Ahead of the Curve:
Designing Prices and Markets When Data is Imperfect

Learn how to use practical analytics to make better pricing and selling decisions, resulting in improved revenue outcomes, even when market data is limited, noisy, or constantly changing.

LUNCH AND LEARN

Thursday, May 7, 2026 | 12pm ET

Thank you for attending!

Featuring
Shixin Wang

Upcoming Course Calendar
scl.gatech.edu/course-calendar

scl.gatech.edu

Supply Chain and Logistics Institute

The Pricing Decision You Make Every Day

Procurement Auctions

Buying freight, materials, services — bidders are anonymous

Multi-Product Pricing

Software bundles, parts catalogs, MRO supplies

New Product Launches

Cold-start SKUs, seasonal items, geographic expansion

What the textbook says:
"Find the optimal price assuming you know the demand distribution."

What you actually have:
Ranges, a few transactions, an AI prediction, or a regulator telling you what data you can use.

Wrong assumptions can cost 30–70% of achievable revenue.

Five Types of Information You Actually Have

Bounds

"No one pays above $200, no one below $20"

Insight 1

Quantile

"60% bought at $40" from a price experiment

Insight 1

Few Samples

Sometimes just one past transaction to learn from

Insight 3

AI Signal

LLM/ML estimate of buyer value — may hallucinate

Insight 3

Auction Data

Only winning bid, not bidder identities

Insight 2

Today: three insights, each addressing these information regimes

Insight 1

Simple Beats Sophisticated

When you don’t know the distribution, a 2-option menu captures most of the value of the optimal mechanism with infinite options.

The Power of Simple Menus in Robust Selling Mechanisms
Management Science, 2025

A Pricing Puzzle

You know:
Customers value the product between $0 and $100

From a price experiment:
At least 60% will buy at $40
That's all you know.

Example from Wang (Management Science, 2025)

The Cost of Guessing Wrong

Distribution A (assumed)

40% at $0, 30% at $40, 30% at $100

Optimal price: $100 | Revenue: $30

Distribution B (actual)

30% at $0, 60% at $40, 10% at $100

Optimal price: $40 | Revenue: $28

The Disaster

If you assumed Distribution A but the real one is B:
Your $100 price yields only $10 (10% × $100)
The right price would have earned $28

64%
revenue lost

Wrong distributional assumptions → severe revenue loss. This motivates robust pricing.

A Two-Option Menu Solves It

Option A: The Sure Thing

Pay $50, get the product for sure

High-value customer ($100) picks this

Utility = $100 – $50 = $50 ✓

Option B: The Lottery

Pay ~$33, get the product with prob. 5/6

Low-value customer ($40) picks this

Utility = 5/6×$40 – $33 = $0 (participates!)

Result: Revenue = $25 under EITHER distribution

The menu eliminates worst-case loss. No distributional guess needed.

A simple 2-option menu hedges against distributional uncertainty.

You Already Use Menus Everywhere

SaaS Tiers

Basic
$9/mo

Pro
$29/mo

Enterprise
$99/mo

Self-select tier

Shipping

$5 Standard

3–5 days

$15 Priority

next-day

Probability of speed

Insurance

Deductible | Premium

$500 $180/mo

$1,500 $120/mo

$5,000 $60/mo

Risk vs. price trade-off

Blind Boxes

Legendary 1%

Rare 9%

Common 90%

Fixed price, random reward

Airline Fares

Business
$1,200

Flexible
$520

Basic
$280

Risk + flex menu

All menus of lotteries in disguise — the theory validates what practitioners do intuitively.

How Many Options Do You Actually Need?

