamic resource allocation and control
  - BAESystems; BTExact; HP Labs; IBM Labs; Xerox PARC; etc.
- The result of applying evolutionary computation techniques to market-based systems has generated apparently novel microeconomic mechanisms; attracted positive press coverage (e.g. in The Economist); and from that came the attention and interaction of major companies in the financial services industry
- Summer 2004: UK EPSRC announce a major (£1.6m) grant to an academic/industrial MBC research consortium:
  - University of Birmingham (Yao et al);
  - University of Liverpool (Phelps et al);
  - University of Southampton (Jennings et al);
  - BAESystems (Figuereido et al);
  - BTExact (Marrow et al);
  - HP Labs (Cliff et al);
  - IBM Labs (Khepart et al).

It's a once-a-decade kind of thing

- Mainframes
- Minicomputers
- Micros/PCs
- LAN/Distributed
- Internet/Web
- Utility/Service...

Example: Markets in the Utility Data Centre…

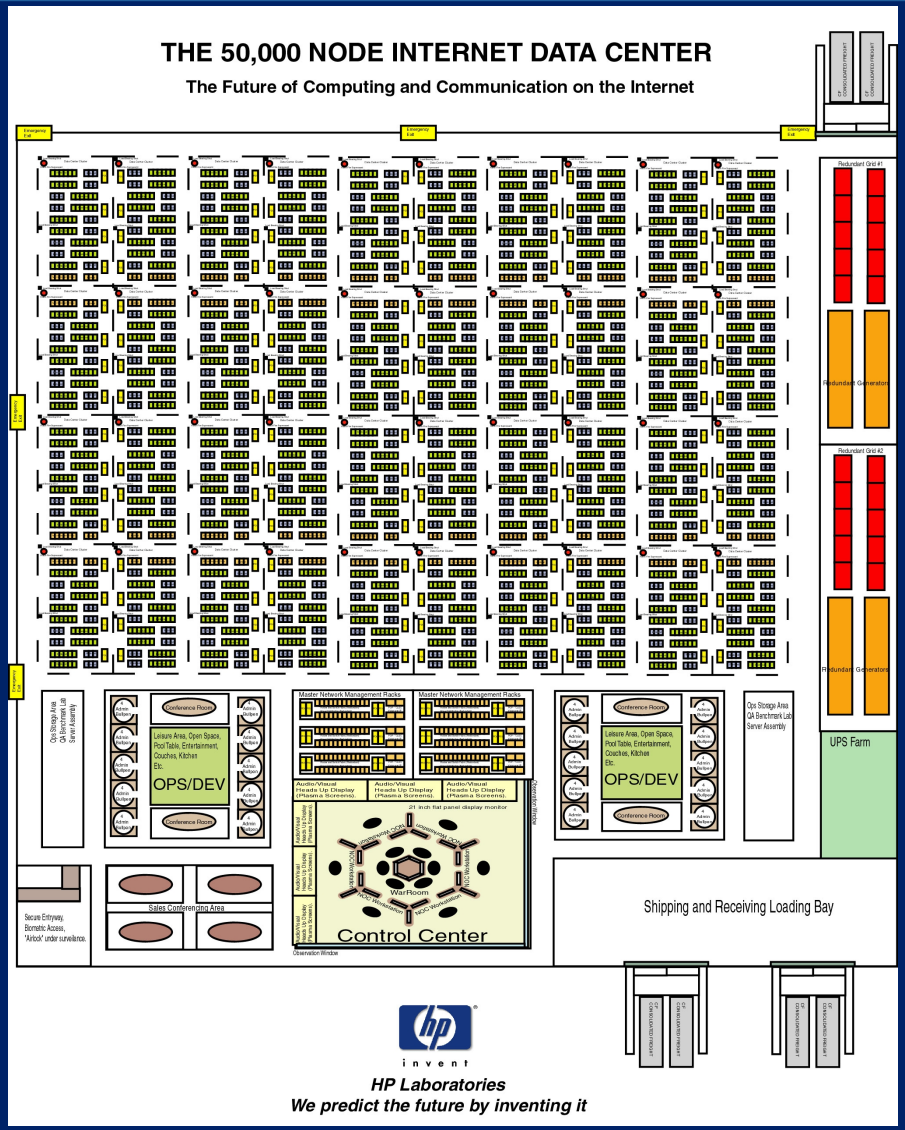

50,000 blade-servers

in a very big shed

with a very big air-conditioning system

Utility Data Center (UDC)

Just one shed in a global network…

If you can't stand the heat…

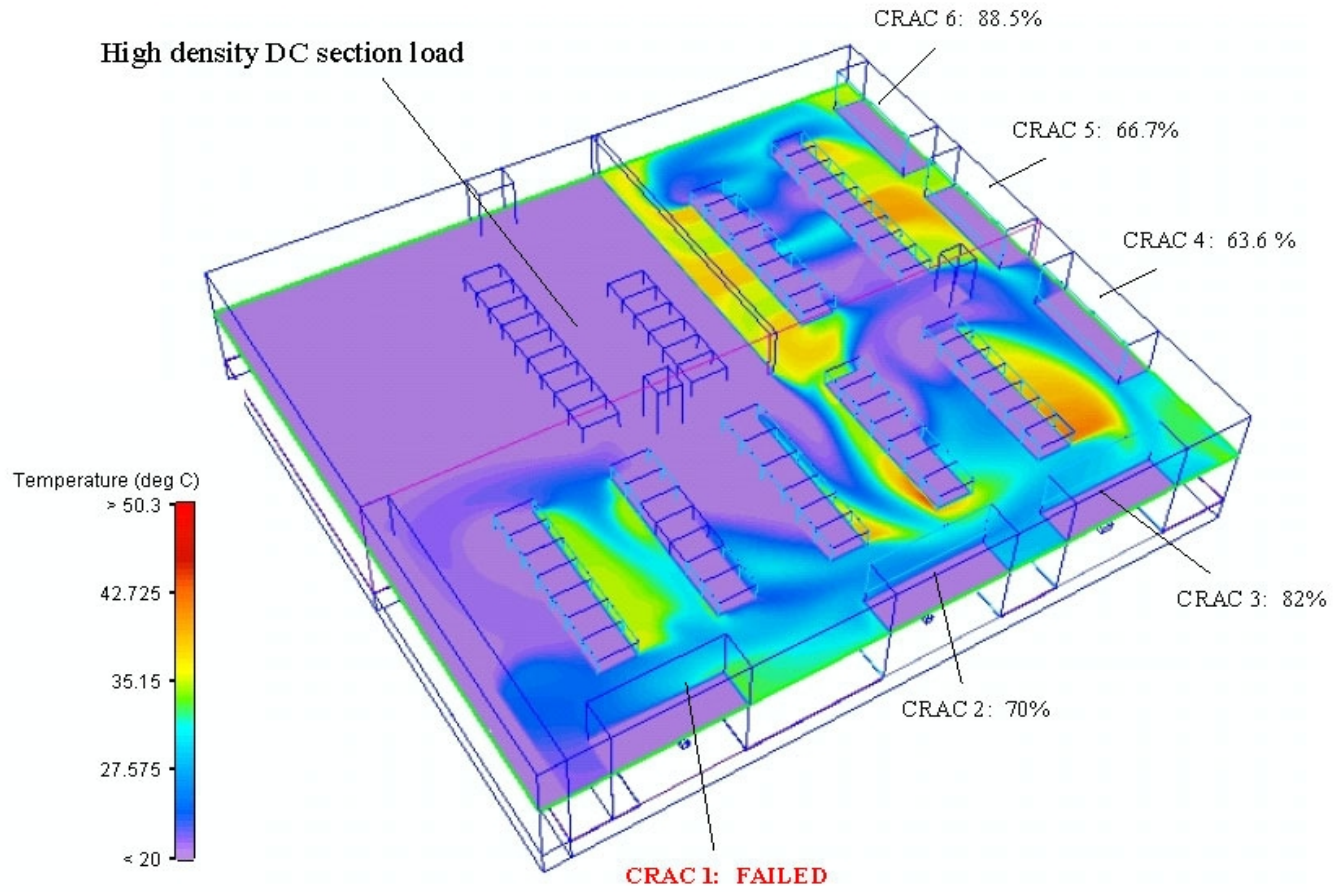

FIG. 5. Temperature Plane Plot after CRAC unit failure and load redistributed.

146.

%

Modular Data Centres In Big Steel Boxes

Google nabs patent for Sun's Project Blackbox? | The ...

Biting the hand that feeds IT

Cash 'n' Carrion Events Jobs

The Register®

Hardware | Software | Music & Media | Comms | Security | Management | Science | Odds & Sods

PCs & Chips | Servers | Storage | Data Networking | Semi-Coherent Computing | Virtualization

The Register » Hardware »

Google nabs patent for Sun's Project Blackbox?

'I own mobile data centers'

By Cade Metz in San Francisco → More by this author

Published Wednesday 10th October 2007 08:12 GMT

Download free whitepaper - Effective collaboration for SMEs

Google now owns a patent for data centers stuffed into shipping containers. You know, data centers like Sun's Project Blackbox.

The Mountain View outfit first filed for this "Modular Data Center" patent in December 2003, and today it was rubber stamped by the US Patent Office.

The patent describes a data center based on an "intermodel shipping container". That's a shipping container that's shipped by multiple carriers. "Rack or shelf mount computing systems," the patent explains, are "mounted within

Search
Reg Hardware
Reg Developer
Channel Register
Whitepapers
Reg Jobs
News Tools
Newsletters & Feeds
Reg Mobile
Reg Desktop News Alerts
Reg Shops
Reg Merchandise
Reg Books
Mobile Gadgets
Hosting
Top Stories
El Reg | Hardware
Remembering the IBM PC
Google builds very own

It's about scalability, not mobility

And MSFT too (*FT* Feb 25, 2008)

Microsoft predicts rise of the datacentre

Consolidation to blur distinctions

Central processing to replace desktops

By Richard Waters in San Francisco

A handful of US companies is set to dominate the emerging market for "cloud computing", which will assume a central role in the information technology world, according to Steve Ballmer, chief executive of Microsoft.

