Market Power in Artificial Intelligence

Centre for Competition Economics
October 2024

Joshua Gans, University of Toronto

Technologies usually improve competition

- transportation
- energy
- information technology
- the Internet

Technologies usually improve competition

- transportation
- energy
- information technology
- the Internet

But THIS TIME IT'S DIFFERENT

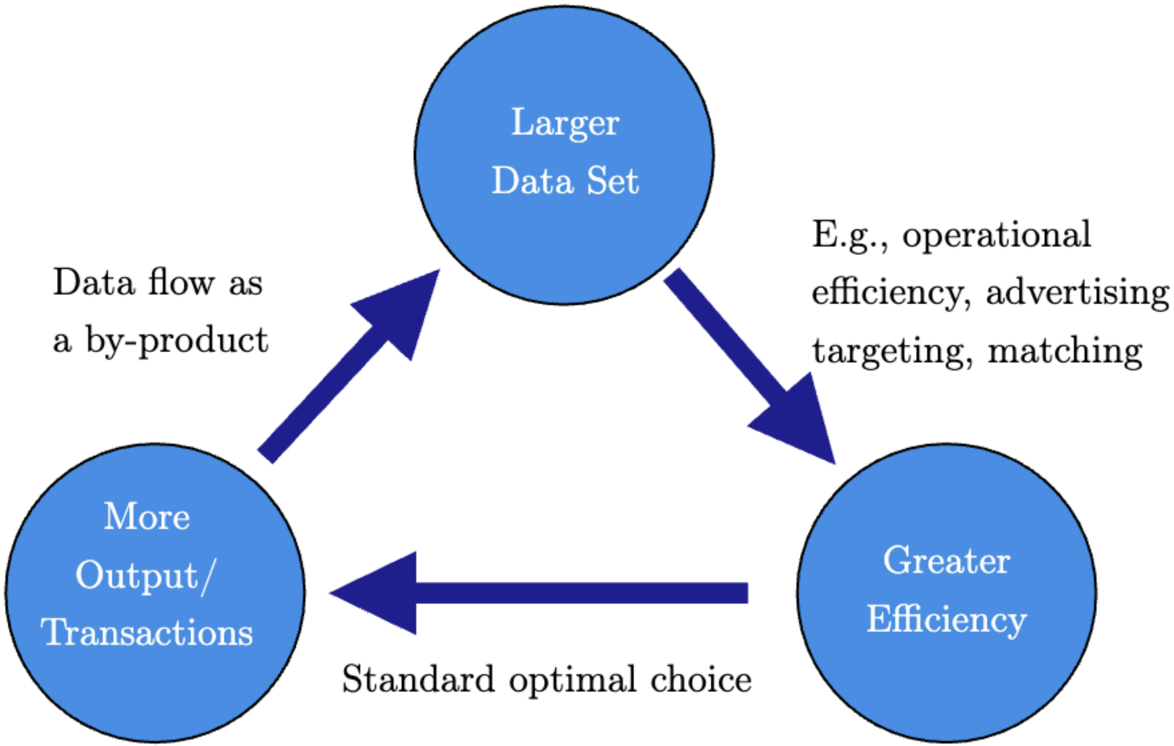

Figure 16.1: The Data Feedback Loop

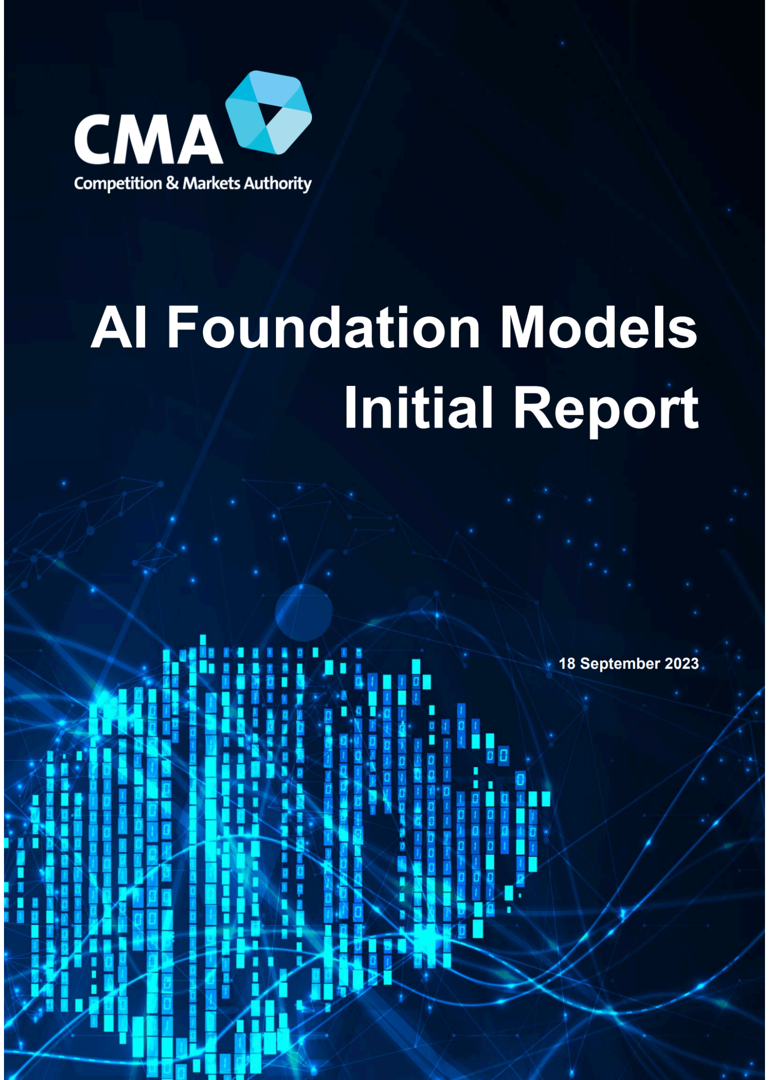
CMA
Competition & Markets Authority
AI Foundation Models
Initial Report
18 September 2023

Outline

Outline

- Calibration: What is AI?

Outline

- Calibration: What is AI?
- Algorithmic Collusion

Outline

- Calibration: What is AI?
- Algorithmic Collusion
- Market Power in Training Data

Outline

- Calibration: What is AI?
- Algorithmic Collusion
- Market Power in Training Data
- Market Power in Input Data

Outline

- Calibration: What is AI?
- Algorithmic Collusion
- Market Power in Training Data
- Market Power in Input Data
- Market Power in Prediction

Outline

- Calibration: What is AI?
- Algorithmic Collusion
- Market Power in Training Data
- Market Power in Input Data
- Market Power in Prediction
- Diverse Business Models

Outline

- Calibration: What is AI?
- Algorithmic Collusion
- Market Power in Training Data
- Market Power in Input Data
- Market Power in Prediction
- Diverse Business Models
- Diverse Governance

What is AI?

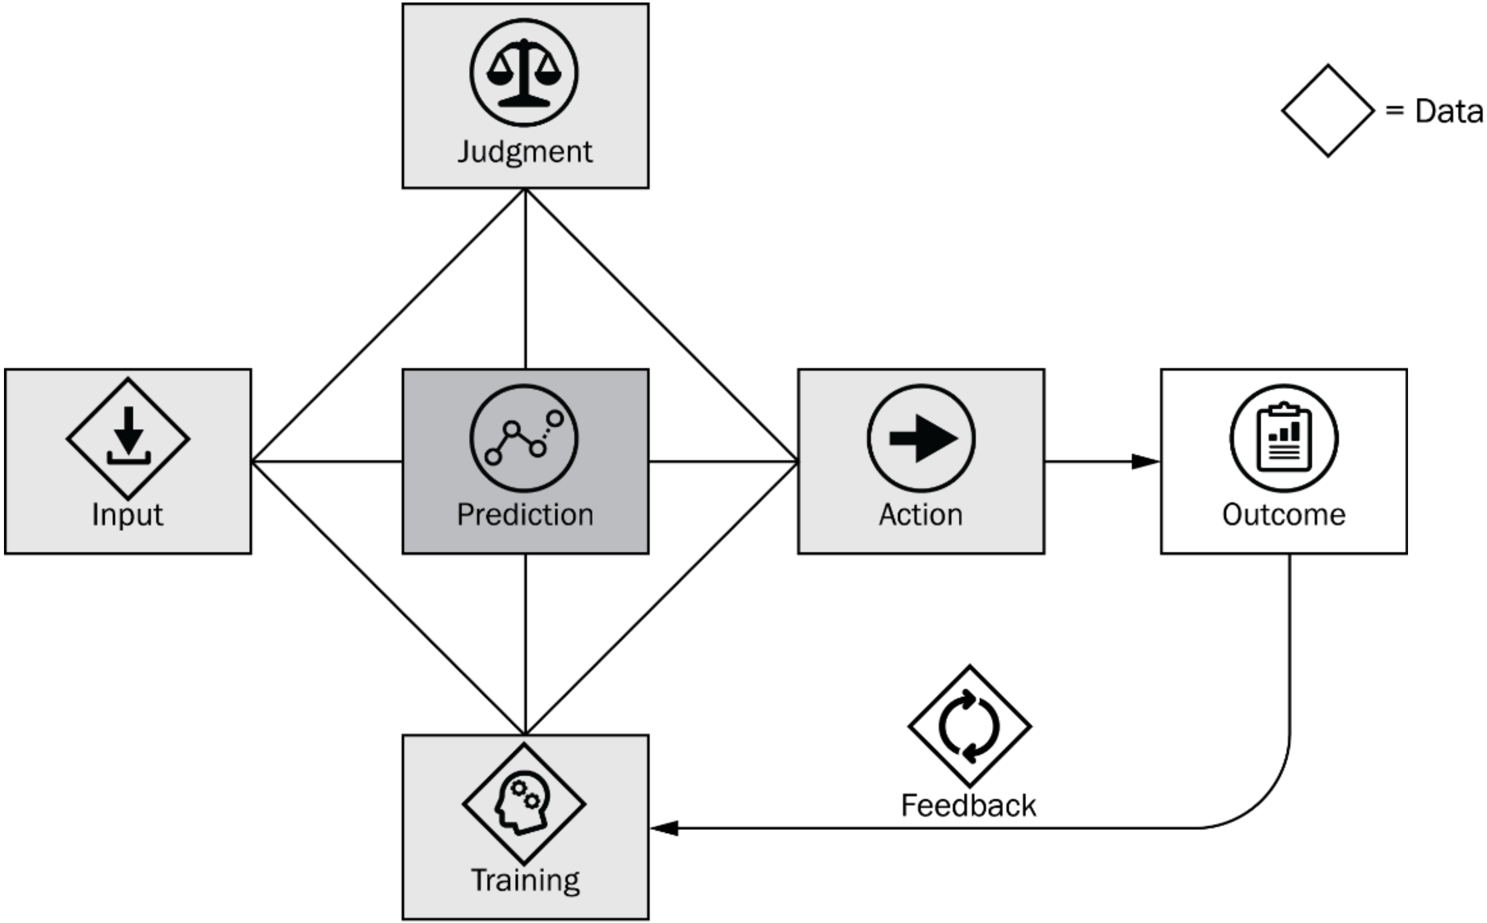
Judgment
= Data
Input
Prediction
Action
Outcome
Feedback
Training

