CS395T
Agent-Based Electronic Commerce
Fall 2006

Peter Stone

Department or Computer Sciences
The University of Texas at Austin

Week 5a

Good Afternoon, Colleagues

Are there any questions?

Logistics

- Next week's readings

TAC Supply Chain Management Scenario

- 6 agents are PC manufacturers
- 220 simulated game days
- suppliers and customers modeled by game server

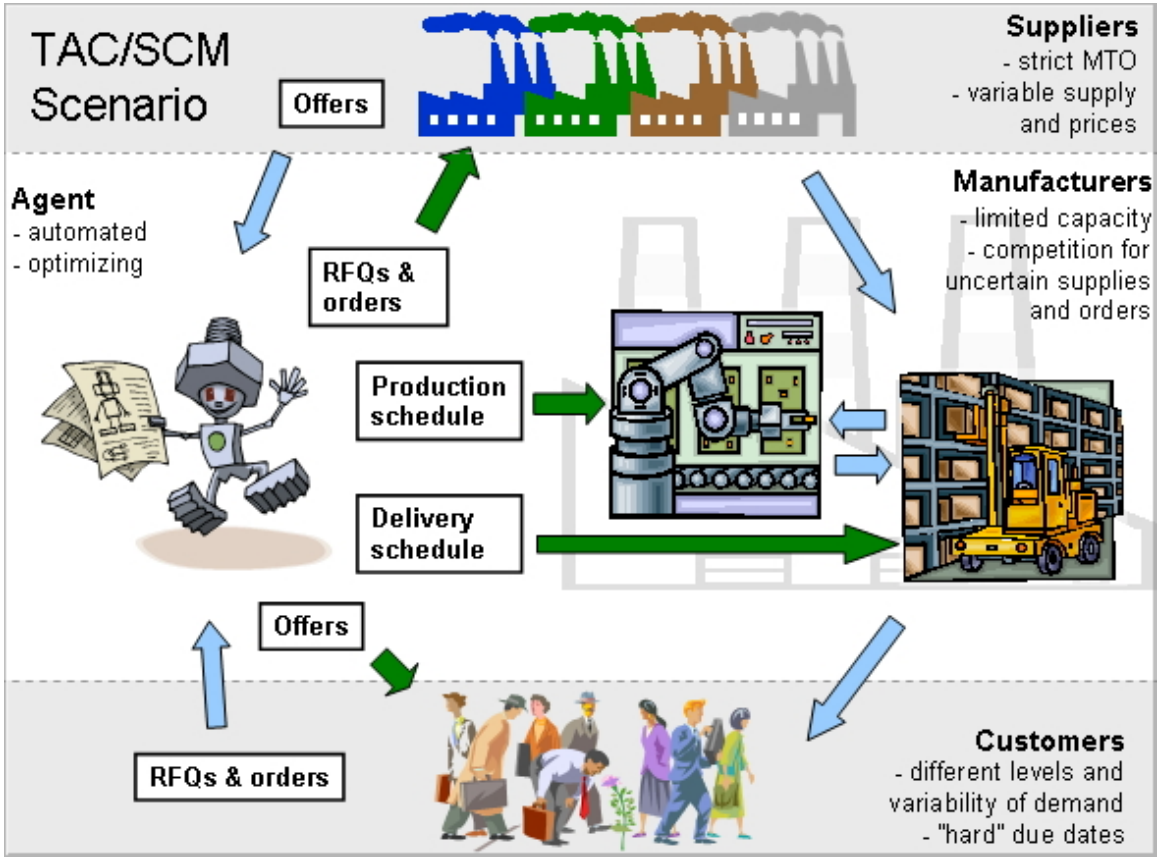

TAC Supply Chain Management Scenario

- 6 agents are PC manufacturers
- 220 simulated game days
- suppliers and customers modeled by game server