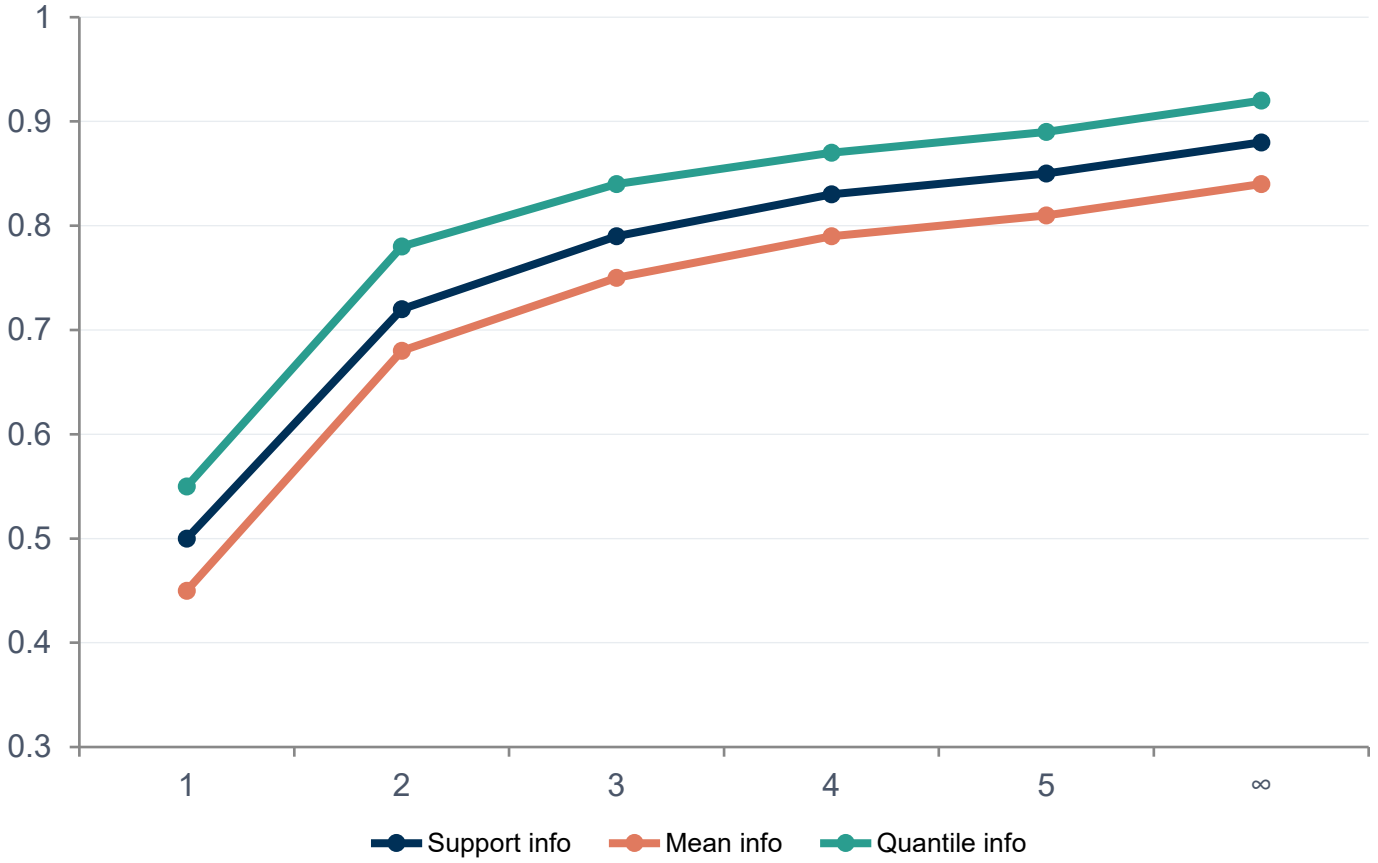

Key Finding
The elbow is at n = 2

A 2-option menu captures most of the gap between deterministic pricing and the infinite-menu optimum.

Diminishing returns kick in fast: n = 3 or 4 is rarely worth the operational cost.

Why Complexity Is a Tax

Operational Cost

Sophisticated systems, complex inventory management, demand forecasting across infinite options

Customer Confusion

Decision paralysis when presented with too many options — "paradox of choice" hurts conversion

Fairness Scrutiny

Continuum menus enable extensive price discrimination — regulatory and reputational risk

Each option you add costs you across all three dimensions — but buys very little robustness past 2 or 3.