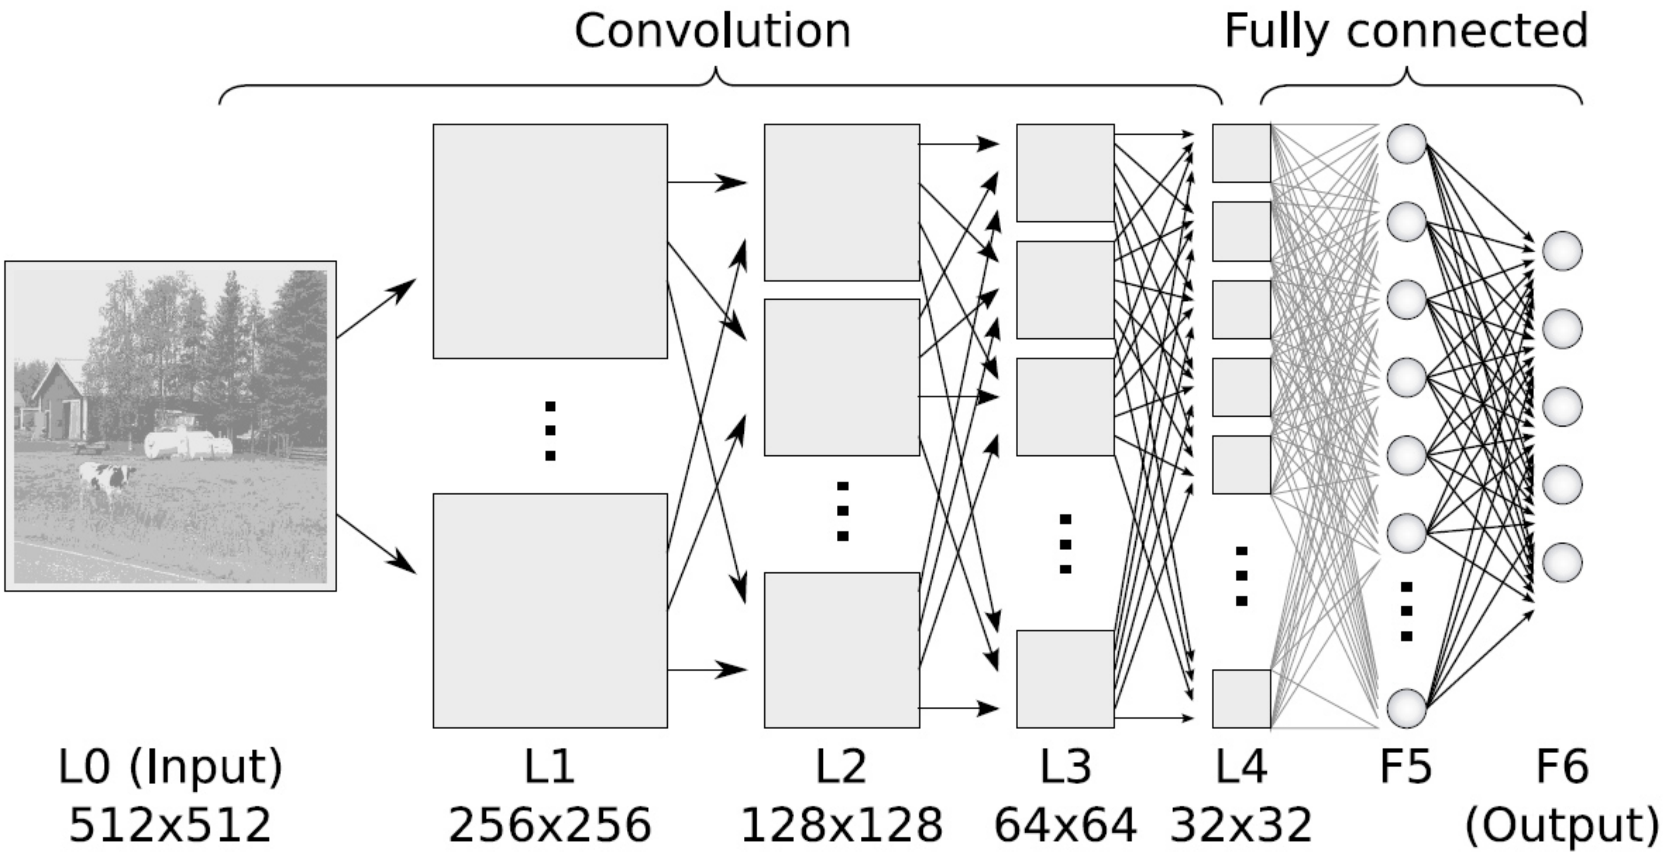

Source: https://www.ais.uni-bonn.de/deep_learning/

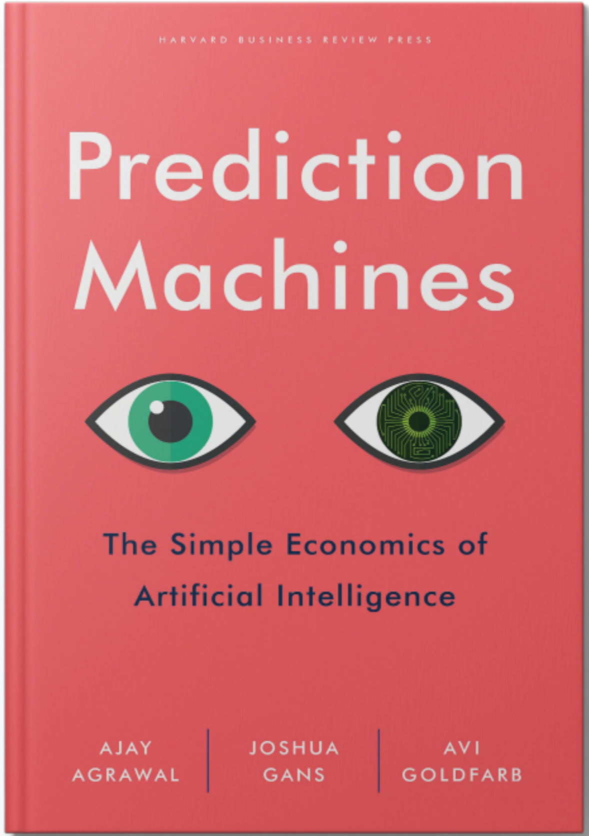
HARVARD BUSINESS REVIEW PRESS
Prediction Machines
The Simple Economics of Artificial Intelligence
AJAY AGRAWAL
JOSHUA GANS
AVI GOLDFARB