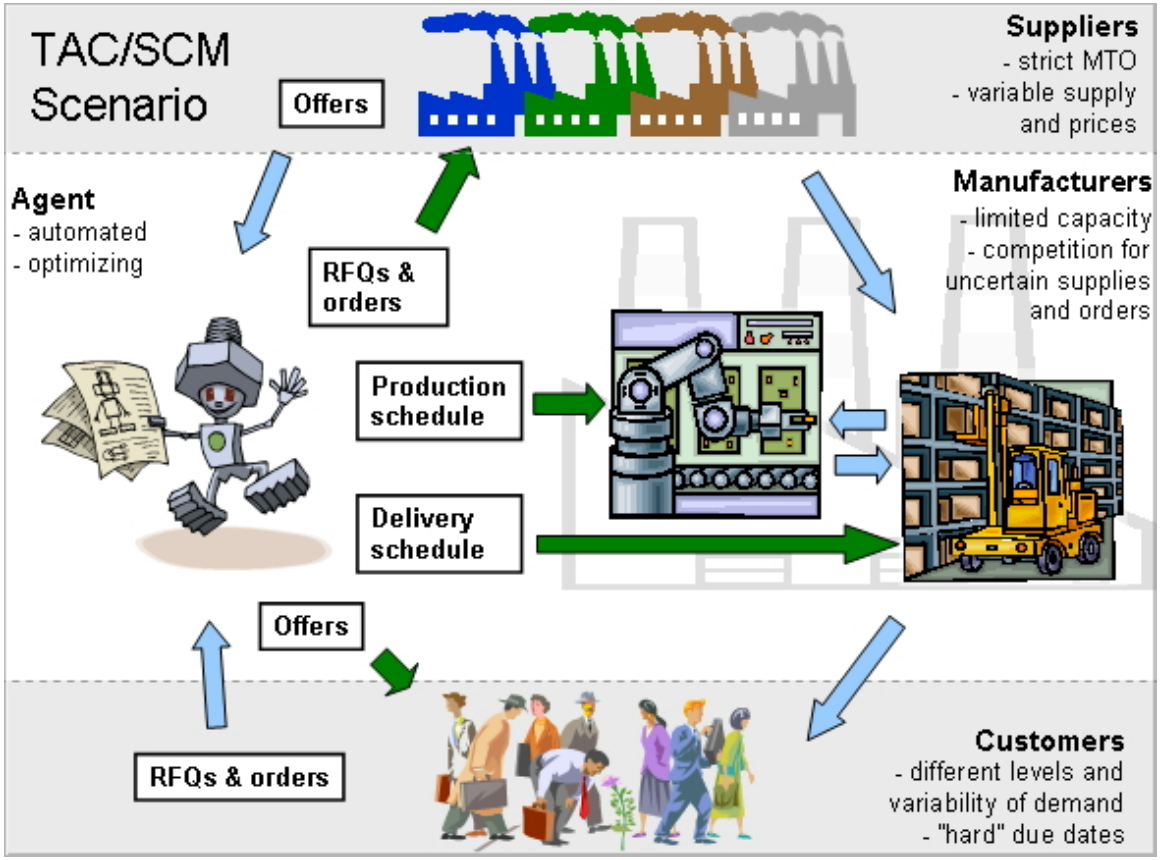

Challenges:

- Incomplete information
- Time constraints: each simulated day lasts 15 seconds

Agents' Daily Decisions

- Issue RFQs to suppliers
- Accept/reject supplier offers
- Plan days production mix
- Select completed orders to ship
- Bid on customer RFQs