His prediction – in an interview with the FT – points to a future in which the distinctions between software, hardware and internet companies fall away and a small number of big technology companies, each running a collection of vast datacentres, controls the IT landscape.

Cloud computing involves the centralised storage and processing of information – a shift that could reduce the role of desktop computers and the servers and other equipment run by many companies. This is forcing companies such as Microsoft to rethink their strategies.

Mr Ballmer predicted that a new super-group of tech companies would dominate the cloud computing market, each of them managing what amounts to a giant centralised computer made up of a number of big datacentres.

"Amazon has one. Rumours are Google will have one. We've said we're going to have one," Mr Ballmer said.

The predicted emergence of super-powers in corporate computing points to a higher concentration of influence in the technology world. The winners would have the resources to operate at large scale, bringing down unit costs for computing to levels others cannot match.

Internet retailer Amazon was the first to start renting out capacity in its datacentres, selling storage and charging companies based on the number of transactions it processes. While Google has started offering online applications to companies, it has yet to start selling computing capacity in this way.

Mr Ballmer was speaking before the launch this week of one of Microsoft's most important new products this year, its latest software for servers known as Windows Server 2008. The software marks Microsoft's entry into the market for "virtualisation" technology.

The launch represents Microsoft's attempt to catch up with VMware, a software rival that has dominated the first phase of virtualisation. However, a key piece of Microsoft's own virtualisation software will not be ready. The software company has promised to deliver the piece, known as a hypervisor, by August.

See Editorial Comment

Empire faces a next-generation adversary

FT INTERVIEW

Steve Ballmer is priming his forces to repel rival VMware in the battle to dominate virtualisation, says **Richard Waters**

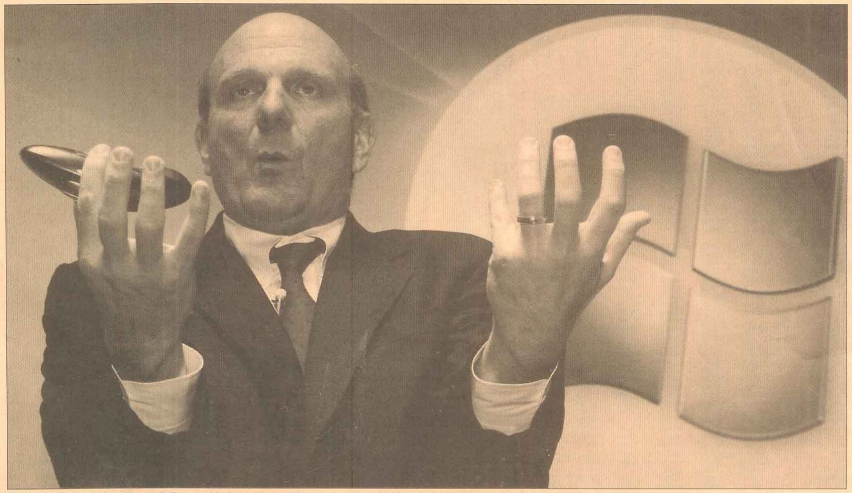

Revolution is coming: Steve Ballmer, chief executive of Microsoft, has said that there is no stopping the arrival of virtualisation Bloomberg

Microsoft will this week fire the opening shot in a battle set to be every bit as important to its long-term future as its internet showdown with Google.

While the search war will help determine who rules the consumer web, this one will decide who reigns supreme in corporate datacentres and on office desktops – the core of Microsoft's existing business.

At stake, says Steve Ballmer, chief executive, is nothing less than "the way computers get managed and deployed". That shift, he tells the Financial Times, "is a bigger opportunity for us than anyone – though, any time you have a major shift in this industry, it can be a threat if you don't move fast enough".

Others put it more bluntly. According to some, the big changes under way in corporate computing could lead to a new technology architecture where Microsoft's traditional source of dominance – its position in desktop and server operating systems – carries far less weight.

"It's incredibly dangerous to Microsoft longer term," says Michael Cherry, an analyst at Directions on Microsoft, an independent research firm.

The initial battleground is a technology known as virtualisation. This involves separating computing tasks from the machines on which they run. Rather than having a single application running on a single operating system on each server, virtualisation makes it possible to run several applications and operating systems simultaneously – each is known as a "virtual machine". Those virtual machines can then be moved between actual, real world servers.

That creates a new efficiency since most corporate servers are estimated to run at only 15-20 per cent of capacity. It also greatly improves the manageability of computing tasks since they can easily be backed up or moved to other machines when hardware fails.

For a company that has made its money selling single instances of Windows each time a new computer is installed, this represents a potentially significant change in business practice.

So far, Microsoft has responded by dropping the price for multiple instances of Windows running on single machines. As this new way of running computers continues to spread, it will have negative as well as positive impact on Microsoft's business model, says Mr Ballmer although he adds: "You can't fight City Hall – virtualisation is going to happen."

Yet the challenge to Microsoft extends beyond the adjustments needed to the way it charges for its products. At the heart of virtualisation is a piece of software known as a hypervisor, which sits between the virtual machines and the server, making it possible for the multiple computing tasks to all draw on the resources of the same piece of hardware. That new base level of software could assume a more central role, taking on more of the functions currently undertaken by operating systems.

In much the same way, Microsoft's own operating system began as a small "microkernel" before expanding to become a dominant computing platform, says Mr Cherry. It is this potential that represents the biggest long-term challenge to Microsoft, he adds.

"Any piece of software can morph into any other piece of software," concedes Mr Ballmer, though he largely brushes off the threat.

As in the fight with Google, this latest battle also involves an adversary whose early dominance of a new market has been turning heads in Silicon Valley. VMware was already gaining a reputation in corporate datacentres before its red-hot initial public offering last summer brought it to wider attention. The possibility that it could one day assume a central role in the new computing architecture has made it the world's fourth most valuable independent software company.

Also echoing its fight with Google, Microsoft starts well behind as it takes on VMware – although in this case it stands a far better chance of catching up. Microsoft's response comes with the launch this week of Windows Server 2008. The software company has two big advantages in its favour as it squares up against VMware. One is the ability to bundle the software with the new Windows server so it gets wide distribution.

"It goes in Windows, so it makes it hard for the IT guys to pay for a third party," says Mr Cherry.

Microsoft will add $40 to the price of the new server software to include the virtualisation feature. All in, the cost of Microsoft's software will be a third or less than that of VMware, says Mr Ballmer.

A second advantage for Microsoft is its existing position in corporate datacentres. Thanks to the market share that Windows has already built up on servers, its software tools are used by many IT staffs to manage their real-world machines. Now, those tools will also be able to manage the "virtual machines" that run in virtualised servers. Though VMware dominates the virtual world, it cannot match those real-world capabilities.

Yet with fewer than 10 per cent of corporate servers estimated to have been virtualised and with VMware moving fast to build on its early lead, the race is still at a very early stage.

MBC4UDC: market-based control for UDCs

- Distributed heterogeneous computer systems such as UDCs require controllers to implement allocation policies for scarce resources.
- Centralized controllers are vulnerable & scale poorly: we want decentralized control
- Resources include: CPU time; RAM space; hard disk space; network bandwidth
- These resources are…
  - demanded by users/jobs (the demand varies dynamically & is unpredictable)
  - supplied by the UDC computer nodes (supply varies dynamically & is unpredictable)
- So…
  - attach buyers to jobs and give them some money to buy resources from sellers
  - attach sellers to resources and let them compete to sell resources to buyers
  - when demand exceeds supply, price will rise (seller's market)
  - when supply exceeds demand, price will fall (buyer's market)

Markets in the Utility Data Centre…

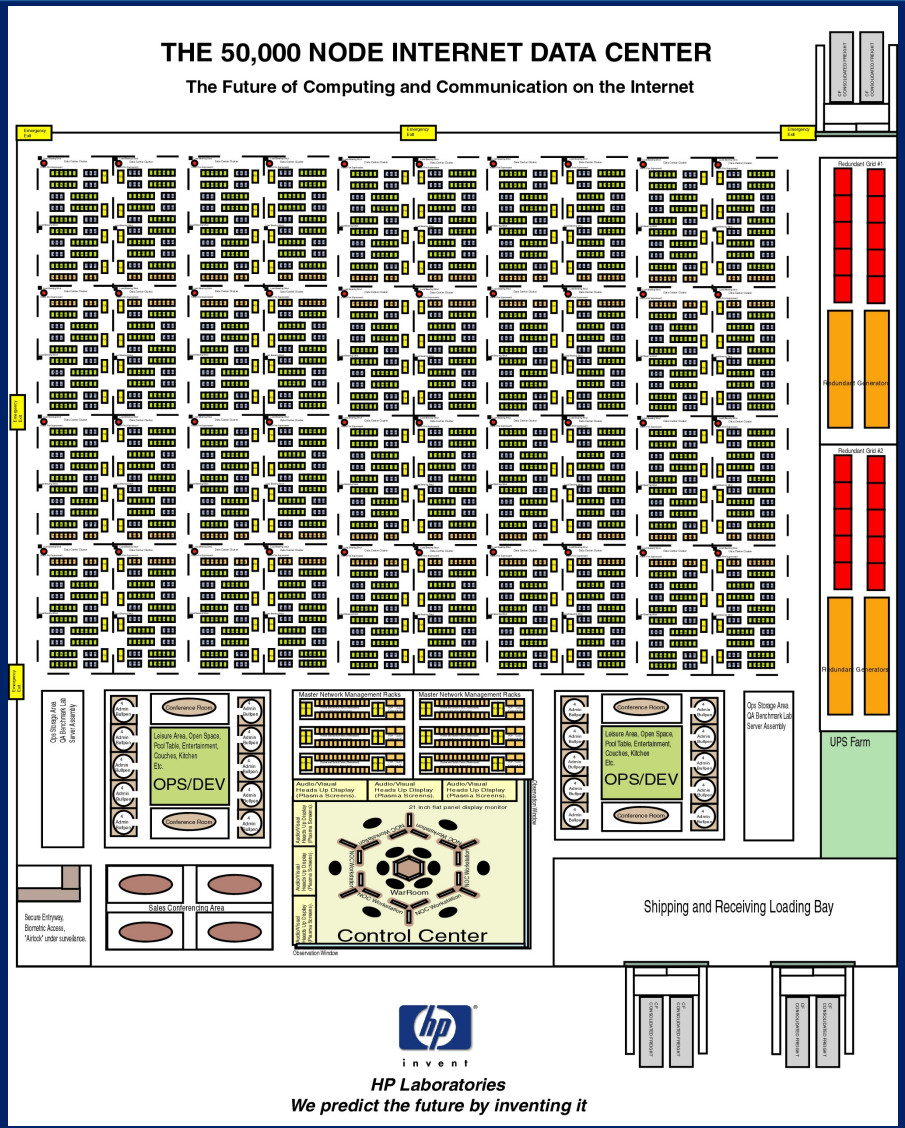

market-based control for UDC dynamic resource allocation

balance shifting supply and demand using ideas from free-market microeconomics

use automated optimisation to fine-tune the trader strategies and the auction-market mechanisms

trade, arbitrage, & hedge in a federated global commodity network…

Medium-term strategic relevance

- Medium-to-long term utility computing vision includes multi-vendor provision of commodity compute resources on a utility charging basis