Artificial Intelligence

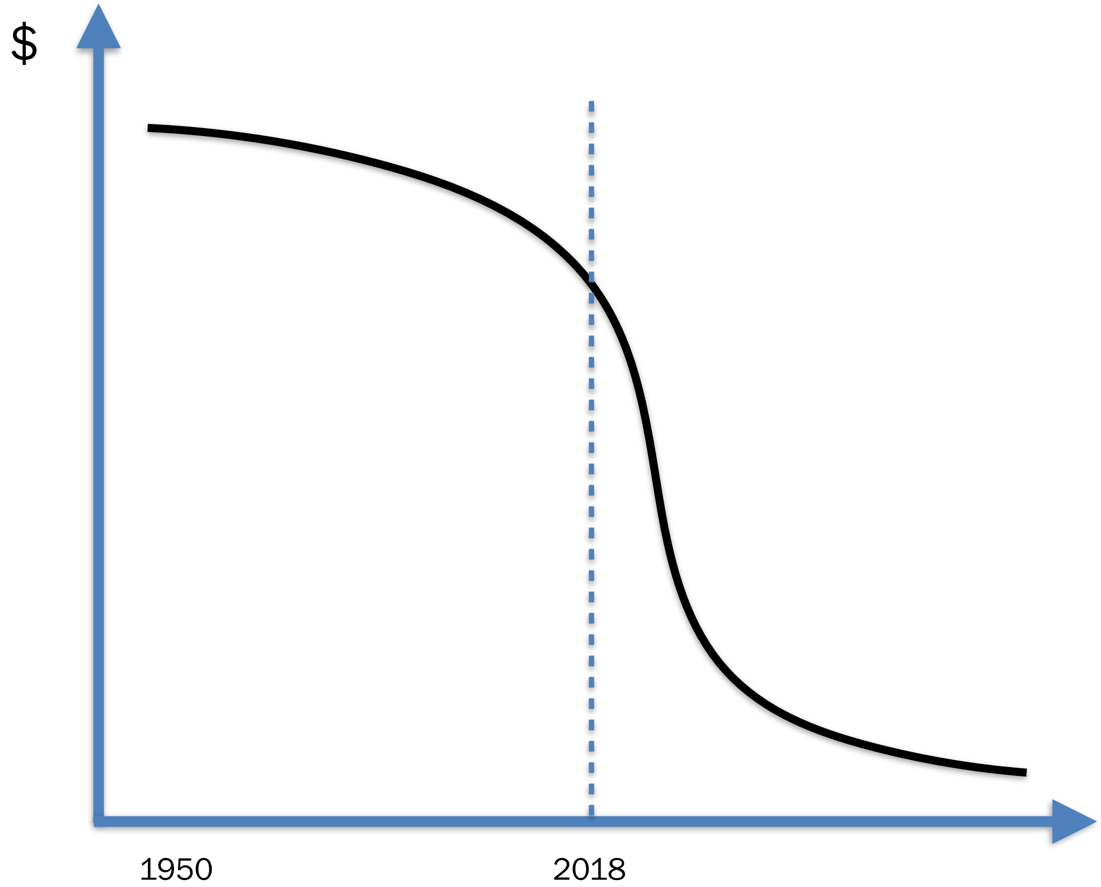

Artificial Intelligence

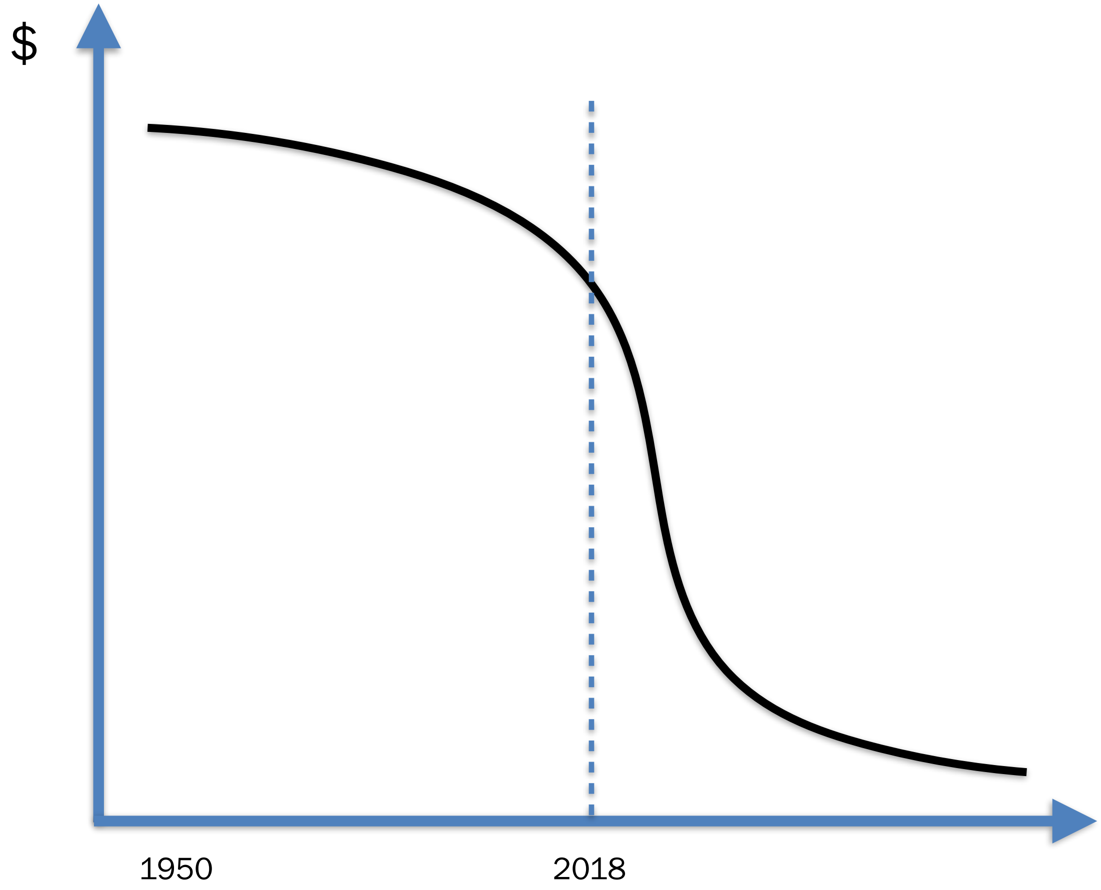

Artificial Intelligence

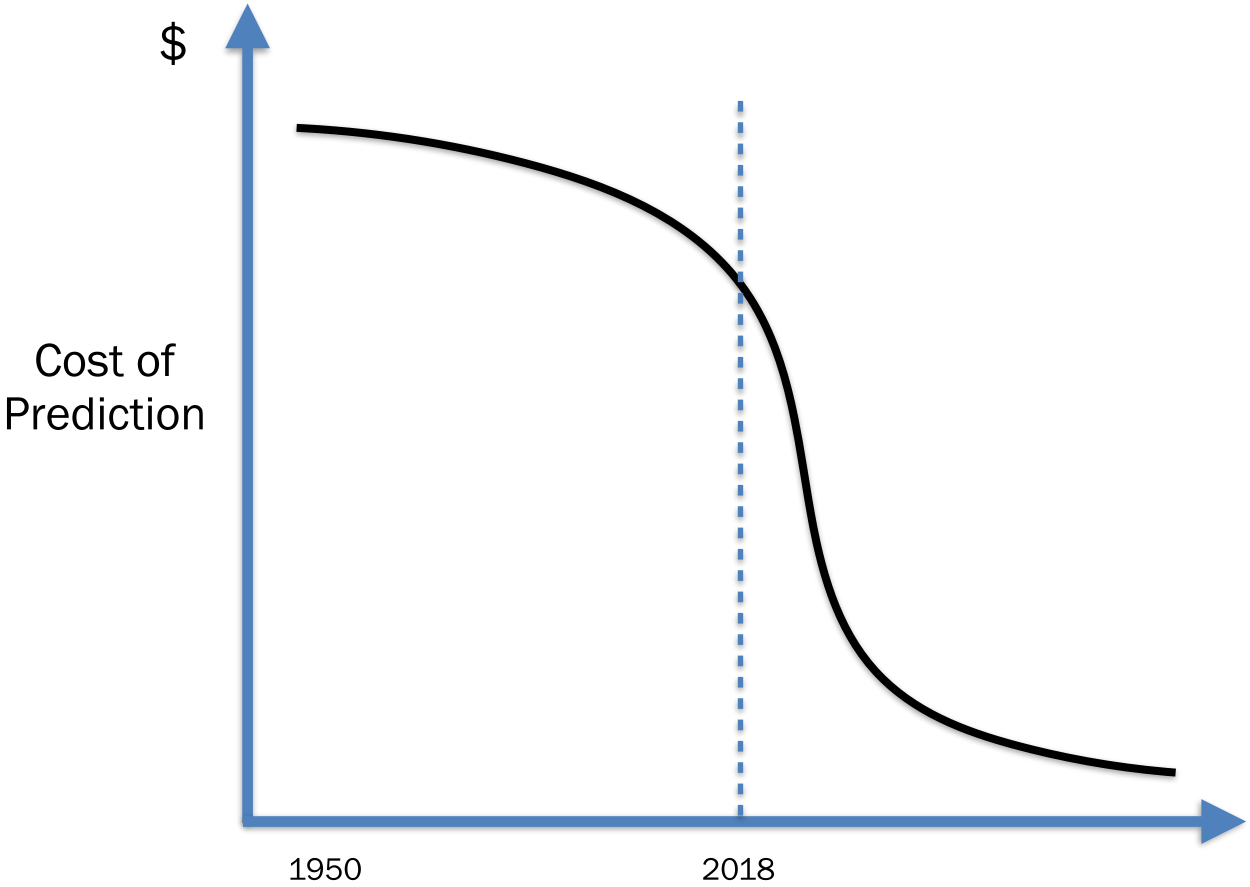

PREDICTION:

Using information that you *do* have to generate information that you *don't* have

Market Power in Training Data

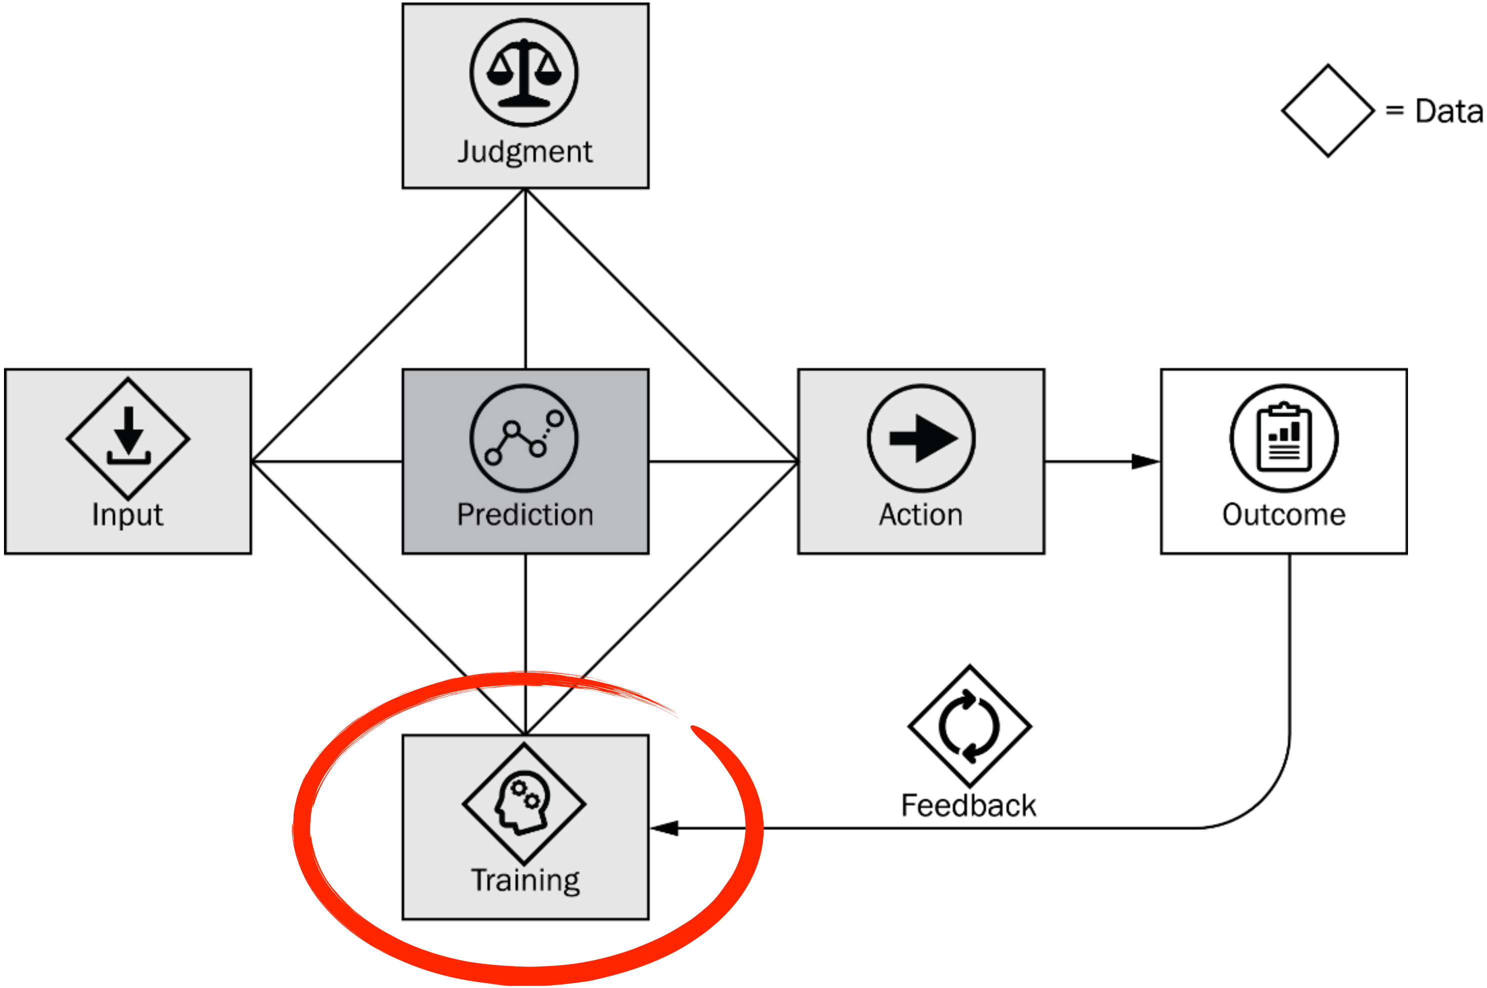
= Data
Judgment
Input
Prediction
Action
Outcome
Feedback
Training

