Adventures in Demand Analysis Using AI

Philipp Bach, Victor Chernozhukov, Sven Klaassen, Martin Spindler, Jan Teichert-Kluge, Suhas Vijaykummar

13th ECB Conference on Forecasting Techniques
ECB, 23/3/2026

Discussion by **Dimitris Georgarakos** (ECB)

Disclaimer

The views expressed in the presentation are my own and do not reflect the views of the European Central Bank or of the euro system

The paper in a 'nutshell'

Research Question

- *Can AI-generated product embeddings improve demand estimation and causal inference?*

Data

- Amazon toy car products (≈9,600 items)
- Multimodal inputs: sales ranks, prices, *text*, *images* (captures quality, branding etc.)
- Sales rank used as proxy for (relative) quantity

Methodology

- Transformer-based embeddings (text and images) for predicting price and demand
- Dynamic panel model + Double Machine Learning (DML)
- Estimate price elasticity (average and heterogeneous)

Paper: key findings and contribution

Key Findings

- AI embeddings: strong determinants/ modifiers of the price elasticity function
- Embeddings act mainly as **effect modifiers**, not confounders
- Strong heterogeneity in price elasticity across products

Contribution

- Integrates AI representations into demand analysis
- Links between ML and causal inference

Endogeneity

Endogeneity of Prices:

- Prices may respond to demand shocks
    - “Sticky prices” argument
    - Lagged controls
- No credible instrumental variable; dynamic controls may not fully solve simultaneity
- Latent, time-varying factors could bias elasticity estimates

Implication:

- Results are **causal under relatively strong assumptions** (otherwise: high-quality correlations)

Suggestion:

- Explore (alternative) IV strategies and sensitivity analysis:
    - **Omitted variable bias** bounds (e.g., Cinelli and Hazlett (2020): strength of unobserved confounders based on observed covariates; Oster (2019): coefficient stability)
    - Richness of the embeddings might actually help tighten **Manski-type bounds** (e.g., use in MIV: *potential outcomes* are ordered in a monotone way)

Embeddings

What do embeddings capture?

- Quality? Branding? Visibility? Unobservables (incl. consumer psychological outlook)?

Is there an ‘optimal’ number/ type of embeddings?

- Customers’ rating, feedback?

‘Black box’ concern

- Embeddings: not strong confounders, but strong **effect modifiers**
- Economic interpretability?
- Risk of “black box heterogeneity”
- Hard to translate findings into:
  - Policy insights
  - Structural demand models

Suggestion

- Decompose embeddings (e.g., interpretable dimensions)
- Link clusters more explicitly to economic characteristics (**matched/ combined datasets**: sellers, consumers, etc.)

External validity

- How these embeddings **generalize across markets** or **platforms?**
  - May encode platform equilibrium (incl. consumer tastes, seller strategies and ranking algorithms)/ customer **selection** based on unobservables
  - **Amazon:** (relatively) stable prices, quality focused, algorithmic ranking **vs. Temu** (heavy promotions, dynamic discounts, gamified pricing)
- sales rank → quantity transformation (relies on power law assumption)
- Rank ≠ quantity (nonlinear, noisy mapping); ME likely non-trivial
- Estimated elasticities sensitive to distributional assumptions?

Suggestion

•Validate with external data that provide some direct information on prices, quantity *and* key product characteristics (e.g., **scanner data**?)

•Sensitivity to alternative mappings

Overview

- **Important** and **innovative** paper
- AI can significantly enhance demand estimation
- some challenges:
    - Causal inference
    - Interpretability
- Are embeddings a substitute for economic structure or a complement?

Thank you!

dimitris.georgarakos@ecb.europa.eu