  … a global federated network of utility fabric

- an international market in utility compute resources?
  - Spot markets for near-instant access to raw resources
  - Forward contracts and derivative markets to hedge risk
  - Any market volatility will attract wave-riders & day-traders
  - Market data (prices, volumes, etc) will be valuable info
  - Market regulations may vary across countries
  - National security issues too

http://www.economist.com/displaystory.cfm?story_id=2352183

  - Will commodity producers like HP, IBM, or Sun diversify to become compute-utility market-makers/specialists or international exchange operators?

…or will existing international financial institutions move in on this new global commodity market? …or will an IT/finance collaboration work best?

MBC in future computing: Grid

- A single utility data center (UDC) might internally resemble a computational grid
- UDCs can also be nodes on computational grids (cf. global networks of UDCs)

- Much "grid" research assumes altruistic caring/sharing resource allocation policies
  - makes minimum/peak quality-of-service guarantees hard to deliver

- Many grid researchers assume that omnipotent centralized resource planners/controllers will suffice
  - These people are basically advocating the construction of centralized command economies
  - They should remind themselves what happened to the Eastern Bloc at the end of the 20th Century

- Others seem to think that local control by the applications will suffice
  - Known to have stability problems as applications adapt to compete for resources

- Increasing numbers of grid researchers now turning their attention to MBC/MBRA

On the trading floors…

- It happens that these automatically-optimized robot traders and market mechanisms, originally intended for control of data centers, are of significant interest to major investment banks, buyside funds, and exchange operators

D.Hart-Davis/DHD Photo Gallery (http://gallery.hd.org)

Background Economics #1

- From everyday experience, we're used to the idea that if the demand for some tradable item exceeds the supply, then the price for that item will rise… colloquially we call this a shortage, or a sellers' market
- Similarly if the quantity supplied is greater than the quantity demanded, then the price will fall… a surplus, or a buyers' market
- So we're familiar with the idea of the price of a tradable item being a function of the quantity demanded or supplied

- Formally, the price that buyers are prepared to pay at each possible quantity is referred to as the demand: plotted on a graph of price *vs* quantity as a demand curve
  - usually slopes downward (as prices increase, quantity demanded decreases)
- …and the price that sellers are prepared to sell for at each possible quantity is referred to as the supply: plotted on a graph of price *vs* quantity as a supply curve
  - usually slopes upward (as prices increase, quantity supplied increases)