Simplicity isn't a compromise — it's a feature.

Insight 1: Takeaway

Audit your menu/tier complexity. A two-option menu is a strong default.

- If you have more than 3 pricing tiers, you’re paying a complexity tax for marginal robustness gain.
- A two-option menu is provably near-optimal across support, mean, and quantile information regimes.
- Where this lands: SaaS pricing, B2B contract structures, shipping/service-level options, retail bundling.

Insight 2

Design for the Data You Actually Observe

When you can’t see individual buyers, the right mechanism depends on which aggregate statistic you can observe.

Robust Mechanism Design with Anonymous Information (2026)
Multi-Item Screening with a Maximin-Ratio Objective (2025)

Procurement Reality: The Data You Don’t Have

Sealed-bid procurement: you only see the winning bid

English auctions: winner never reveals true max valuation

Ad exchanges, eBay, real estate: bidder IDs anonymized by design

GDPR, CCPA: bidder-level data is increasingly regulated

False-name bidding makes individual data unreliable anyway

Aggregate data isn’t a weakness — it’s the modern reality.

Two Versions of the Same Question

What if you only see...

Highest Bid Distribution

From posted-price experiments or sealed-bid auctions

E-commerce A/B pricing tests

What if you only see...

Intermediate Order Statistic

Second-price auctions on ad exchanges

Google Ad Exchange winning prices

For each: what's the provably best mechanism?

Result 1: Highest Bid Known → Post a Price

Optimal mechanism: **Posted price calibrated to the highest-bid distribution.**

REAL-WORLD SCENARIOS

E-commerce A/B Pricing

$29	$39	$49	$59

Test prices → observe purchase rates

✓ **Optimize the posted price**

SaaS Trial Conversion

Highest converters → max-bid signal

✓ **Single posted price wins**

Ride-Share Surge

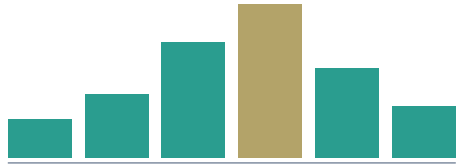

Peak demand = max willingness

✓ **Time-varying posted price**

Result 2: Second-Price Data → 2nd-Price with Reserve

Optimal mechanism: **Second-price auction with optimal reserve.**

The implied i.i.d. distribution is regular above its reserve — second-price-with-reserve is robustly optimal.

REAL-WORLD SCENARIOS

Programmatic Ad Auction

AD SLOT

Bids: $2.40, $2.10, $1.80

Winner pays $2.10 (2nd-price)

✓ Add reserve from data

Online Marketplace

Bid history visible:

$220 (winner)
$215 (winner pays)
$200, $195, …

✓ Tune reserve from clearing

B2B Vendor Selection

4 sealed bids → keep winner + runner-up

Historical 2nd-bid data

✓ Sets reserve, keeps fairness

Privacy & Fairness Are Built In, Not Bolted On

CAUSE

The observable data is an aggregate (e.g., highest bid) — symmetric across bidders by construction.

FORCES

Any robustly optimal mechanism must treat every bidder identically — it can't condition on what it can't see.

DELIVERS THREE PROPERTIES SIMULTANEOUSLY

Privacy

No bidder-level data used or required

Fairness

Identical treatment of every bidder

Revenue Optimal

Worst-case best across all consistent priors

Pivot: Multi-Product Pricing Under Limited Info

New question: You sell multiple products. Should you sell separately, bundle, or mix?

The hard part: even with full information, the optimal mechanism for n products can require an infinite menu of lotteries.

What if you only know value ranges for each product?

✔ **Separate Selling**

Price each item independently with randomized prices

$p_1 = 2$

$p_2 = 2$

vs.

? **Bundling?**

Sell everything together at one price

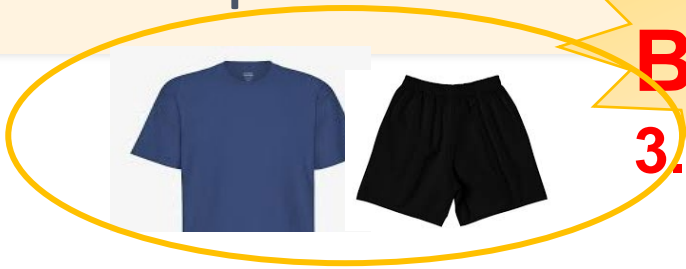

Buy both at 3.