Agents' Daily Decisions

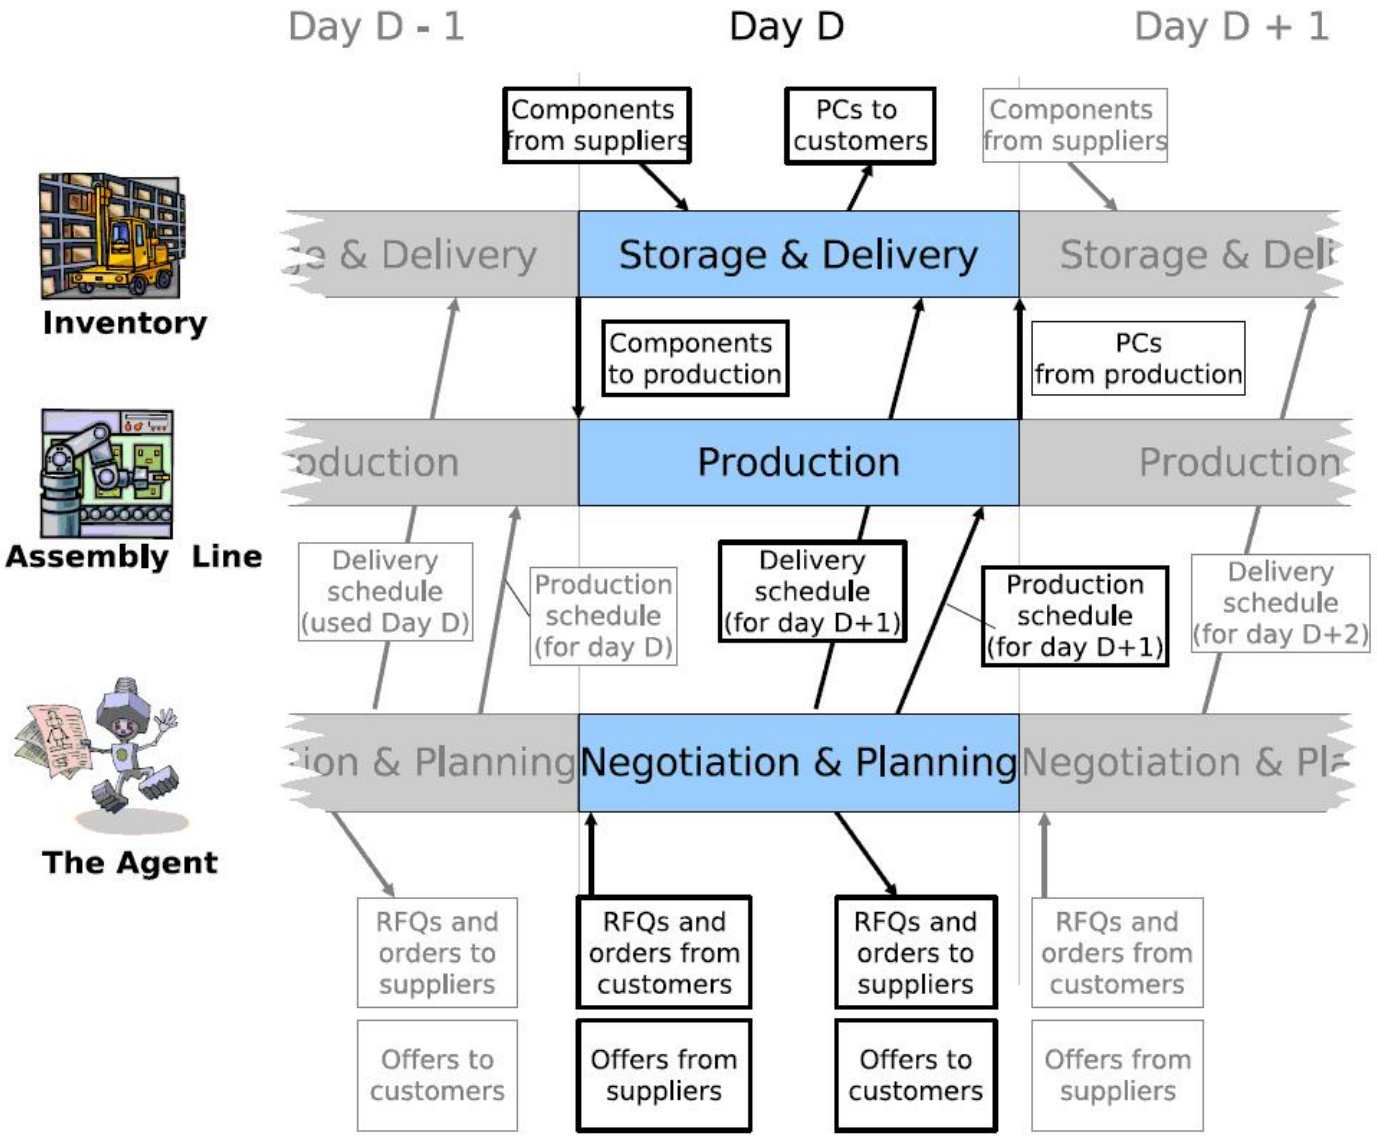

TAC SCM Problems and Techniques

- Dynamic optimization under uncertainty
- Price prediction
- learning and adaptivity
- Multiattribute negotiation
- Strategic bidding and procurement
- Experimental methodology

TAC SCM Problems and Techniques

- Dynamic optimization under uncertainty
- Price prediction
- learning and adaptivity
- Multiattribute negotiation
- Strategic bidding and procurement
- Experimental methodology

TAC experience has yielded contributions to AI literature on these and other topics.