Background Economics #2a

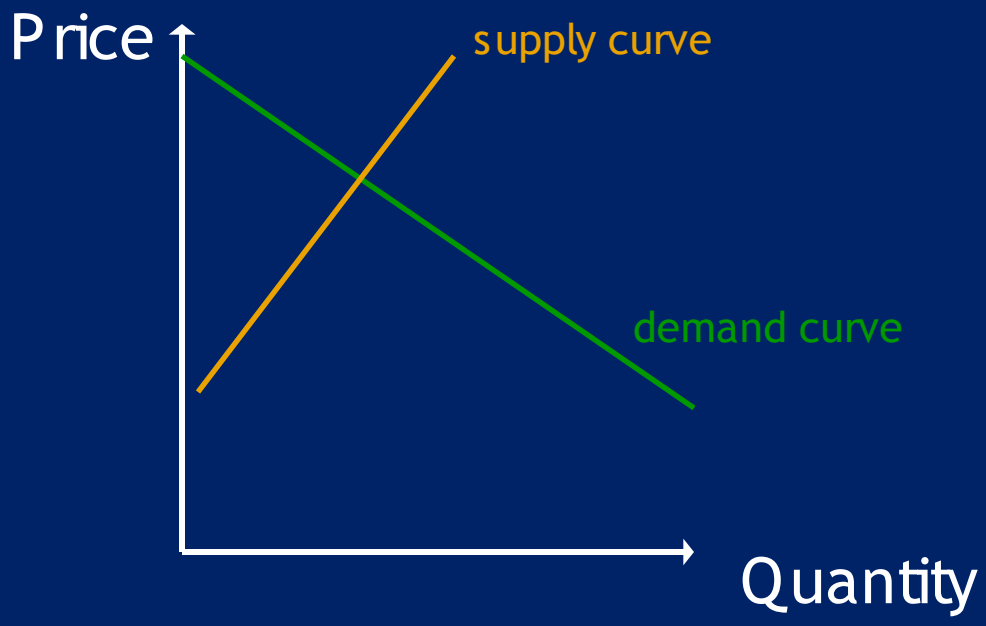

Background Economics #2b

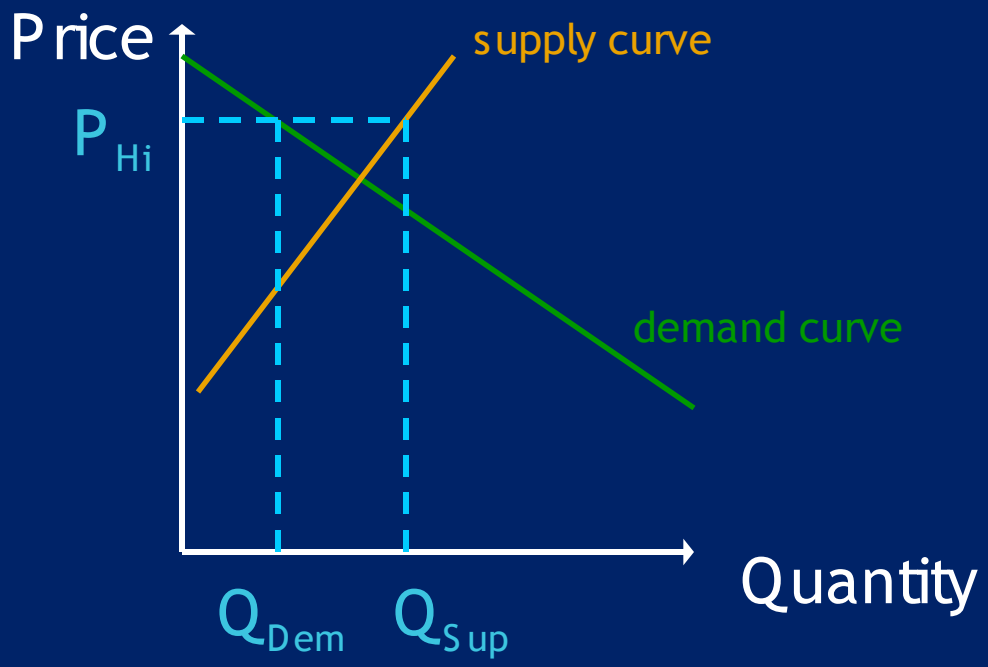

At high prices the quantity supplied will exceed the quantity demanded

Background Economics #2c

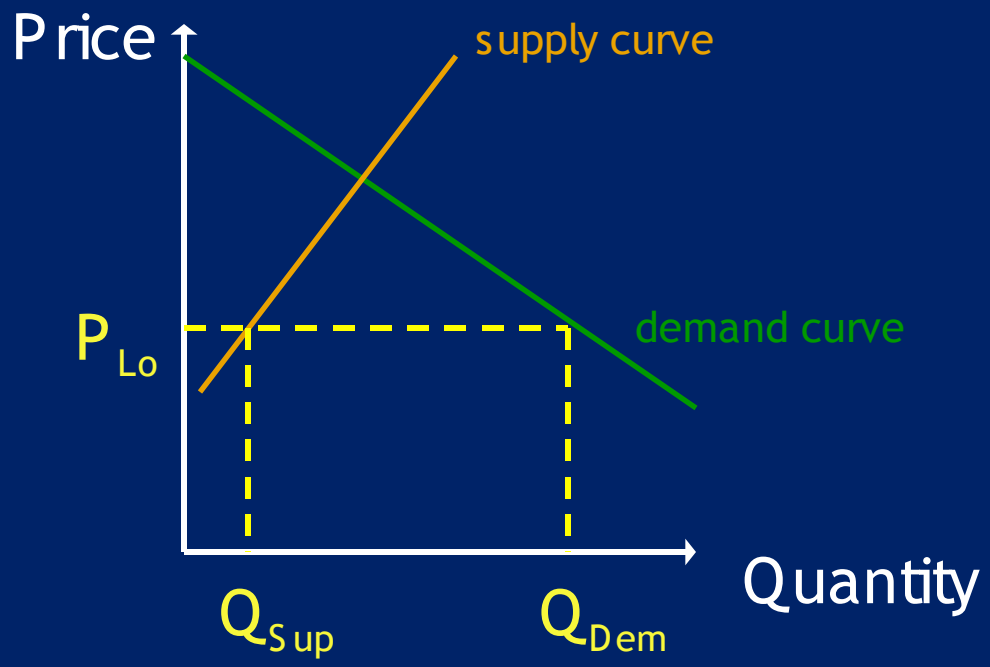

At low prices the quantity demanded will exceed the quantity supplied

Background Economics #2d

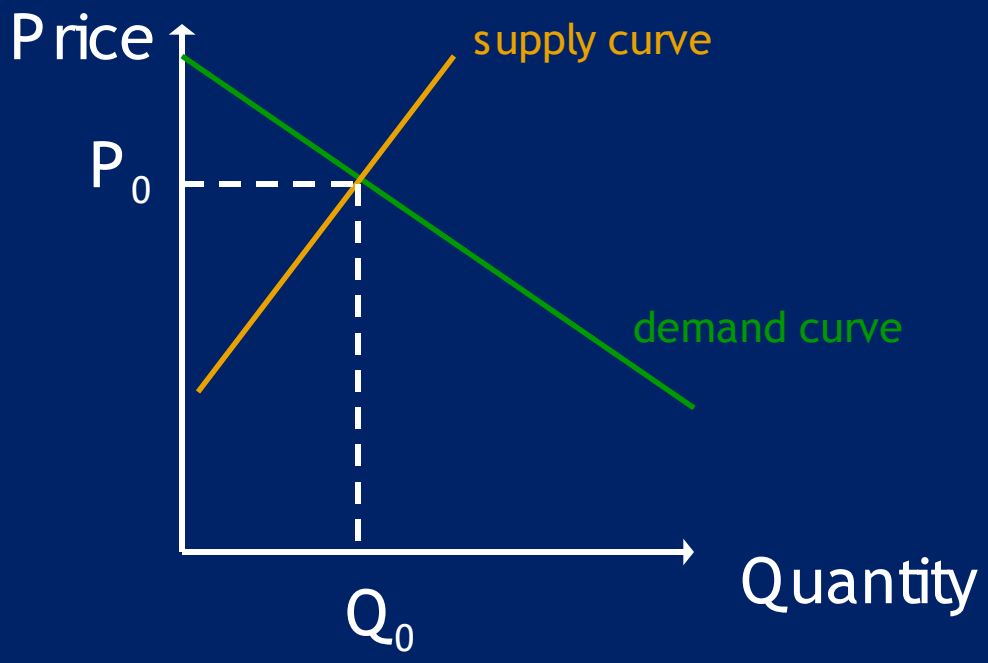

At an intermediate price, quantity supplied will equal quantity demanded
This is the Equilibrium Price P_0 (and corresponding Equilibrium Quantity Q_0).

Background Economics #3

- Free markets are self-correcting: if transactions take place at prices away from equilibrium, competition among buyers or sellers will move prices back toward the equilibrium; at the equilibrium, traders have no incentive to change their prices: a competitive equilibrium is achieved and maintained, until supply or demand change
- Such market mechanisms can give efficient allocation of resources without a central controller or external intervention. A common ideal of efficient allocation is the notion of Pareto Efficiency
  - An allocation is Pareto efficient if no-one can be made better-off without someone being made worse-off
  - Pareto efficient allocations can arise from free markets despite the fact that each trader in the market is acting only to serve his or her self-interest…
  - …the traders appear to be led by an invisible hand
  - Free markets are not guaranteed to achieve optimal allocations: conditions in which they fail are well known

Auctions and Traders

- In any economy, buyers & sellers meet and trade in auctions
  - there are many types of auction e.g.
    - English (ascending bid)
    - Dutch (descending offer)
    - retail (posted offer)
  - The continuous double auction (CDA) is very interesting
- traders in a CDA deal with noisy asynchronous data from multiple sources and must react in real-time to maximize their utility
- Classical microeconomic theory breaks down when we try and analyze this kind of market
- Without any central control or coordination, human traders in a CDA rapidly and reliably converge on the market's theoretical equilibrium price
- How much intelligence do they need to do that?
- D Cliff & J Bruten. Animat market-trading interactions as collective social adaptive behavior. *Adaptive Behavior* **7**(3):385-414, 1999. http://www.hpl.hp.com/techreports/97/HPL-97-91.html

Experimental Economics

- Laboratory-style studies of human market-trading behaviours
- A small number of human subjects divided into "buyers" and "sellers"
- All traders given a private-value limit price
  - buyers given sums of cash
  - sellers given units of a commodity
- Traders interact within some market mechanism

- Demonstrated rapid equilibration in CDA with very small numbers of traders (e.g. 12 in total) – and there is a lot going on in and around their heads
- NB: professional traders tend to do only slightly better than naïve subjects
- V L Smith. An experimental study of competitive market behavior. *Journal of Political Economy*. 70:111-137, 1962

Economics 101 vs Reality: it ain't smooth

Economics 101

- intersection of smooth supply and demand curves gives theoretical equilibrium price P_0 and equilibrium quantity Q_0

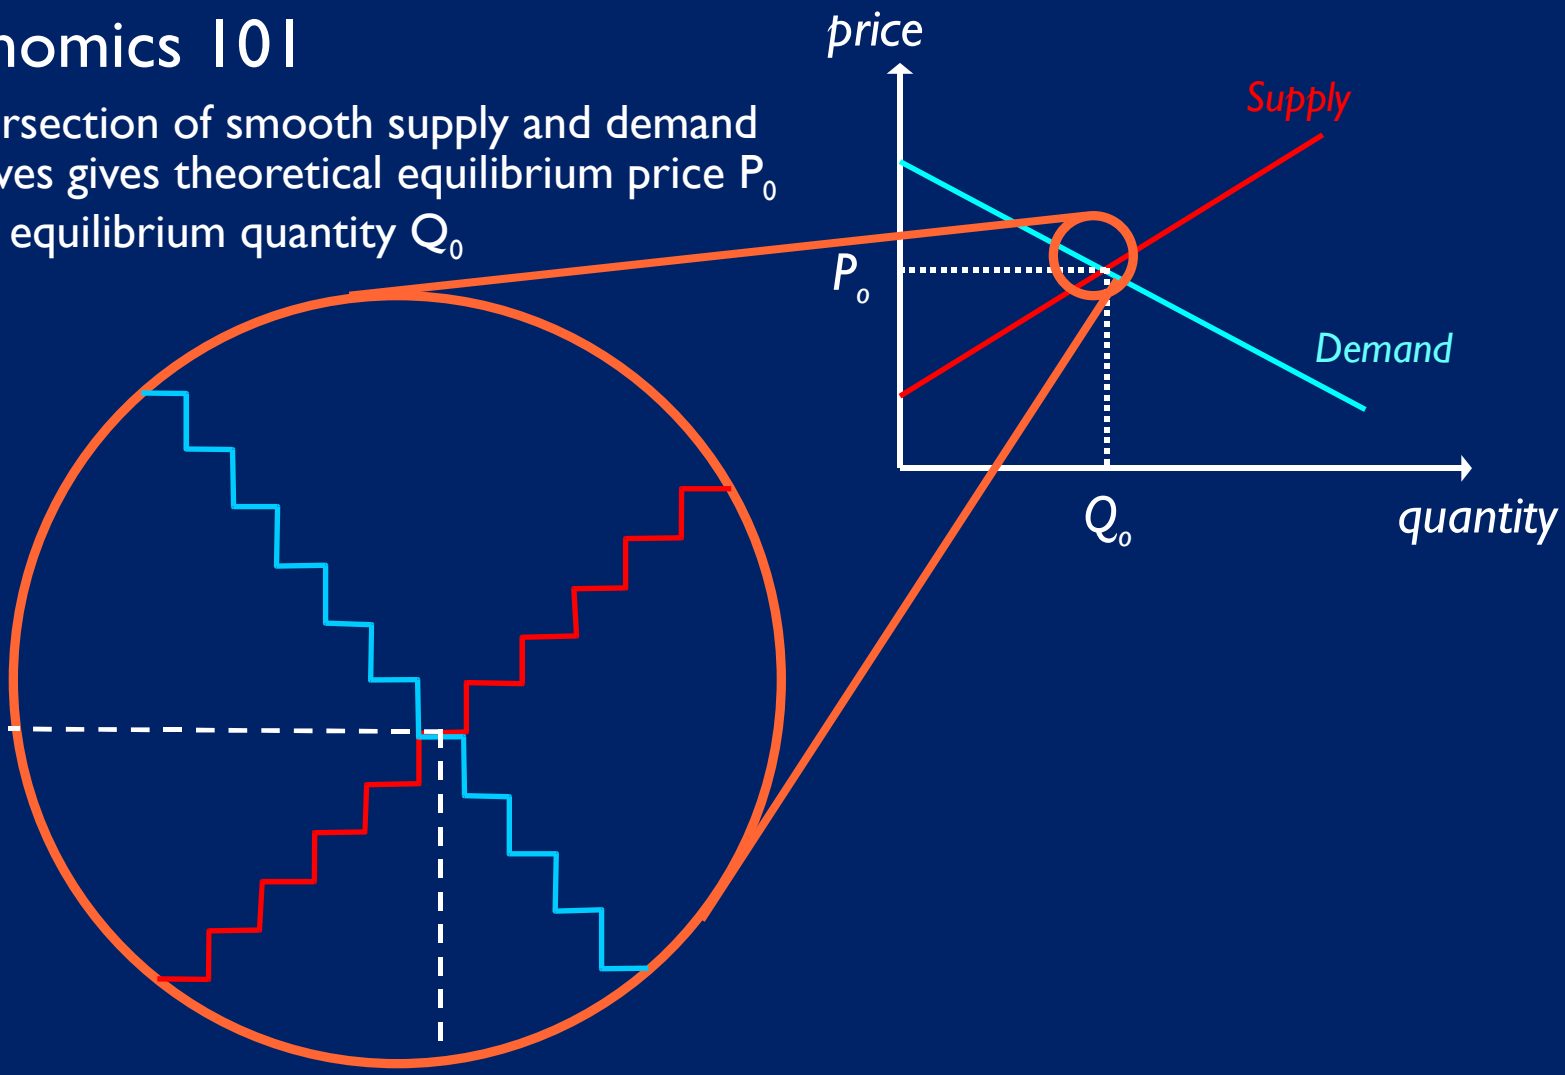

Economics 101 vs Reality: it ain't smooth

Economics 101

- intersection of smooth supply and demand curves gives theoretical equilibrium price P_0 and equilibrium quantity Q_0