Firms' incentives to sell training data are constrained by fears of creative destruction

American Economic Review 2020, 110(9): 2819–2858
https://doi.org/10.1257/aer.20191330

Nonrivalry and the Economics of Data[†]

By Charles I. Jones and Christopher Tonetti*

Data is nonrival: a person's location history, medical records, and driving data can be used by many firms simultaneously. Nonrivalry leads to increasing returns. As a result, there may be social gains to data being used broadly across firms, even in the presence of privacy considerations. Fearing creative destruction, firms may choose to hoard their data, leading to the inefficient use of nonrival data. Giving data property rights to consumers can generate allocations that are close to optimal. Consumers balance their concerns for privacy against the economic gains that come from selling data broadly. (*JEL* C80, D11, D21, D83, E22, K11, O34)

In recent years, the importance of data in the economy has become increasingly apparent. More powerful computers and advances in algorithms such as machine learning have led to an explosion in the usefulness of data. Examples include self-driving cars, real-time language translation, medical diagnoses, product recommendations, and social networks.

This paper develops a theoretical framework to study the economics of data. We are particularly interested in how different property rights for data determine its use in the economy, and thus affect output, privacy, and consumer welfare. The starting point for our analysis is the observation that data is nonrival. That is, at a technological level, data is infinitely usable. Most goods in economics are rival: if a person consumes a kilogram of rice or an hour of an accountant's time, some resource with a positive opportunity cost is used up. In contrast, existing data can be used by any number of firms or people simultaneously, without being diminished. Consider a collection of one million labeled images, the human genome, the US Census, or the data generated by 10,000 cars driving 10,000 miles. Any number of firms, people, or machine learning algorithms can use these data simultaneously without reducing the amount of data available to anyone else.

Firms' incentives to sell training data are constrained by fears of creative destruction

American Economic Review 2020, 110(9): 2819–2858
https://doi.org/10.1257/aer.20191330

Nonrivalry and the Economics of Data†

By Charles I. Jones and Christopher Tonetti*

Data is nonrival: a person's location history, medical records, and driving data can be used by many firms simultaneously. Nonrivalry leads to increasing returns. As a result, there may be social gains to data being used broadly across firms, even in the presence of privacy considerations. Fearing creative destruction, firms may choose to hoard their data, leading to the inefficient use of nonrival data. Giving data property rights to consumers can generate allocations that are close to optimal. Consumers balance their concerns for privacy against the economic gains that come from selling data broadly. (*JEL* C80, D11, D21, D83, E22, K11, O34)

In recent years, the importance of data in the economy has become increasingly apparent. More powerful computers and advances in algorithms such as machine learning have led to an explosion in the usefulness of data. Examples include self-driving cars, real-time language translation, medical diagnoses, product recommendations, and social networks.

This paper develops a theoretical framework to study the economics of data. We are particularly interested in how different property rights for data determine its use in the economy, and thus affect output, privacy, and consumer welfare. The starting point for our analysis is the observation that data is nonrival. That is, at a technological level, data is infinitely usable. Most goods in economics are rival: if a person consumes a kilogram of rice or an hour of an accountant's time, some resource with a positive opportunity cost is used up. In contrast, existing data can be used by any number of firms or people simultaneously, without being diminished. Consider a collection of one million labeled images, the human genome, the US Census, or the data generated by 10,000 cars driving 10,000 miles. Any number of firms, people, or machine learning algorithms can use these data simultaneously without reducing the amount of data available to anyone else.

beneficiaries of business stealing: they obtain the property rights to the varieties that suffer from creative destruction.[3] Then, the free entry condition is

$$\chi w_t = V_{it} + \frac{\int_0^{N_t} \delta(\tilde{x}_{it}) V_{it}\, di}{\dot{N}_t}. \tag{48}$$

The left side χw_t is the cost of the χ units of labor needed to create a new variety. The right side has two terms. The first is the value of the new variety that is created. The second is the per-entrant portion of the rents from creative destruction.

Market Power in Input Data

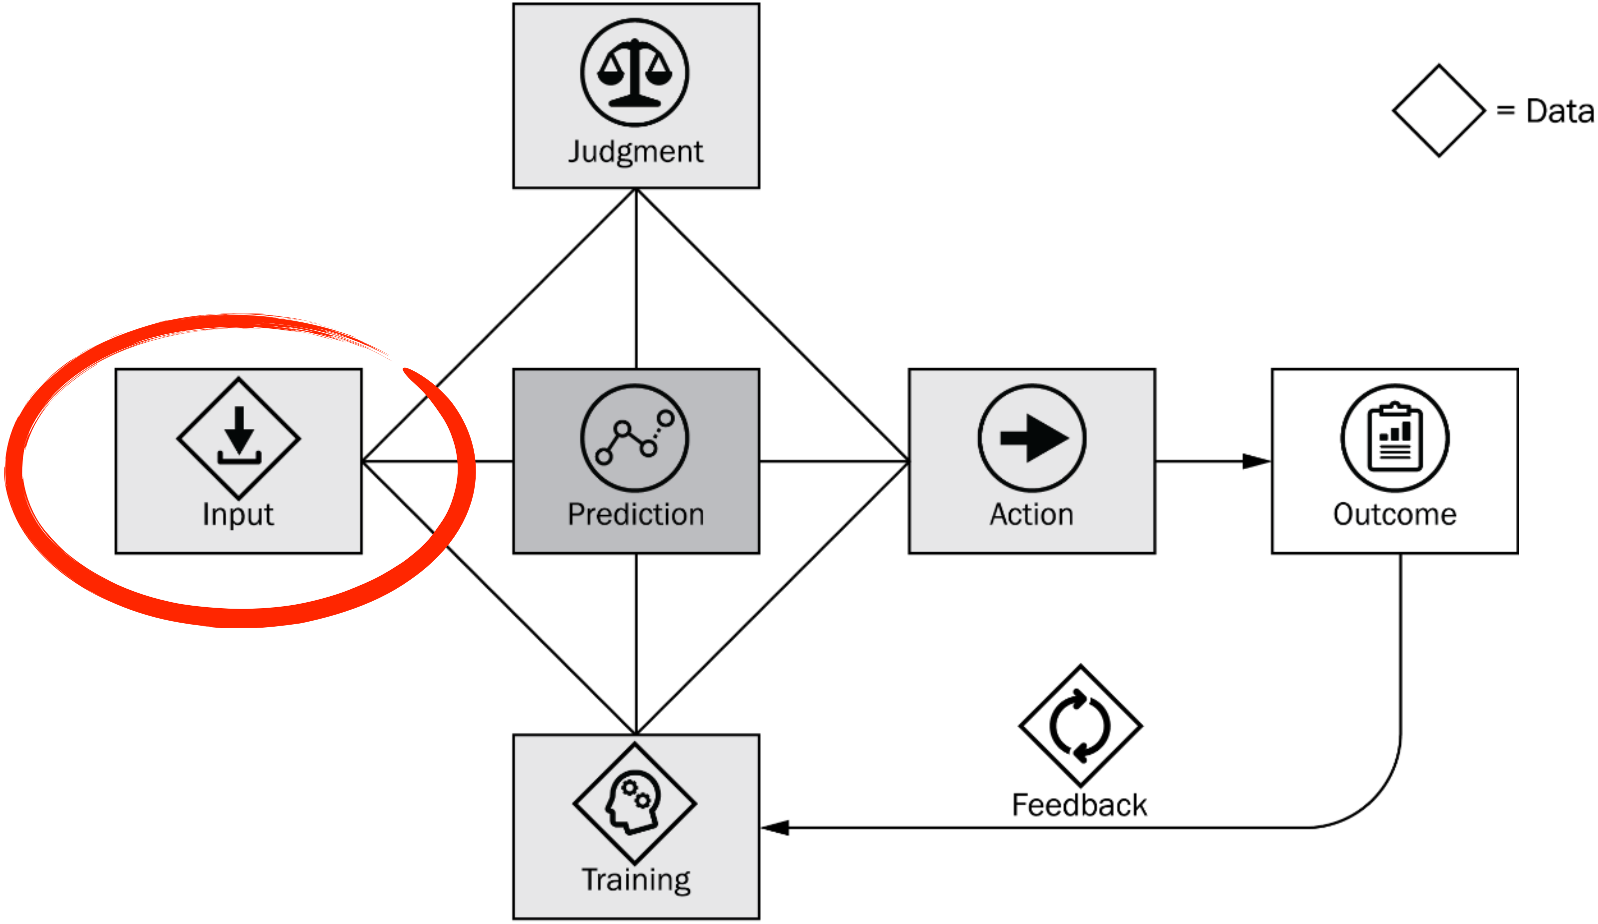
Judgment
= Data
Input
Prediction
Action
Outcome
Feedback
Training