Multi-Item Result: Separate Selling Wins

Theorem: With only support information (value ranges) for each product, the robustly optimal mechanism is **separable.**

Simple to Implement

Post randomized prices for each item independently — no mixed-bundling complexity

Exploits Joint Info

Tailors each item's price distribution using cross-item support data

Provably Optimal

Strictly beats robust mechanisms that ignore cross-item value ranges

Insight 2: Takeaway

Match your mechanism to your data. Privacy compliance can be revenue-optimal.

You Observe	Optimal Mechanism	Industry Example
Highest bid distribution	**Posted price**	E-commerce pricing tests
Intermediate order statistic	**2nd-price with reserve**	Ad exchanges (Google AdX)
Product value ranges only	**Separate randomized pricing**	B2B catalogs, parts pricing
Total budget known	**Randomized bundling**	Enterprise software suites

Insight 3

The Buyer Knows More Than You Do — Use It

Stop fighting buyer information asymmetry — turn it into a resource.

Pricing with a Hidden Sample (2026)
Optimal Pricing with Unreliable Signals (2026)

The buyer may know the market better than the seller.

Seller
sees one hidden sample

Buyer
knows market context

Cargill®

The Cold-Start Problem

New SKU Launch

No historical demand data available

Fast Fashion

Distribution shifts before you collect enough samples

New Market

Legacy data doesn't transfer to new geography

Dynamic B2B

Each contract is unique; sparse observations

Sometimes you have just ONE past transaction. Can you run sophisticated robust pricing with a single sample?

Yes.

Bridging Two Paradigms

Statistic-Based Pricing

Uses mean, variance, quantiles

- ✓ Strong guarantees
- ✗ Needs ~100+ samples to estimate

Bridge

Sample-Based Pricing

Uses raw observed samples

- ✓ Works with few samples
- ✗ Typically weaker guarantees

Hidden Sample Pricing

Uses buyer's knowledge to bridge both paradigms with ONE sample

Hidden Sample: The Mechanism in One Picture

1

Draw Sample

Seller privately observes one sample s from past data

2

Commit Rule

Seller announces pricing rule h(s) — but keeps s hidden

3

Buyer Decides

Buyer evaluates expected price E[h(s)] using her distribution knowledge

4

Reveal & Charge

If buyer participates, sample is revealed and h(s) is charged

The buyer's expected price depends on the distribution — which she knows but the seller doesn't.

Delayed Public Signal.

Contract
h(s)

Commodity index s
Another transaction
price

From Buyer Knowledge to AI Signals

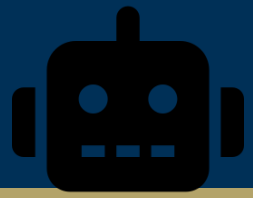

So far: the buyer's knowledge is the seller's resource.
Now: what if the seller has AI signals about the buyer?

Twist: AI signals can be highly informative
or completely hallucinated — and you don't know which.

The Unreliable AI Signal Problem

Accurate Signal

s = v (buyer's true valuation)

The AI/LLM captured accurate info about this buyer's preferences.

Hallucinated Signal

s ~ F (random draw from prior)

The AI had no real info — generated a generic guess.

The seller doesn't know which regime applies. **But the buyer does** — she knows how much data she shared.