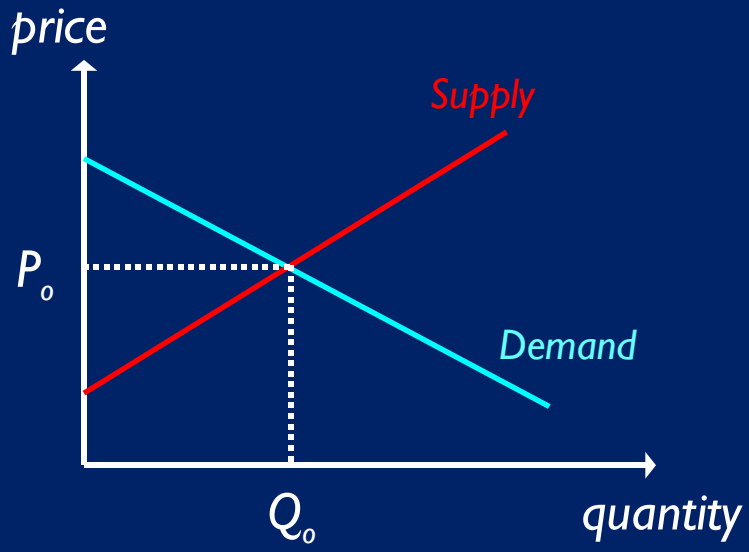

CDA reality #1: it's stepped

- supply and demand curves stepped
- each step in the supply curve represents an additional unit available for sale at the indicated price
- each step in the demand curve represents an additional unit desired at the indicated price

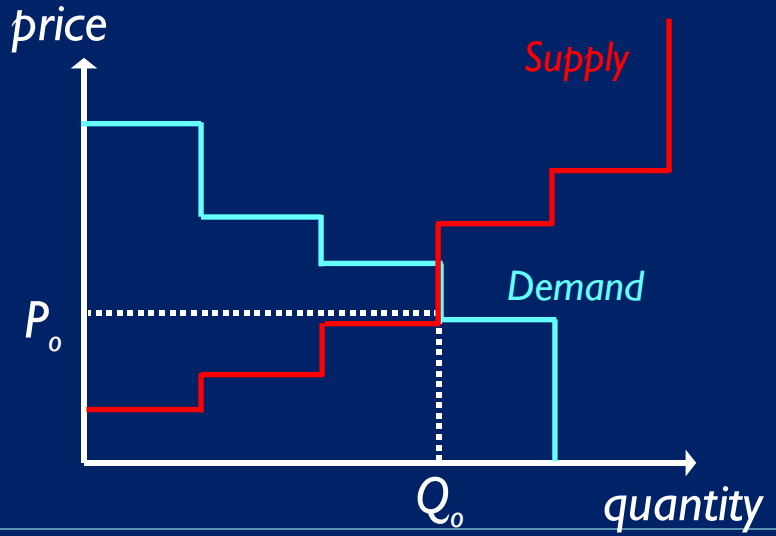

Economics 101 vs Reality: everybody lies

CDA reality #1: it's stepped

- supply and demand curves stepped
- each step in the supply curve represents an additional unit available for sale at the indicated price
- each step in the demand curve represents an additional unit desired at the indicated price

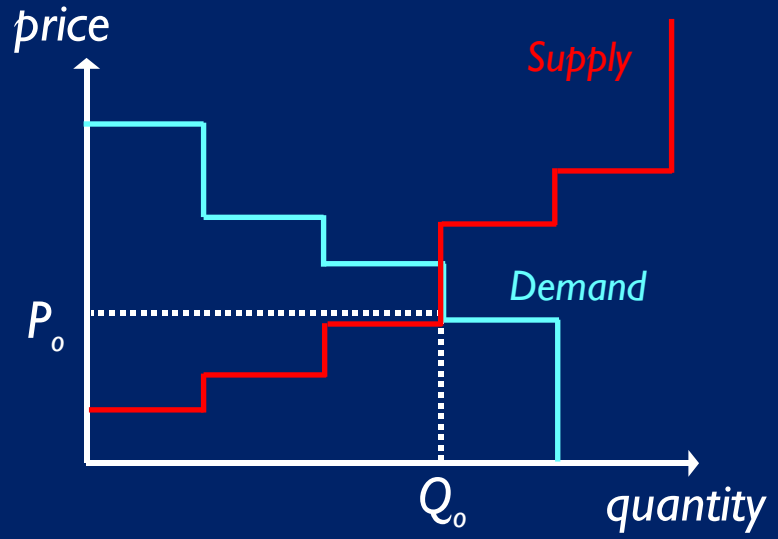

CDA reality #2: they lie

- underlying curves hidden by traders attempting to maximize utility
- the apparent supply and demand may be very different; giving different apparent values of P_0 & Q_0

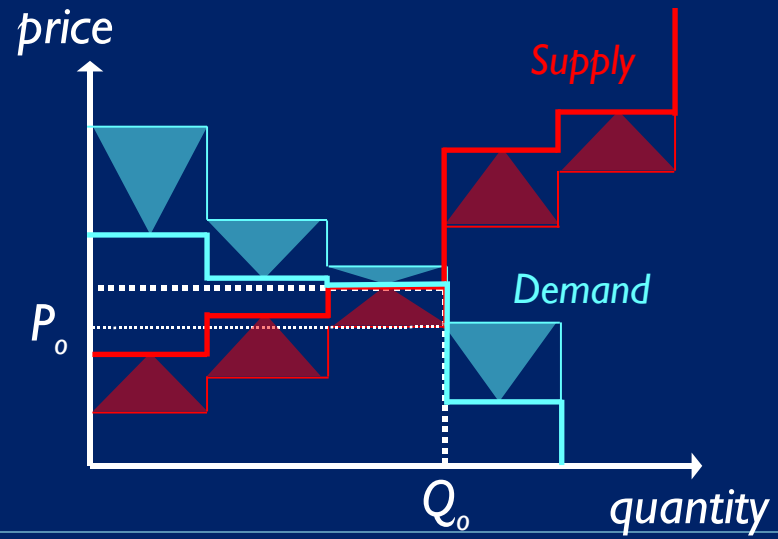

Economics 101 vs Reality: it don't stand still

CDA realities #1& #2

- supply and demand curves stepped
- underlying curves hidden by traders attempting to maximize utility

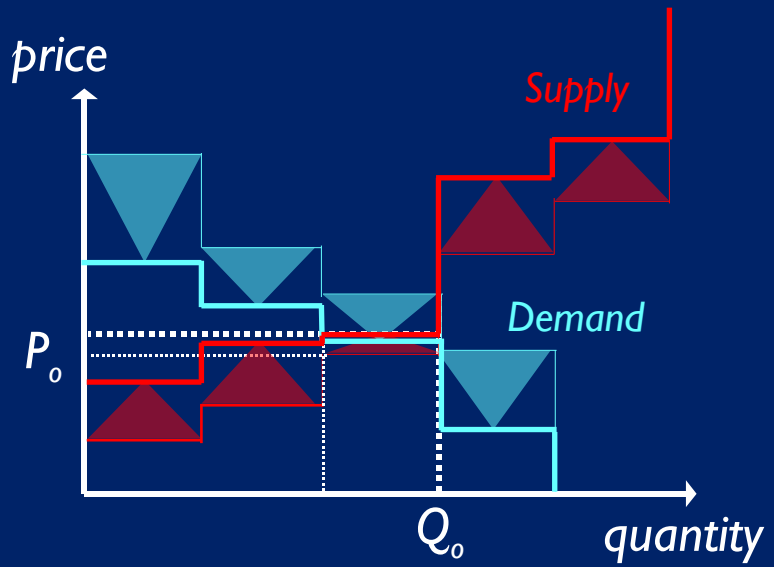

CDA reality #3: it moves

- dynamic variation in both the curves, and so in values of $P_0(t)$ and $Q_0(t)$…
- small number of traders means that one buyer and one seller agreeing a transaction and leaving the market can lead to relatively big changes in the supply and demand curves
- this happens after every transaction