Literature on Sharing Information Under Oligopoly

Literature on Sharing Information Under Oligopoly

- Sharing input data does not change market structure but the intensity of competition within a market

Literature on Sharing Information Under Oligopoly

- Sharing input data does not change market structure but the intensity of competition within a market
- Vives (1984) and Raith (1996): firms might share voluntarily

Literature on Sharing Information Under Oligopoly

- Sharing input data does not change market structure but the intensity of competition within a market
- Vives (1984) and Raith (1996): firms might share voluntarily
- Jansen (2008)

Literature on Sharing Information Under Oligopoly

- Sharing input data does not change market structure but the intensity of competition within a market
- Vives (1984) and Raith (1996): firms might share voluntarily
- Jansen (2008)
  - Precommiting to share input data (on, say, demand)

Literature on Sharing Information Under Oligopoly

- Sharing input data does not change market structure but the intensity of competition within a market
- Vives (1984) and Raith (1996): firms might share voluntarily
- Jansen (2008)
    - Precommiting to share input data (on, say, demand)
    - Without pre-commitment, sharing is reduced

Literature on Sharing Information Under Oligopoly

- Sharing input data does not change market structure but the intensity of competition within a market
- Vives (1984) and Raith (1996): firms might share voluntarily
- Jansen (2008)
  - Precommiting to share input data (on, say, demand)
  - Without pre-commitment, sharing is reduced
  - Sharing is reduced the greater the degree of market power

Literature on Sharing Information Under Oligopoly

- Sharing input data does not change market structure but the intensity of competition within a market
- Vives (1984) and Raith (1996): firms might share voluntarily
- Jansen (2008)
  - Precommiting to share input data (on, say, demand)
  - Without pre-commitment, sharing is reduced
  - Sharing is reduced the greater the degree of market power
- de Cornière-Taylor (2023): general characterisation of impact of data on reaction curves and implications for structural change

Market Power in Prediction

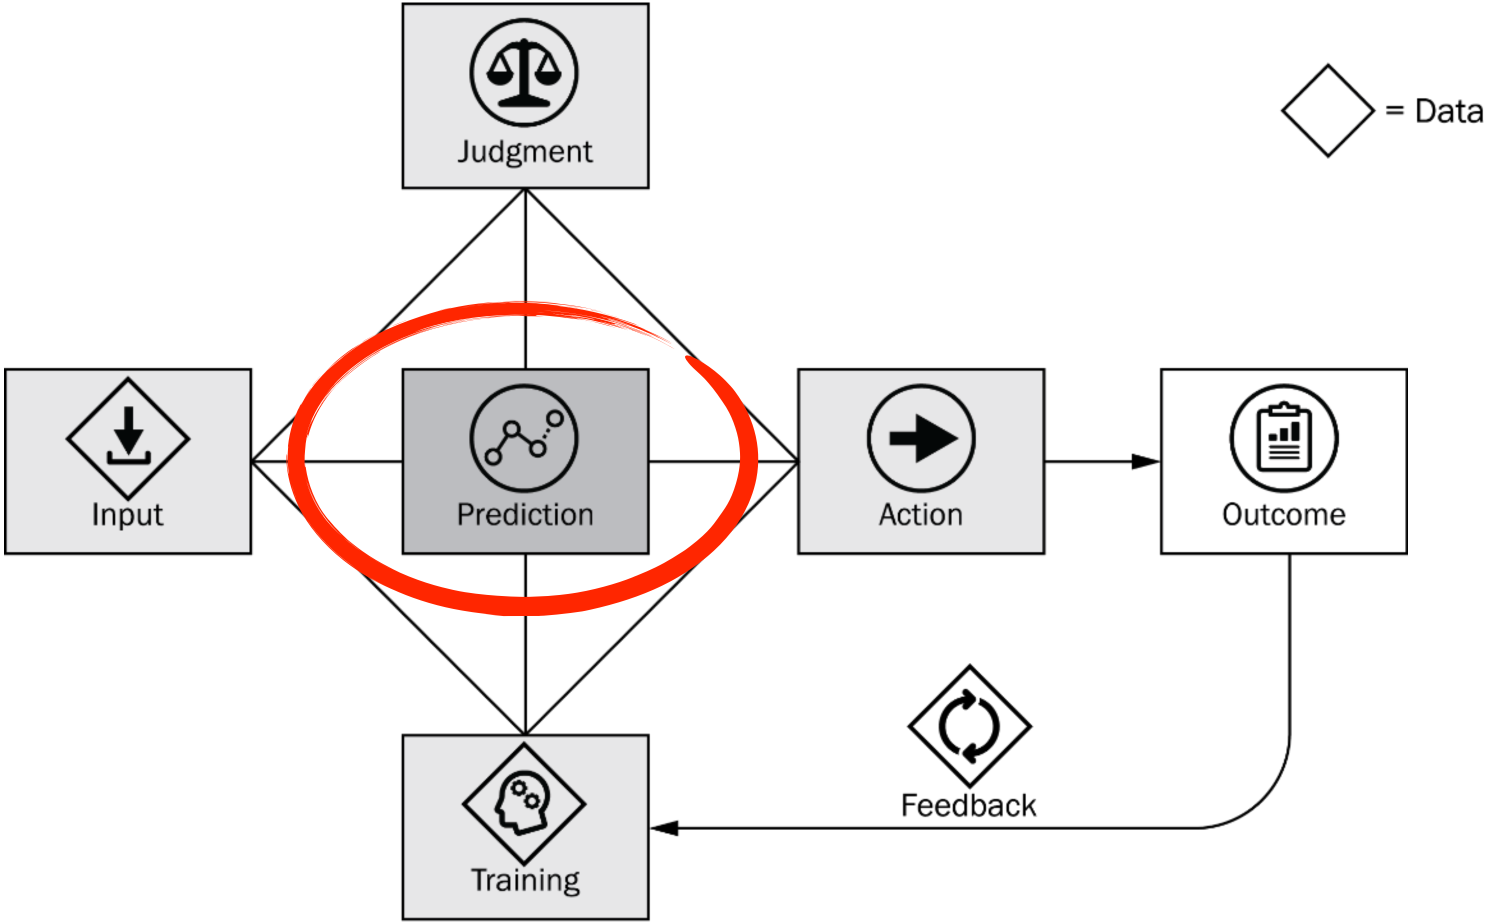
= Data
Judgment
Input
Prediction
Action
Outcome
Feedback
Training

Advertising: Predictions of Match Quality

Advertising: Predictions of Match Quality

- Bergemann and Bonatti (2015, AEJ: Micro) "Selling Cookies"

Advertising: Predictions of Match Quality

- Bergemann and Bonatti (2015, AEJ: Micro) "Selling Cookies"
    - Cookies are basically predictions or key input data for predicting match quality

Advertising: Predictions of Match Quality

- Bergemann and Bonatti (2015, AEJ: Micro) “Selling Cookies”
  - Cookies are basically predictions or key input data for predicting match quality
  - Predictions are most valuable the higher are the stakes

Advertising: Predictions of Match Quality

- Bergemann and Bonatti (2015, AEJ: Micro) “Selling Cookies”
  - Cookies are basically predictions or key input data for predicting match quality
  - Predictions are most valuable the higher are the stakes
  - Stakes are low for matching when ad space (limited advertising) is expensive or cheap (want to advertise widely): prediction only alters the “advertise or not” decision for a small number of consumers.

Advertising: Predictions of Match Quality

- Bergemann and Bonatti (2015, AEJ: Micro) “Selling Cookies”
  - Cookies are basically predictions or key input data for predicting match quality
  - Predictions are most valuable the higher are the stakes
  - Stakes are low for matching when ad space (limited advertising) is expensive or cheap (want to advertise widely): prediction only alters the “advertise or not” decision for a small number of consumers.
  - If prediction is provided by fragmented suppliers, then prediction prices increase — akin to the pricing of complements problem.

Bundling of Attention and Advertising

Bundling of Attention and Advertising

- Athey and Gans (2010)

Bundling of Attention and Advertising

- Athey and Gans (2010)
  - Suppose Nk consumers value one of k products so an advertiser has a $1/k$ probability of the ‘right’ match; their expected return per ad is v/k.

Bundling of Attention and Advertising

- Athey and Gans (2010)
    - Suppose Nk consumers value one of k products so an advertiser has a $1/k$ probability of the 'right' match; their expected return per ad is v/k.
    - If platform can predict matches (perfectly), advertisers value an ad at v; v is distributed U[0,1]

Bundling of Attention and Advertising

- Athey and Gans (2010)
  - Suppose *Nk* consumers value one of *k* products so an advertiser has a 1/*k* probability of the 'right' match; their expected return per ad is *v*/*k*.
  - If platform can predict matches (perfectly), advertisers value an ad at *v*; v is distributed U[0,1]
  - Advertiser is indifferent between buying *Nk* ads (w/o prediction) or *N* ads (w/ prediction)

Bundling of Attention and Advertising

- Athey and Gans (2010)
  - Suppose Nk consumers value one of k products so an advertiser has a $1/k$ probability of the 'right' match; their expected return per ad is v/k.
  - If platform can predict matches (perfectly), advertisers value an ad at v; v is distributed U[0,1]
  - Advertiser is indifferent between buying Nk ads (w/o prediction) or N ads (w/ prediction)
  - If there is ad space of a per consumer, platform profits are $(1-a)Na$ (w/o prediction) and $(1-ka)Nka$ (w/ prediction)

Bundling of Attention and Advertising

- Athey and Gans (2010)
    - Suppose Nk consumers value one of k products so an advertiser has a $1/k$ probability of the 'right' match; their expected return per ad is v/k.
    - If platform can predict matches (perfectly), advertisers value an ad at v; v is distributed U[0,1]
    - Advertiser is indifferent between buying Nk ads (w/o prediction) or N ads (w/ prediction)
    - If there is ad space of a per consumer, platform profits are $(1\text{-}a)Na$ (w/o prediction) and $(1\text{-}ka)Nka$ (w/ prediction)
    - Result: if a unconstrained choice, no difference in profits with versus without prediction as $a = 1/2$ (w/o prediction) and $a = 1/2k$ (w/ prediction). If a constrained, it reduces profits w/o prediction. **Prediction is valuable with scarcity**.