The Consistency-Robustness Trade-Off

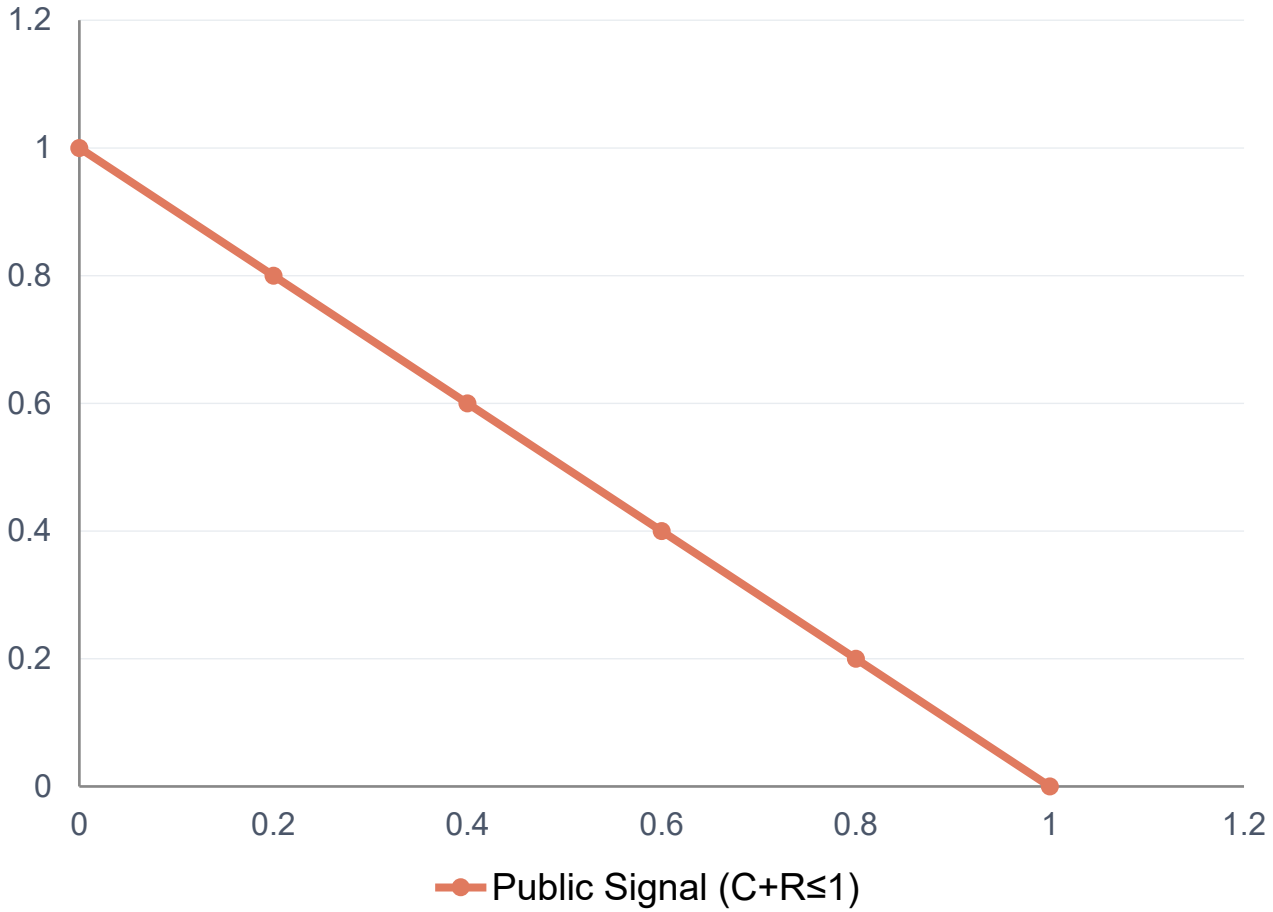

Consistency (C):
How well you do when the AI signal is accurate

Robustness (R):
How well you do when the AI signal is hallucinated

Naive baseline:
Randomize between using signal and ignoring it
Achieves C + R ≤ 1

Every gain in C costs you in R, one-for-one.

The Counterintuitive Move: Keep the Signal Private

Don’t reveal the AI signal to the buyer.