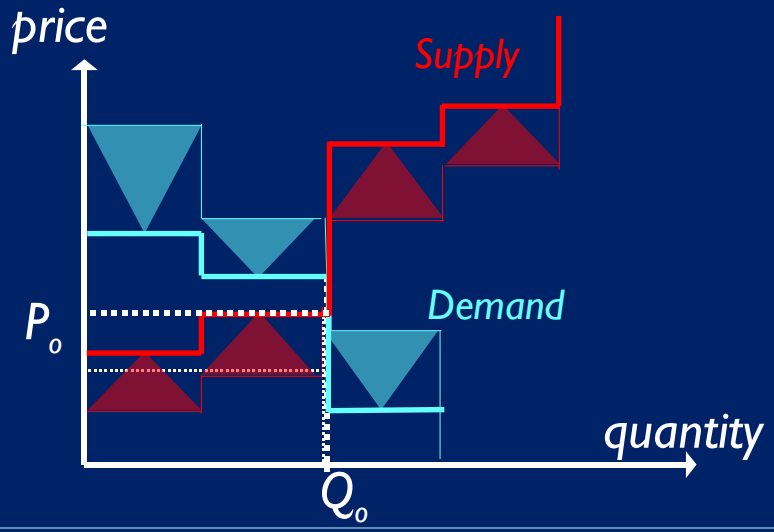

Economics 101 vs Reality: liars tell new lies

CDA realities #1, #2, & #3

- It's stepped, they lie, it moves

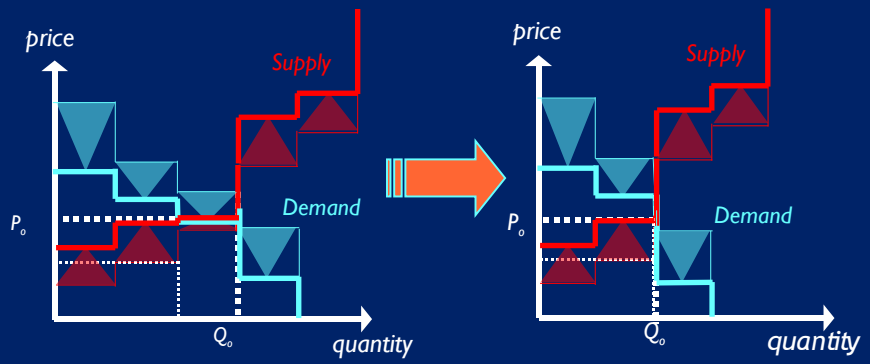

CDA reality #4: the lying liars keep changing their lies

- Even without transactions, traders dynamically alter their margins

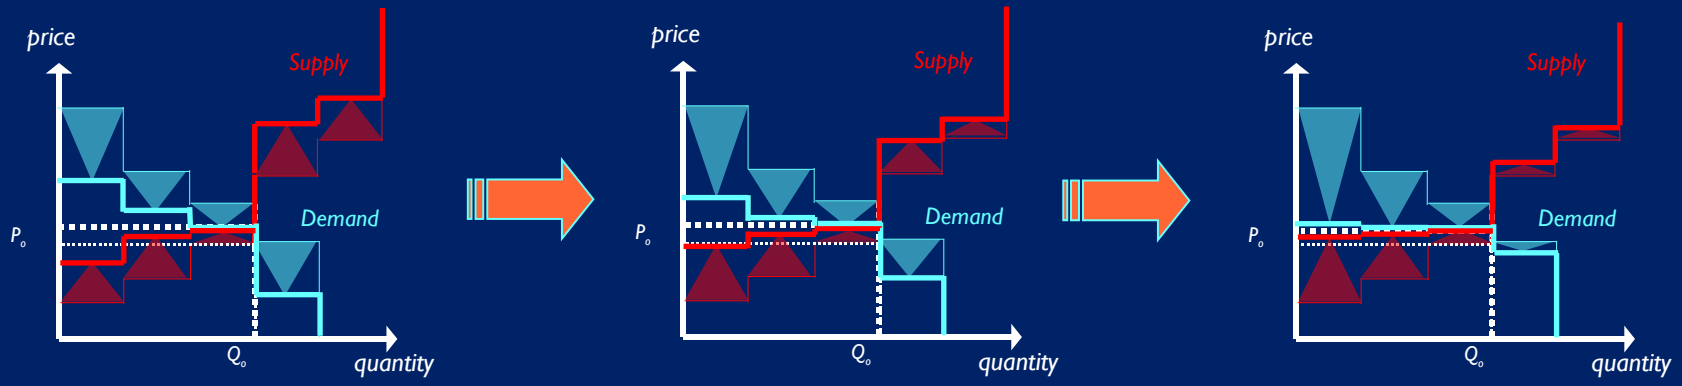

Really Reality: Reuters FX dealing

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

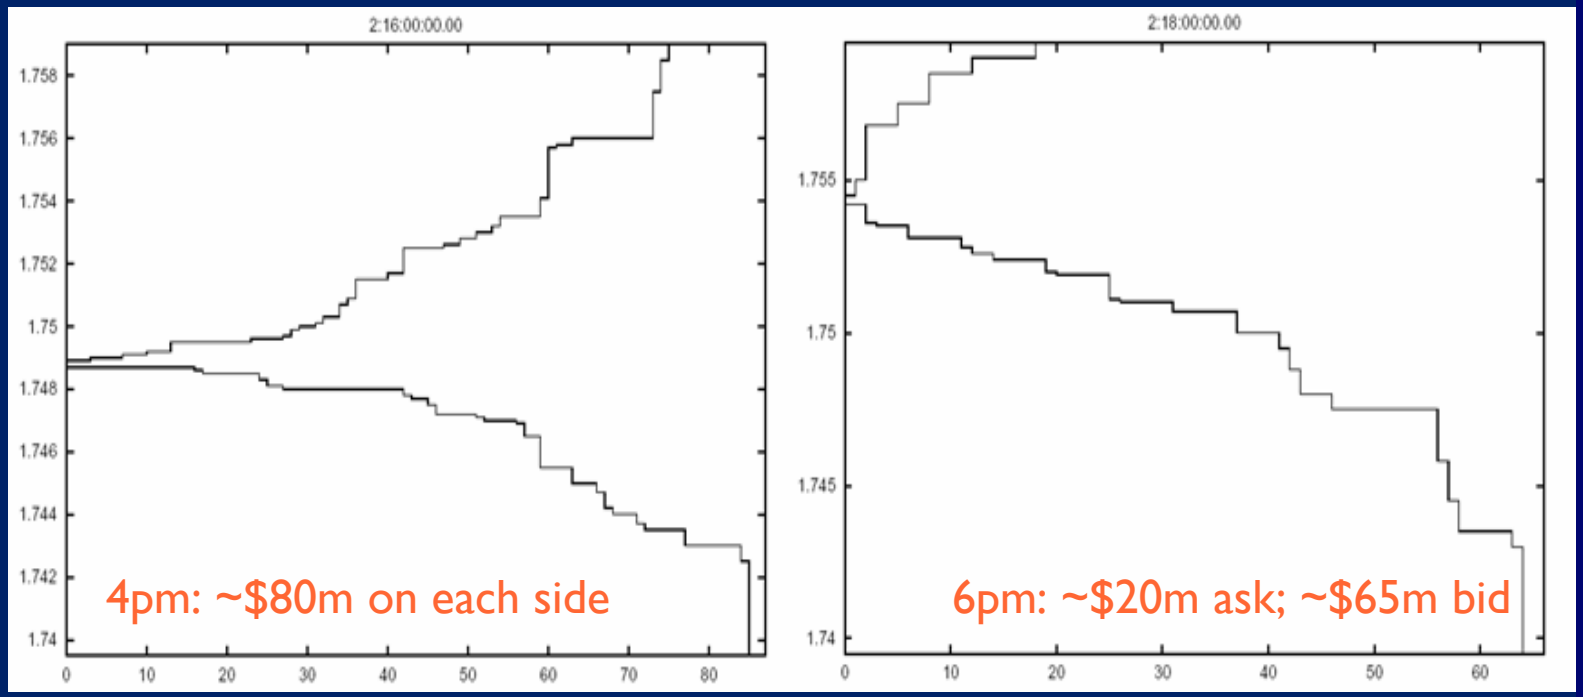

- J. Danielsson & R. Payne "Measuring & explaining liquidity on an electronic limit order book: evidence from Reuters D2000-2". Manuscript, London School of Economics (2003). www.riskresearch.org

Experimental Economics

- laboratory-style studies of human market-trading behaviors
- Pioneered by Vernon Smith (who got a Nobel Prize for it)
- a small number of human subjects are split into "buyer" and "seller" groups
- all traders given a private-value limit price
  - buyers given sums of cash; can't buy above limit price
  - sellers given units of an arbitrary commodity; can't sell below limit price
- traders interact within some market mechanism
  - Buyers may quote bid prices
  - Sellers may quote offer prices
- demonstrated rapid equilibration in CDA with very small numbers of traders
- NB: professional traders tend to do only slightly better than naïve subjects
- V Smith. An experimental study of competitive market behavior. *Journal of Political Economy*. 70:111-137, 1962