Bundling of Attention and Advertising

- Athey and Gans (2010)
    - Suppose Nk consumers value one of k products so an advertiser has a $1/k$ probability of the 'right' match; their expected return per ad is v/k.
    - If platform can predict matches (perfectly), advertisers value an ad at v; v is distributed U[0,1]
    - Advertiser is indifferent between buying Nk ads (w/o prediction) or N ads (w/ prediction)
    - If there is ad space of a per consumer, platform profits are $(1-a)Na$ (w/o prediction) and $(1-ka)Nka$ (w/ prediction)
    - Result: if a unconstrained choice, no difference in profits with versus without prediction as $a = 1/2$ (w/o prediction) and $a = 1/2k$ (w/ prediction). If a constrained, it reduces profits w/o prediction. **Prediction is valuable with scarcity**.
- Prat and Valletti (2022)

Bundling of Attention and Advertising

- Athey and Gans (2010)
  - Suppose *Nk* consumers value one of *k* products so an advertiser has a 1/*k* probability of the 'right' match; their expected return per ad is *v*/*k*.
  - If platform can predict matches (perfectly), advertisers value an ad at *v*; v is distributed U[0,1]
  - Advertiser is indifferent between buying *Nk* ads (w/o prediction) or *N* ads (w/ prediction)
  - If there is ad space of *a* per consumer, platform profits are (1-*a*)*Na* (w/o prediction) and (1-*ka*)*Nka* (w/ prediction)
  - Result: if *a* unconstrained choice, no difference in profits with versus without prediction as *a* = 1/2 (w/o prediction) and *a* = 1/2*k* (w/ prediction). If *a* constrained, it reduces profits w/o prediction. **Prediction is valuable with scarcity**.
- Prat and Valletti (2022)
  - Each of *k* product markets has an incumbent monopolist (earning π_m) and a potential entrant (changes total profits to $2\pi_c$). Entrant needs ad to create awareness.

Bundling of Attention and Advertising

- Athey and Gans (2010)
  - Suppose *Nk* consumers value one of *k* products so an advertiser has a 1/*k* probability of the 'right' match; their expected return per ad is *v*/*k*.
  - If platform can predict matches (perfectly), advertisers value an ad at *v*; v is distributed U[0,1]
  - Advertiser is indifferent between buying *Nk* ads (w/o prediction) or *N* ads (w/ prediction)
  - If there is ad space of *a* per consumer, platform profits are (1-*a*)*Na* (w/o prediction) and (1-*ka*)*Nka* (w/ prediction)
  - Result: if *a* unconstrained choice, no difference in profits with versus without prediction as *a* = 1/2 (w/o prediction) and *a* = 1/2*k* (w/ prediction). If *a* constrained, it reduces profits w/o prediction. **Prediction is valuable with scarcity**.
- Prat and Valletti (2022)
  - Each of *k* product markets has an incumbent monopolist (earning π_m) and a potential entrant (changes total profits to $2\pi_c$). Entrant needs ad to create awareness.
  - Incumbent may advertise to block entrant from reaching customers. If there are *J* platforms, incumbent will only advertise if $\pi_m - \pi_c > J\pi_c$. Platform wants to create incumbent demand for ads and could not care less about product market competition.

Bundling of Attention and Advertising

- Athey and Gans (2010)
  - Suppose *Nk* consumers value one of *k* products so an advertiser has a 1/*k* probability of the 'right' match; their expected return per ad is *v*/*k*.
  - If platform can predict matches (perfectly), advertisers value an ad at *v*; v is distributed U[0,1]
  - Advertiser is indifferent between buying *Nk* ads (w/o prediction) or *N* ads (w/ prediction)
  - If there is ad space of *a* per consumer, platform profits are (1-*a*)*Na* (w/o prediction) and (1-*ka*)*Nka* (w/ prediction)
  - Result: if *a* unconstrained choice, no difference in profits with versus without prediction as *a* = 1/2 (w/o prediction) and *a* = 1/2*k* (w/ prediction). If *a* constrained, it reduces profits w/o prediction. **Prediction is valuable with scarcity**.
- Prat and Valletti (2022)
  - Each of *k* product markets has an incumbent monopolist (earning π_m) and a potential entrant (changes total profits to $2\pi_c$). Entrant needs ad to create awareness.
  - Incumbent may advertise to block entrant from reaching customers. If there are *J* platforms, incumbent will only advertise if $\pi_m - \pi_c > J\pi_c$. Platform wants to create incumbent demand for ads and could not care less about product market competition.
  - New Result: without prediction, both incumbent and entrant have a 1/k chance of a correct match. With prediction, an incumbent benefits from the ads placed by other incumbents. This free riding means that they will purchase fewer ads with prediction. Thus, platforms have a reduced incentive to adopt better prediction of match quality.

Diverse Business Models


MASTER VOLUME
–
+
INPUT SELECTOR

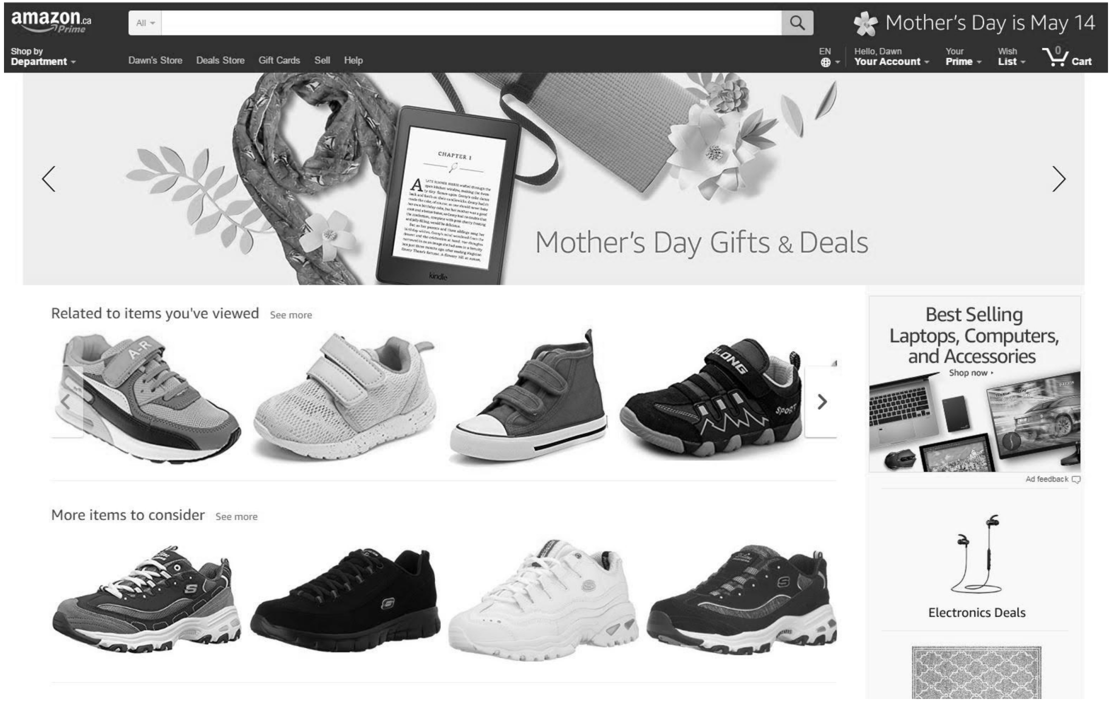
amazon.ca Prime
All
Mother's Day is May 14
Shop by Department
Dawn's Store
Deals Store
Gift Cards
Sell
Help
EN
Hello, Dawn Your Account
Your Prime
Wish List
0 Cart
Mother's Day Gifts & Deals
Related to items you've viewed
See more
More items to consider
See more
Best Selling Laptops, Computers, and Accessories
Shop now
Ad feedback
Electronics Deals