Class Discussion

- Jeremy Stober on linear programming

Learning in TAC

- Travel
- SCM

TacTex-05

- A **predictive** agent
  - Predicts customer demand: Bayesian modeling
  - Predicts order probability: linear model
  - Predicts supplier prices: price probes

TacTex-05

- A **predictive** agent
  - Predicts customer demand: Bayesian modeling
  - Predicts order probability: linear model
  - Predicts supplier prices: price probes
- An **adaptive** agent
  - Adaptive first day ordering:
    supply prices depend on other agents' bids
  - Adaptive end-of-game bidding:
    computer prices depend on other agents' inventory

TacTex-05 overview

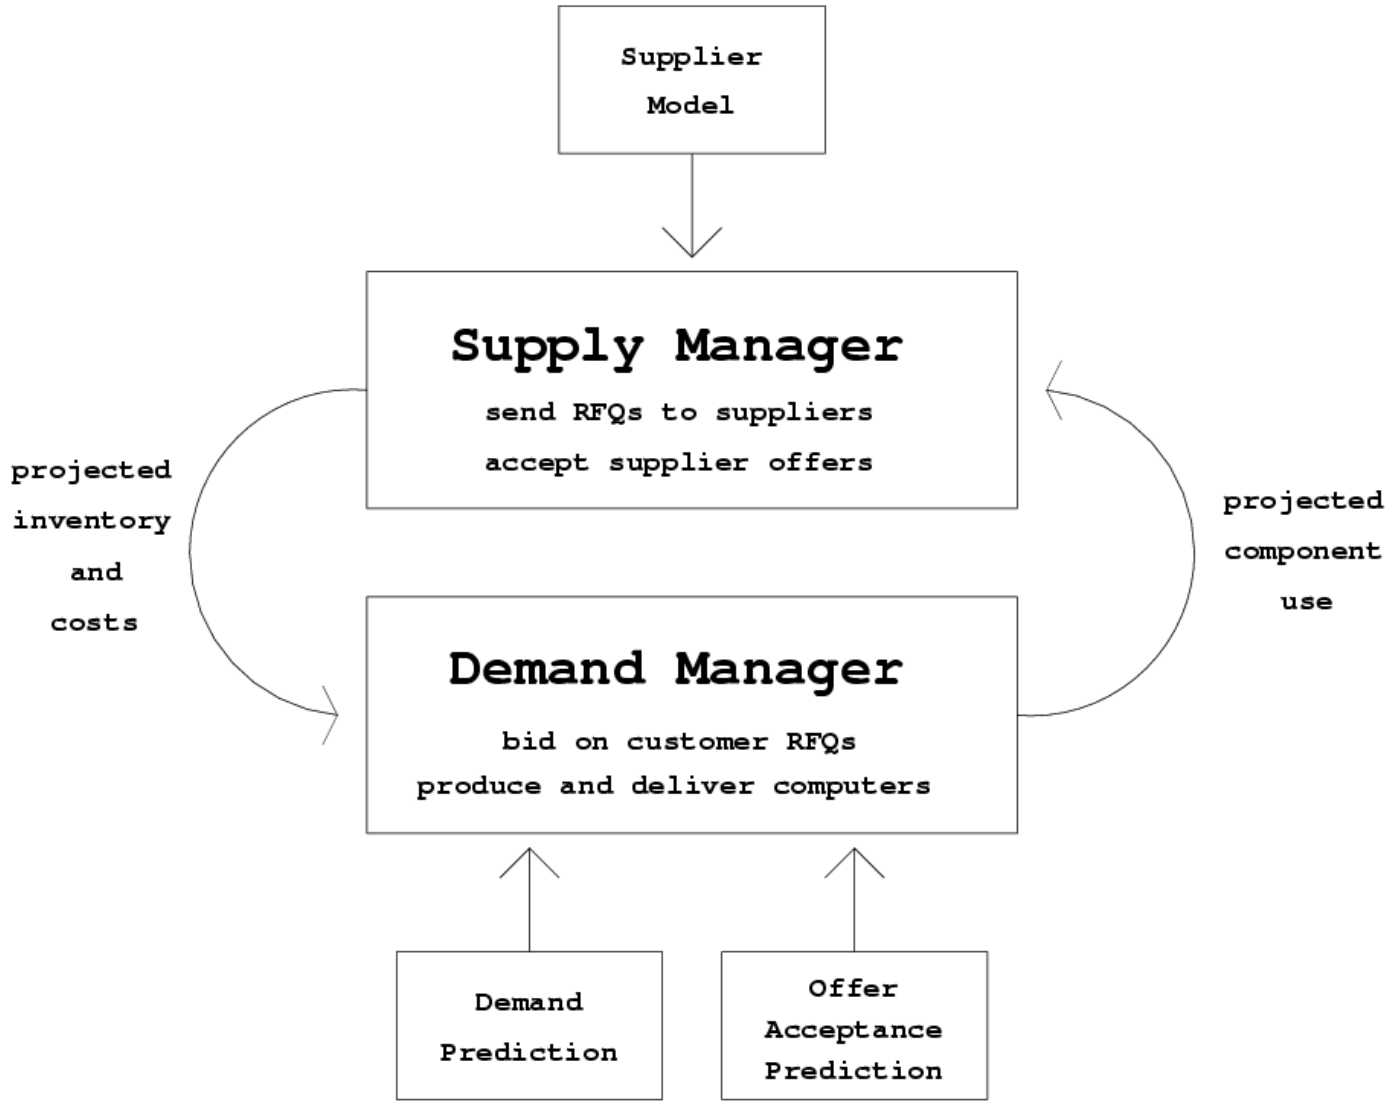

TacTex-05 results

- Won the seeding round, but that's **no guarantee**

TacTex-05 results

- Won the seeding round, but that's **no guarantee**
- Won the finals too! **Adaptation** evident

TacTex-05 results

- Won the seeding round, but that's **no guarantee**
- Won the finals too! **Adaptation** evident
- **Controlled testing** in progress

Bidding for Customer Orders

Customers send Requests for Quotes (RFQs) consisting of:

Bidding for Customer Orders

Customers send Requests for Quotes (RFQs) consisting of:

- the **type** of computer desired (1 of 16 types)
- the **quantity** of computer desired (1 – 20)
- the **due date** (3 – 12 days in the future)
- a **reserve price** indicating the maximum the customer will pay
- a **penalty** that must be paid for each day the delivery is late

Bidding for Customer Orders

Customers send Requests for Quotes (RFQs) consisting of:

- the **type** of computer desired (1 of 16 types)
- the **quantity** of computer desired (1 – 20)
- the **due date** (3 – 12 days in the future)