Experimental Economics: VSmith's results

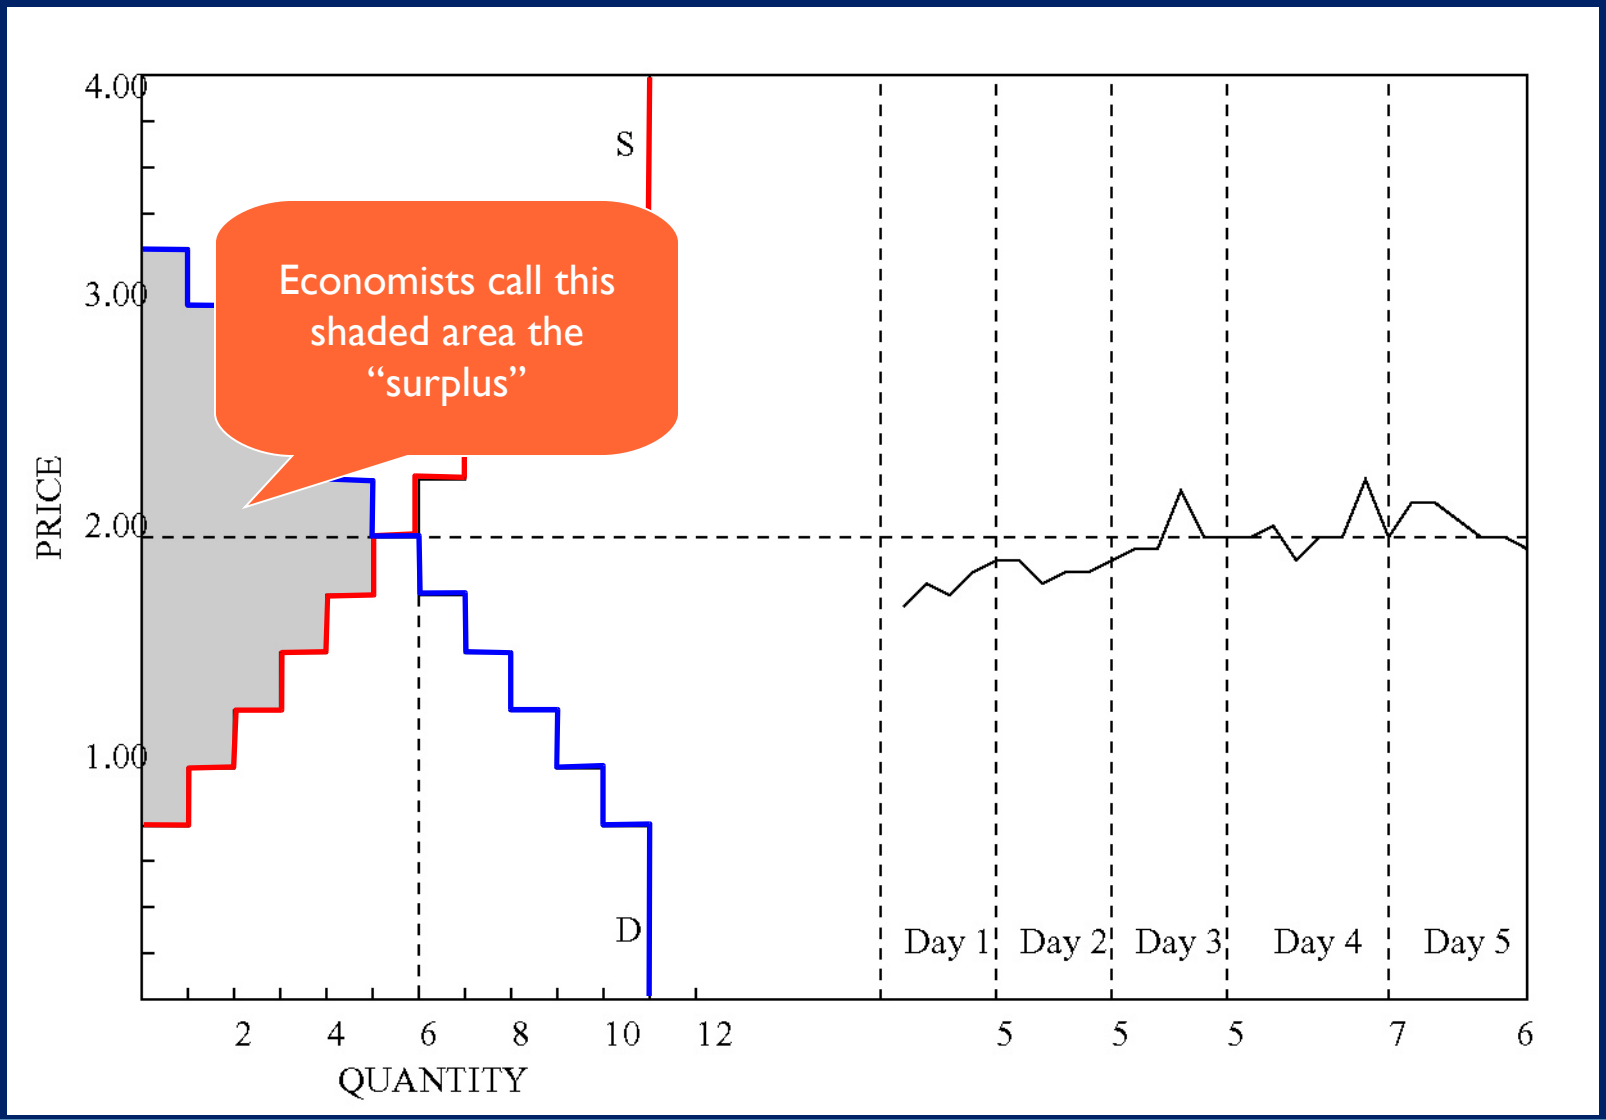

Experimental Economics: VSmith's results

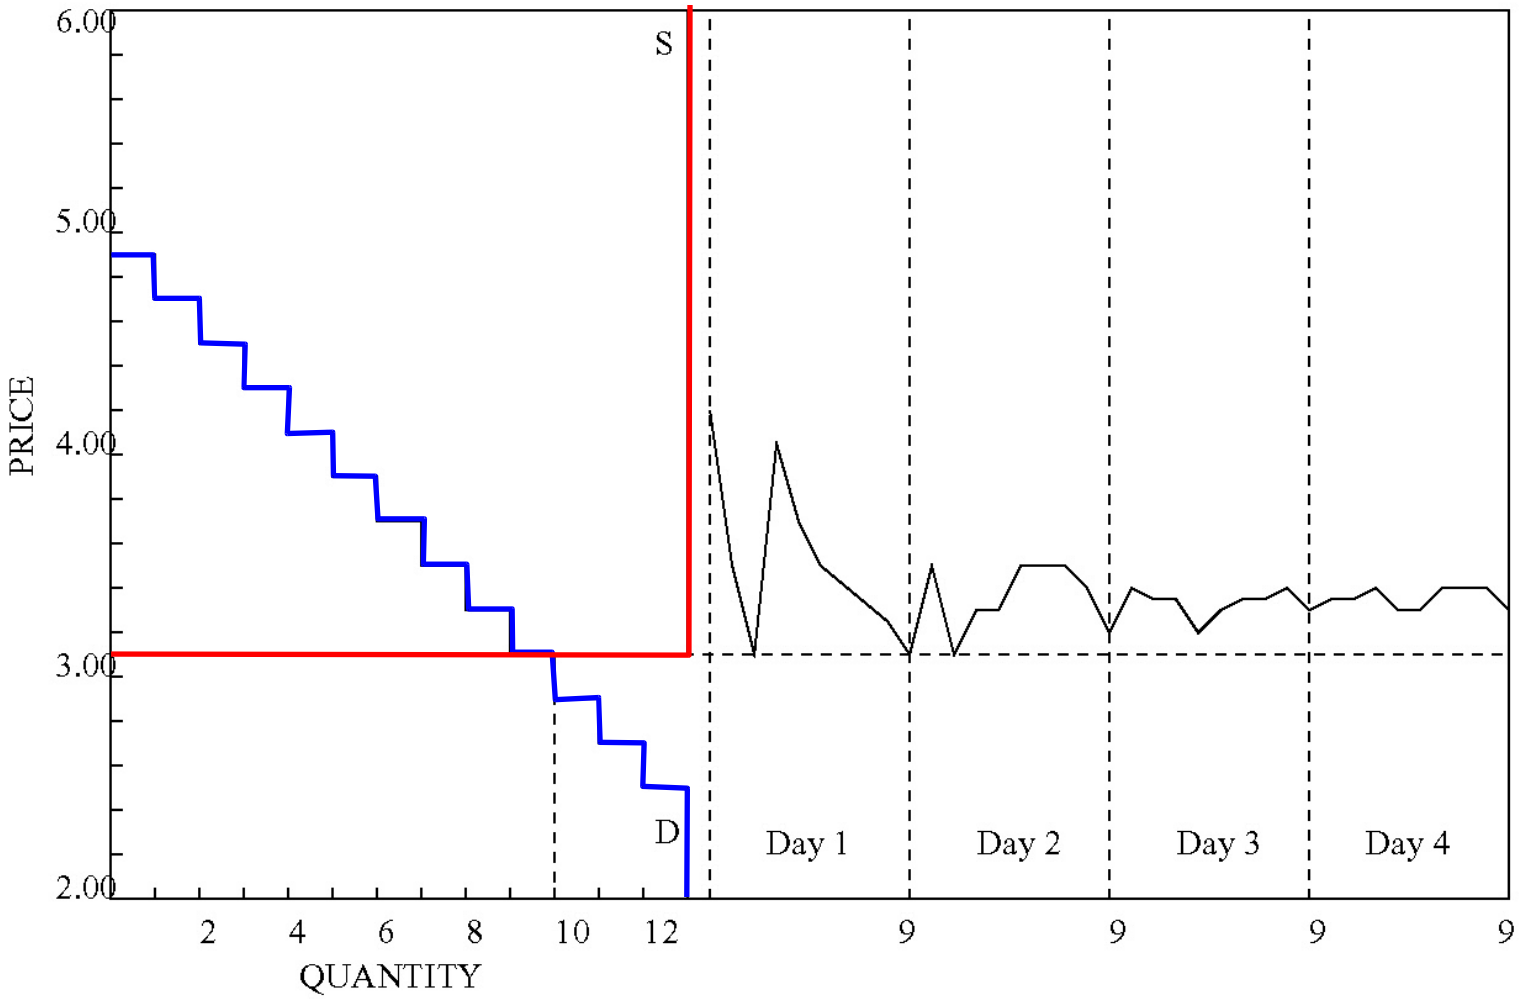

How much intelligence: zero, perhaps?

- How much of the efficiency of a CDA is due to the intelligence of the traders, and how much of it is due to the organisation of the market?
- Gode & Sunder '93 (*Journal of Political Economy* 101(1):119-137) introduced zero-intelligence (ZI) trading agents for CDA markets:
  - ZI-U: unconstrained: generate random bid/offer quote prices
  - ZI-C: random quote prices, but constrained not to trade at a loss
- Ran experimental economics tests with ZI-U, ZI-C, and human traders in 5 different markets, monitored their efficiency
  - ZI-U traders were plain useless
  - ZI-C traders were surprisingly human-like
- Conclusions:
  - most of the intelligence is in the market, not in the traders