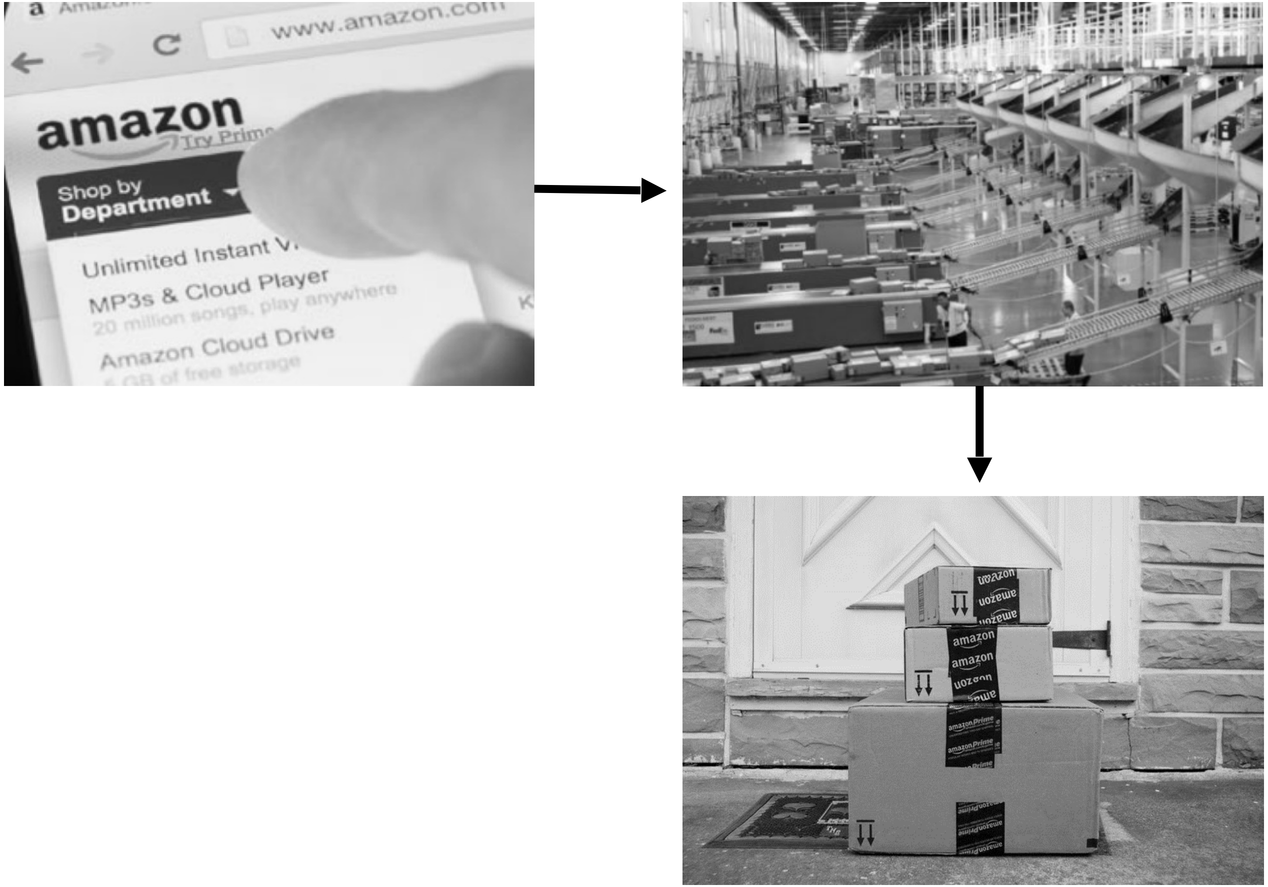
www.amazon.com
amazon
Shop by
Department
MP3s & Cloud Player
20 million songs, play anywhere
Amazon Cloud Drive

US008615473B2

(12) **United States Patent**
Spiegel et al.

(10) **Patent No.: US 8,615,473 B2**
(45) **Date of Patent: Dec. 24, 2013**

(54) **METHOD AND SYSTEM FOR ANTICIPATORY PACKAGE SHIPPING**

(75) Inventors: **Joel R. Spiegel**, Woodinville, WA (US); **Michael T. McKenna**, Bellevue, WA (US); **Girish S. Lakshman**, Issaquah, WA (US); **Paul G. Nordstrom**, Seattle, WA (US)

(73) Assignee: **Amazon Technologies, Inc.**, Reno, NV (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **13/594,195**

(22) Filed: **Aug. 24, 2012**

(65) **Prior Publication Data**

US 2012/0323645 A1 Dec. 20, 2012

Related U.S. Application Data

(62) Division of application No. 13/305,611, filed on Nov. 28, 2011, now Pat. No. 8,271,398, which is a division of application No. 11/015,288, filed on Dec. 17, 2004, now Pat. No. 8,086,546.

(51) **Int. Cl.**
G06Q 99/00 (2006.01)

(52) **U.S. Cl.**
USPC **705/332**; 705/330; 705/333; 705/336; 705/337

(58) **Field of Classification Search**
USPC 705/332, 330, 333, 336, 337
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,055,520	A	4/2000	Heiden et al.
6,394,354	B1	5/2002	Wilz et al.
6,827,273	B2	12/2004	Wilz et al.
6,994,253	B2	2/2006	Miller et al.
7,006,989	B2	2/2006	Bezos et al.
7,130,803	B1	10/2006	Couch et al.
7,222,081	B1	5/2007	Sone
7,610,224	B2	10/2009	Spiegel
7,664,653	B2	2/2010	Dearing
8,086,546	B2	12/2011	Spiegel et al.
2001/0037316	A1	11/2001	Shiloh

(Continued)

FOREIGN PATENT DOCUMENTS

JP	2002109263	4/2002
JP	2003067645	3/2003

(Continued)

OTHER PUBLICATIONS

Office Action from Application No. 2007-546877, mailed Apr. 26, 2011, Amazon Technologies, Inc., 8 pages.

(Continued)

Primary Examiner — Akiba Allen
(74) *Attorney, Agent, or Firm* — Robert C. Kowert; Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.

(57) **ABSTRACT**

A method and system for anticipatory package shipping are disclosed. According to one embodiment, a method may include packaging one or more items as a package for eventual shipment to a delivery address, selecting a destination geographical area to which to ship the package, shipping the package to the destination geographical area without completely specifying the delivery address at time of shipment, and while the package is in transit, completely specifying the delivery address for the package.

24 Claims, 11 Drawing Sheets

US008615473B2

(12) **United States Patent**

(10) **Patent No.: US 8,615,473 B2**

(45) **Date of Patent: Dec. 24, 2013**

(54) METHOD AND SYSTEM FOR ANTICIPATORY PACKAGE SHIPPING

(US); **Girish S. Lakshman**, Issaquah,

(73) Assignee: **Amazon Technologies, Inc.**, Reno, NV (US)

U.S.C. 154(b) by 0 days.

(21) Appl. No.: **13/594,195**

(22) Filed: **Aug. 24, 2012**

(65) **Prior Publication Data**

US 2012/0323645 A1 Dec. 20, 2012

Related U.S. Application Data

(62) Division of application No. 13/305,611, filed on Nov. 28, 2011, now Pat. No. 8,271,398, which is a division of application No. 11/015,288, filed on Dec. 17, 2004, now Pat. No. 8,086,546.

(51) **Int. Cl.**
G06Q 99/00 (2006.01)

(52) **U.S. Cl.**
USPC **705/332**; 705/330; 705/333; 705/336; 705/337

(58) **Field of Classification Search**
USPC 705/332, 330, 333, 336, 337
See application file for complete search history.

References Cited

U.S. PATENT DOCUMENTS

6,055,520	A	4/2000	Heiden et al.
6,394,354	B1	5/2002	Wilz et al.
6,827,273	B2	12/2004	Wilz et al.
6,994,253	B2	2/2006	Miller et al.
7,006,989	B2	2/2006	Bezos et al.
	B1	10/2006	Couch et al.
		5/2007	Sone
		10/2009	Spiegel
		2/2010	Dearing
		12/2011	Spiegel et al.
	A1	11/2001	Shiloh

(Continued)

FOREIGN PATENT DOCUMENTS

JP	2002109263	4/2002
JP	2003067645	3/2003

(Continued)

OTHER PUBLICATIONS

Office Action from Application No. 2007-546877, mailed Apr. 26, 2011, Amazon Technologies, Inc., 8 pages.

(Continued)

Primary Examiner — Akiba Allen

(74) *Attorney, Agent, or Firm* — Robert C. Kowert; Meyertons, Hood, Kivlin, Kowert & Goetzel, P.C.

(57) **ABSTRACT**

A method and system for anticipatory package shipping are disclosed. According to one embodiment, a method may include packaging one or more items as a package for eventual shipment to a delivery address, selecting a destination geographical area to which to ship the package, shipping the package to the destination geographical area without completely specifying the delivery address at time of shipment, and while the package is in transit, completely specifying the delivery address for the package.