If signal is accurate, buyer already knows (s = v)

If signal is hallucinated, buyer ignores it

Seller doesn’t know which → buyer’s behavior under this uncertainty creates a wedge the seller can exploit

Headline: Strict Improvement Over Any Public Mechanism

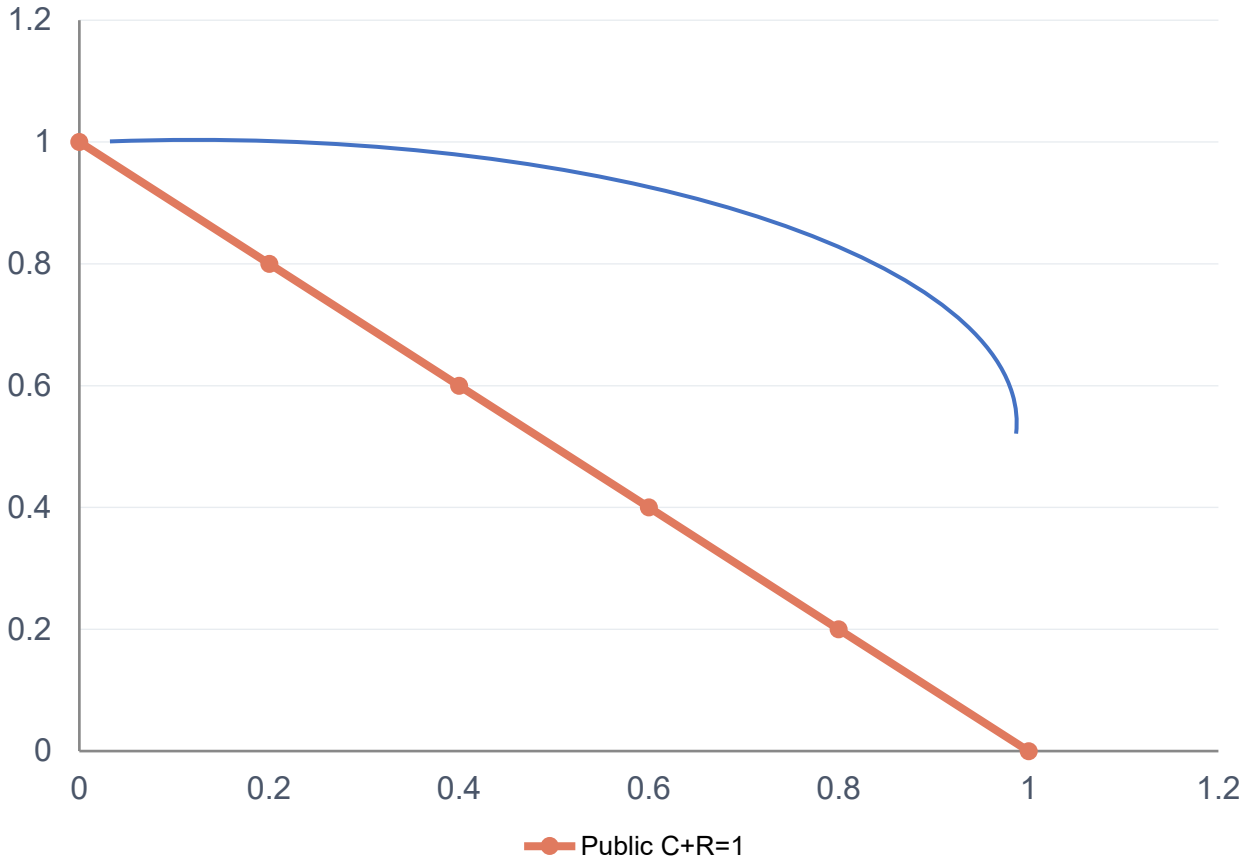

Private > Public

The Pareto frontier shifts outward.

You get higher consistency AND higher robustness simultaneously.

This is free value from keeping your AI signal private.

Perfect Consistency + Meaningful Robustness

C = 1

Perfect

Full surplus extraction when the AI signal is accurate

R ≥ ½

Protected

At least 50% of optimal no-signal revenue under hallucination

You don’t have to give up the upside to protect against AI hallucination.

When You Get Both: 1-Consistent and 1-Robust

Special cases: if the prior has infinite mean OR mean ≤ monopoly price, there exists a mechanism that is simultaneously **1-consistent AND 1-robust.**

Full revenue under accurate signal + full no-signal revenue under hallucination. Best of both worlds.