- a **reserve price** indicating the maximum the customer will pay
- a **penalty** that must be paid for each day the delivery is late

Agents submit **sealed bids**; customers accept lowest offers

Daily reports indicate yesterday's high and low prices

The Bidding Problem

What bids should an agent place on the RFQs it receives each day?

The Bidding Problem

What bids should an agent place on the RFQs it receives each day?

Our approach:

- bid on a large number of RFQs in hopes of winning some fraction
- BUT avoid receiving more orders than can be filled

The Bidding Problem

What bids should an agent place on the RFQs it receives each day?

Our approach:

- bid on a large number of RFQs in hopes of winning some fraction
- BUT avoid receiving more orders than can be filled

This requires:

- a method of **predicting the probability** of winning with a certain bid
- a means of **using these predictions** to maximize expected profit

Learning Bid Acceptance Probabilities

Problem: given the attributes of an RFQ and knowledge of the game conditions, predict the probability of winning the order as a function of the price bid - a **conditional density estimation problem**.

Learning Bid Acceptance Probabilities

Problem: given the attributes of an RFQ and knowledge of the game conditions, predict the probability of winning the order as a function of the price bid - a **conditional density estimation problem**.

We take an approach used previously in a different TAC scenario: (Schapire, Stone, McAllester, Littman, and Csirik 2002)

- Divide the price range into several bins
- Train a **separate predictor** for each endpoint with a regression learner
- Interpolate to derive a function