Gode & Sunder Results: Humans

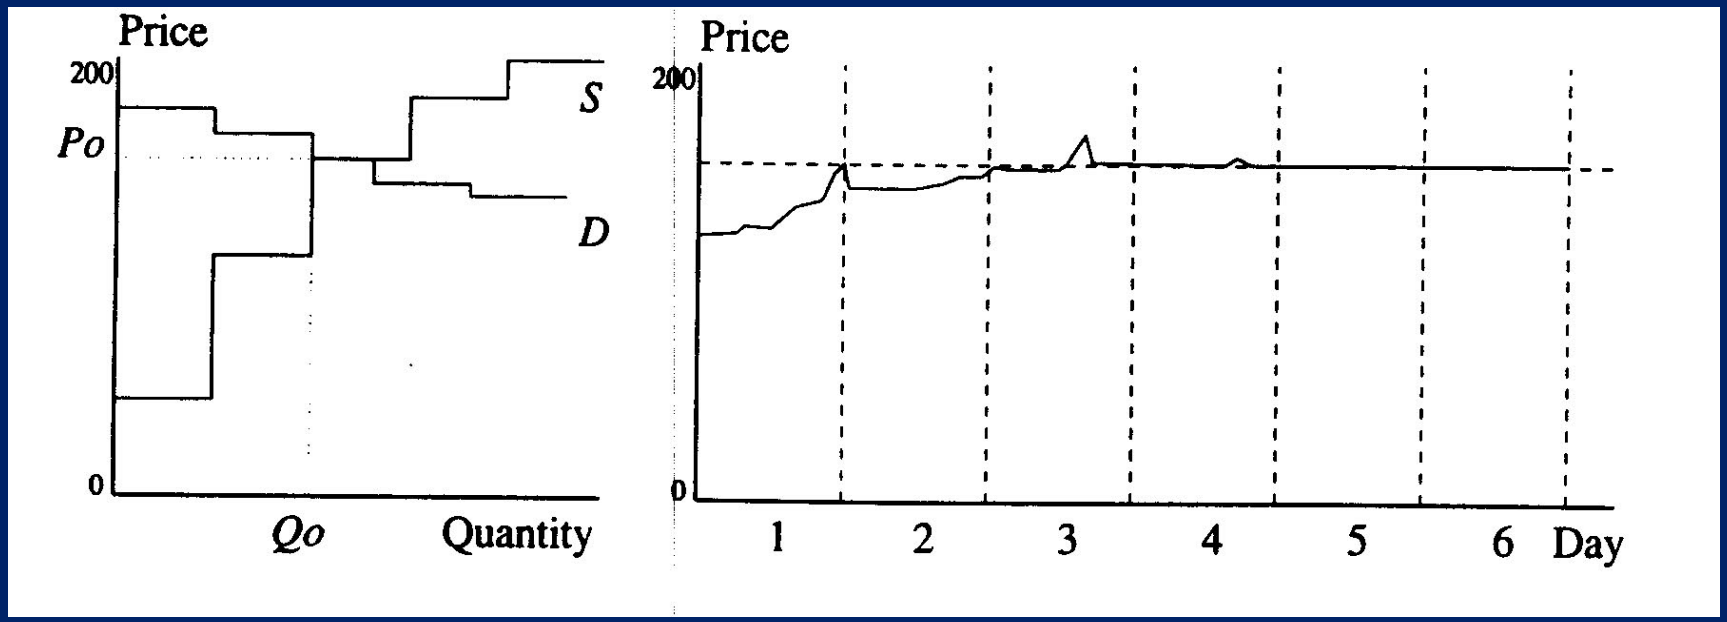

Gode & Sunder Results: ZI-U

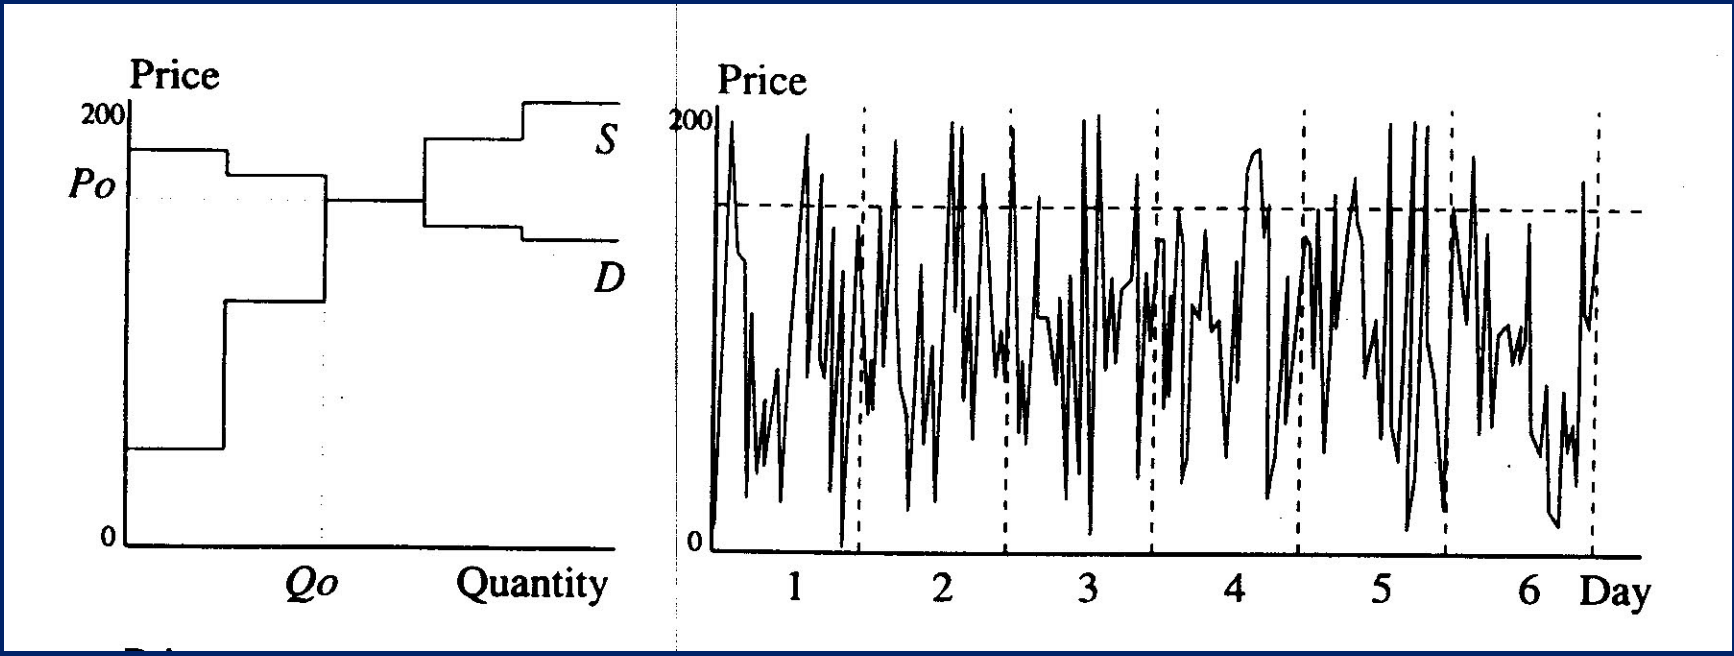

Gode & Sunder Results: ZI-C

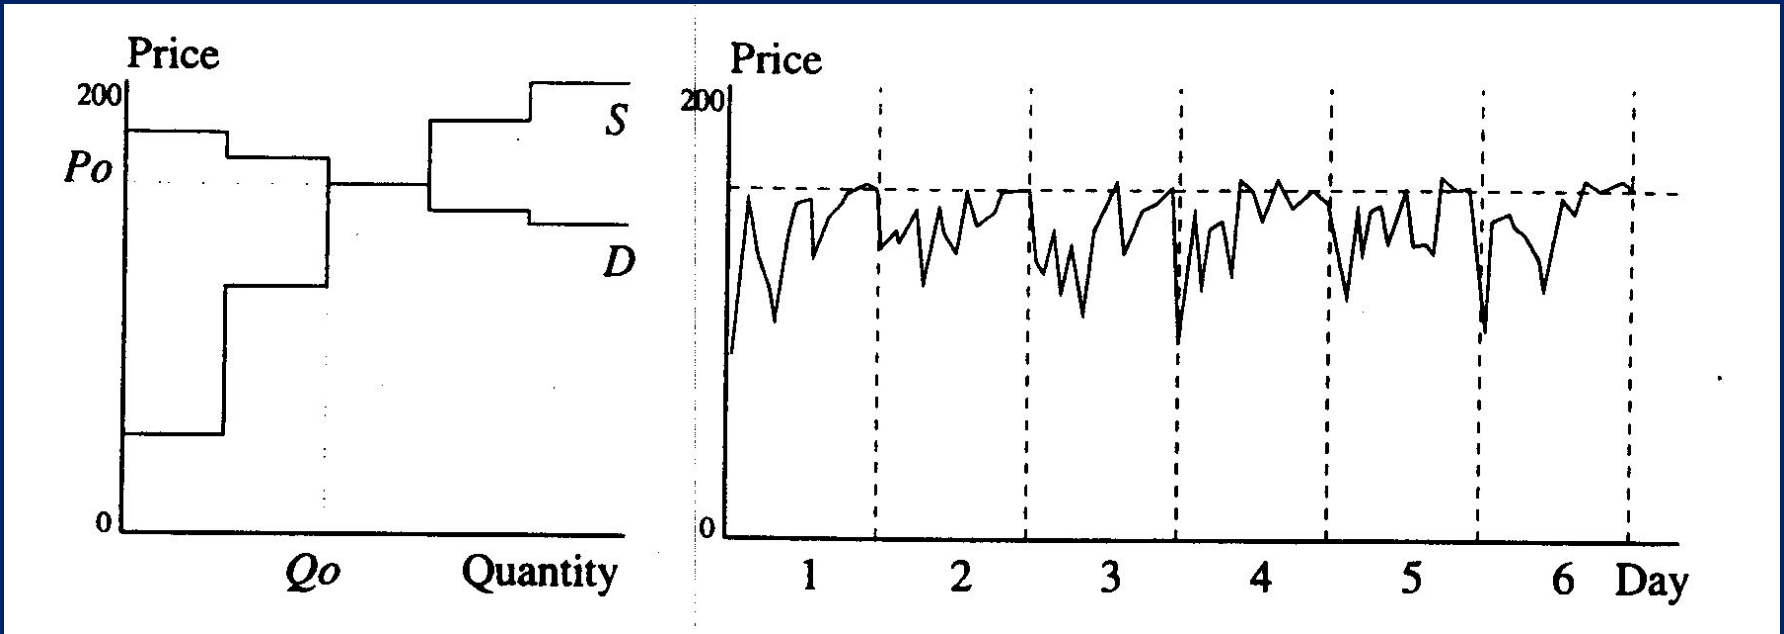

Lecture 1: Summary

- BlahBlahBlah