Heavy-Tailed Luxury

Jewelry, art, collectibles — valuations can be extreme (infinite mean)

Commoditized Products

Low-margin items where mean ≤ monopoly price holds naturally

Long-Tail Markets

Digital goods, online services with Pareto-distributed valuations

Where AI-Signal Pricing Lands

Each scenario: AI signal feeds the seller — but is kept private from the buyer

LLM Personalization

"This buyer is high-value"

↓ private ↓

Mechanism

Chatbot → seller signal

✓ Don't reveal to buyer

Recommendations

Similar users

Avg paid: $42

Hallucinated when too few users → garbage signal

✓ Robust to bad lookups

Third-Party Data

Buyer score from vendor

Quality varies per vendor

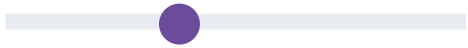

low ← reliability → high

✓ Hedge unreliable vendors

A New Mechanism Design Paradigm

Traditional

Design mechanisms based on what the seller knows

Our Framework

Design mechanisms that leverage what the other side knows about the seller's information

When integrating AI, treat signals as private inputs whose value comes from information asymmetry.

Insight 3: Takeaway

A single transaction is more informative than you think. Keep AI signals private.

- A single transaction can drive sophisticated pricing — if the mechanism elicits the buyer's knowledge.
- Treat AI/ML signals as private inputs, not as prices to publish to the buyer.
- You can have both AI-driven precision and robustness against hallucination — they aren't a strict trade-off.
- Applications: new product launches, dynamic pricing, AI-assisted procurement, B2B value-based pricing.

Five Takeaways for Monday Morning

1. Audit your menu/tier complexity. More than 3 tiers is usually a complexity tax.
2. Match your mechanism to your data. Posted pricing, 2nd-price-with-reserve, separate randomized pricing — each provably right for a specific data regime.
3. Privacy & fairness compliance can be revenue-optimal, not a constraint.
4. Treat AI buyer signals as private inputs, not as prices to publish.
5. A single transaction can drive sophisticated pricing — if designed to elicit buyer knowledge.

Applications

Procurement Auction Redesign

Privacy/fairness/regulatory constraints — bring me your aggregate clearing-price data

Cold-Start Pricing

New products with ≤10 historical transactions — bring me a launch scenario

AI-Assisted Pricing

Signal reliability is uncertain — bring me your LLM/recommender output and ground-truth data

Shixin Wang | shixin.wang@isye.gatech.edu

Open to practicum projects, case studies, and joint applied research

Thank You

Questions & Discussion

"What's the pricing decision keeping you up at night?"

Shixin Wang | shixin.wang@isye.gatech.edu
ISyE, Georgia Institute of Technology

Co-authors: Zhihao Gavin Tang, Yixin Tao, Jiawei Zhang

Upcoming Courses

Principles of Transportation Management

June 1-3, 2026 | Virtual (Instructor-led)

Machine Learning Applications for Supply Chain Planning

Sept 14-17, 2026 | Virtual (Instructor-led)

Generative AI Application for Supply Chain Professionals

October 19-21, 2026 | Georgia Tech Savannah Campus

Georgia AIM Discount

All Georgia residents can receive 50% off our Professional Education courses
— but only until August 31.

Thanks to a grant from the U.S. Department of Commerce's Economic Development Administration, this limited-time offer supports Georgia learners in gaining career-boosting skills.

✓ GA residency verification required
✓ Use code *SCL-GAAIM* at checkout
✓ Cannot be combined with other discounts

Supply Chain and Logistics Institute

Upcoming SCL Lunch and Learn Opportunities

Ahead of the Curve:
Turning Analysis Into Executable Decisions

w/ Chris Gaffney

Thursday, June 4th | 12-1pm ET | Zoom Registration Link

scl.gatech.edu/jun26-lnl

Supply Chain and Logistics Institute

Supply Chain and Logistics Institute

www.scl.gatech.edu

Follow us

Subscribe to our Newsletter