24 Claims, 11 Drawing Sheets


MASTER VOLUME
–
+
INPUT SELECTOR

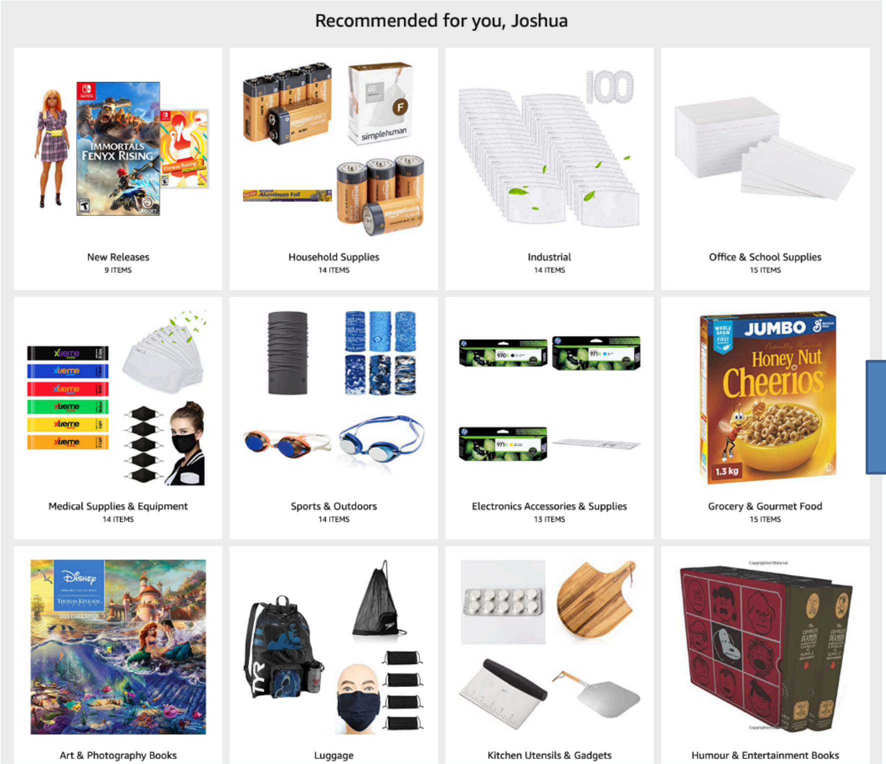

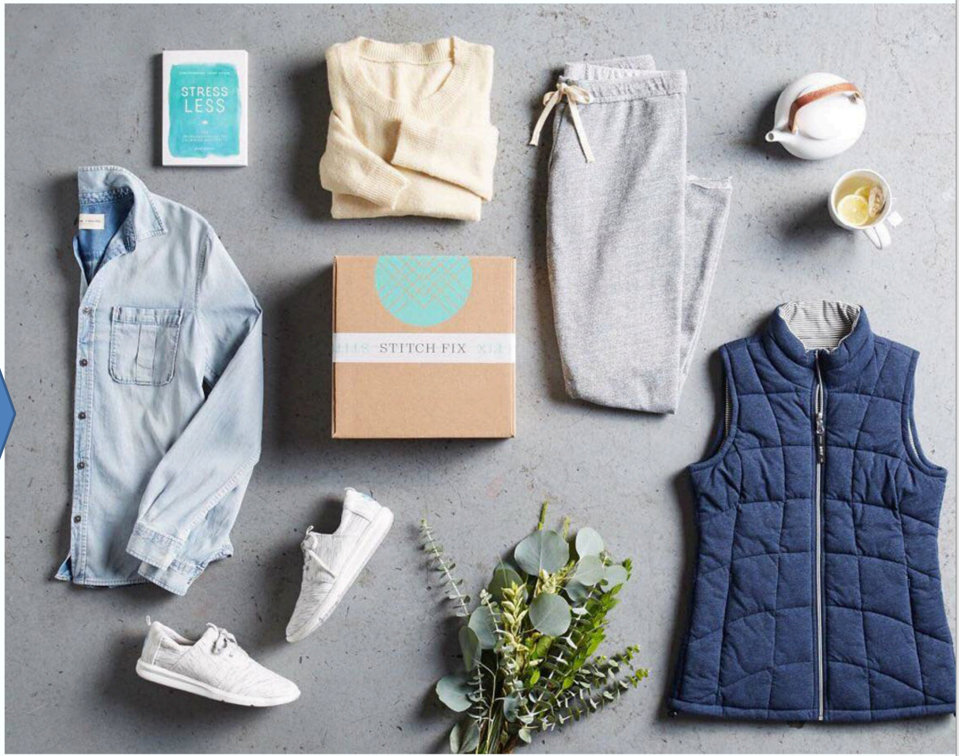

Recommendation Engine

AI Slots into Existing System

Ship without Shopping

AI Drives System

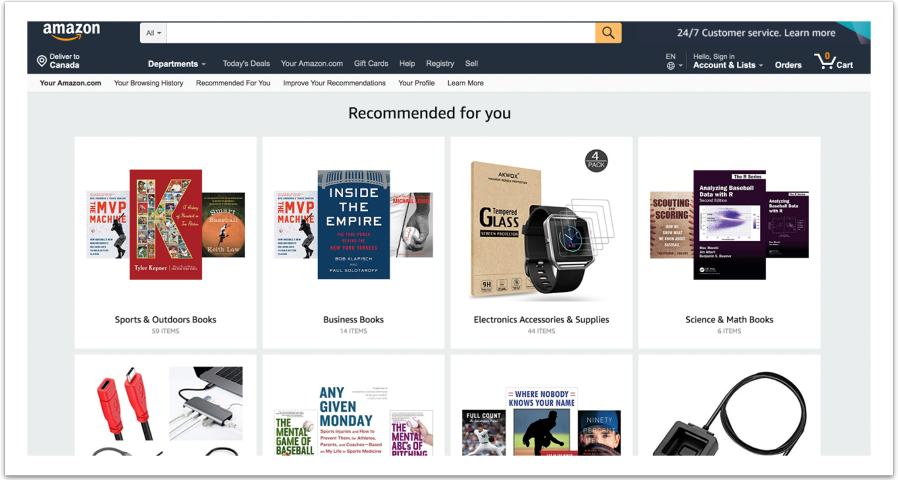
amazon
All
24/7 Customer service. Learn more
Deliver to Canada
Departments
Today's Deals
Your Amazon.com
Gift Cards
Help
Registry
Sell
EN
Hello, Sign in
Account & Lists
Orders
0
Cart
Your Amazon.com
Your Browsing History
Recommended For You
Improve Your Recommendations
Your Profile
Learn More
Recommended for you
THE MVP MACHINE
Tyler Kepner
Smart Baseball
Keith Law
INSIDE THE EMPIRE
BOB KLAPISCH
PAUL SOLOTAROFF
Sports & Outdoors Books
59 ITEMS
Business Books
14 ITEMS
Tempered GLASS
SCREEN PROTECTOR
4 PACK
9H
Electronics Accessories & Supplies
44 ITEMS
SCOUTING SCORING
The R Series
Analyzing Baseball Data with R
Second Edition
Science & Math Books
6 ITEMS
ANY GIVEN MONDAY
THE MENTAL GAME OF BASEBALL
THE MENTAL ABC's OF PITCHING
WHERE NOBODY KNOWS YOUR NAME
FULL COUNT
NINETY PERCENT

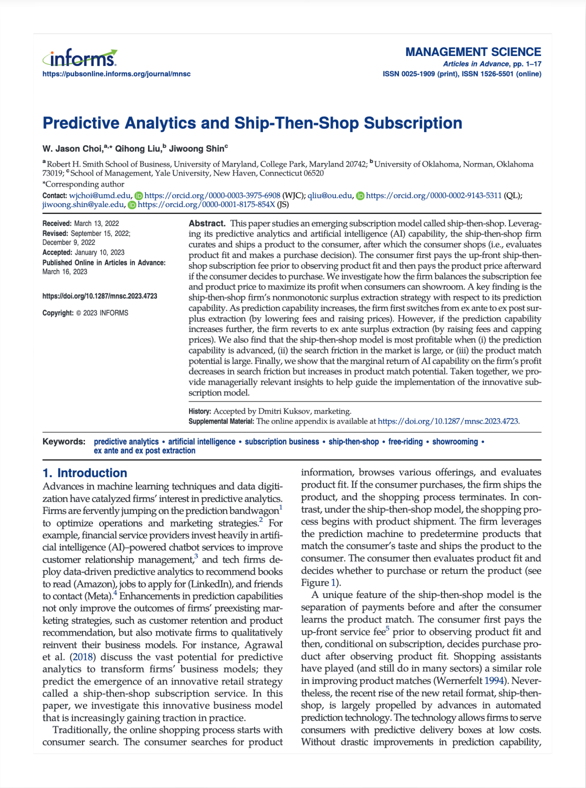

informs
https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–17
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Predictive Analytics and Ship-Then-Shop Subscription

W. Jason Choi,[a,*] Qihong Liu,[b] Jiwoong Shin[c]

[a] Robert H. Smith School of Business, University of Maryland, College Park, Maryland 20742; [b] University of Oklahoma, Norman, Oklahoma 73019; [c] School of Management, Yale University, New Haven, Connecticut 06520

*Corresponding author

Contact: wjchoi@umd.edu, https://orcid.org/0000-0003-3975-6908 (WJC); qliu@ou.edu, https://orcid.org/0000-0002-9143-5311 (QL); jiwoong.shin@yale.edu, https://orcid.org/0000-0001-8175-854X (JS)

Received: March 13, 2022
Revised: September 15, 2022; December 9, 2022
Accepted: January 10, 2023
Published Online in Articles in Advance: March 16, 2023

https://doi.org/10.1287/mnsc.2023.4723

Abstract. This paper studies an emerging subscription model called ship-then-shop. Leveraging its predictive analytics and artificial intelligence (AI) capability, the ship-then-shop firm curates and ships a product to the consumer, after which the consumer shops (i.e., evaluates product fit and makes a purchase decision). The consumer first pays the up-front ship-then-shop subscription fee prior to observing product fit and then pays the product price afterward if the consumer decides to purchase. We investigate how the firm balances the subscription fee and product price to maximize its profit when consumers can showroom. A key finding is the ship-then-shop firm's nonmonotonic surplus extraction strategy with respect to its prediction capability. As prediction capability increases, the firm first switches from ex ante to ex post surplus extraction (by lowering fees and raising prices). However, if the prediction capability increases further, the firm reverts to ex ante surplus extraction (by raising fees and capping prices). We also find that the ship-then-shop model is most profitable when (i) the prediction capability is advanced, (ii) the search friction in the market is large, or (iii) the product match potential is large. Finally, we show that the marginal return of AI capability on the firm's profit decreases in search friction but increases in product match potential. Taken together, we provide managerially relevant insights to help guide the implementation of the innovative subscription model.

History: Accepted by Dmitri Kuksov, marketing.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/mnsc.2023.4723.

Keywords: predictive analytics • artificial intelligence • subscription business • ship-then-shop • free-riding • showrooming • ex ante and ex post extraction

1. Introduction

Advances in machine learning techniques and data digitization have catalyzed firms' interest in predictive analytics. Firms are fervently jumping on the prediction bandwagon[1] to optimize operations and marketing strategies.[2] For example, financial service providers invest heavily in artificial intelligence (AI)–powered chatbot services to improve customer relationship management,[3] and tech firms deploy data-driven predictive analytics to recommend books to read (Amazon), jobs to apply for (LinkedIn), and friends to contact (Meta).[4] Enhancements in prediction capabilities not only improve the outcomes of firms' preexisting marketing strategies, such as customer retention and product recommendation, but also motivate firms to qualitatively reinvent their business models. For instance, Agrawal et al. (2018) discuss the vast potential for predictive analytics to transform firms' business models; they predict the emergence of an innovative retail strategy called a ship-then-shop subscription service. In this paper, we investigate this innovative business model that is increasingly gaining traction in practice.

Traditionally, the online shopping process starts with consumer search. The consumer searches for product information, browses various offerings, and evaluates product fit. If the consumer purchases, the firm ships the product, and the shopping process terminates. In contrast, under the ship-then-shop model, the shopping process begins with product shipment. The firm leverages the prediction machine to predetermine products that match the consumer's taste and ships the product to the consumer. The consumer then evaluates product fit and decides whether to purchase or return the product (see Figure 1).

A unique feature of the ship-then-shop model is the separation of payments before and after the consumer learns the product match. The consumer first pays the up-front service fee[5] prior to observing product fit and then, conditional on subscription, decides purchase product after observing product fit. Shopping assistants have played (and still do in many sectors) a similar role in improving product matches (Wernerfelt 1994). Nevertheless, the recent rise of the new retail format, ship-then-shop, is largely propelled by advances in automated prediction technology. The technology allows firms to serve consumers with predictive delivery boxes at low costs. Without drastic improvements in prediction capability,

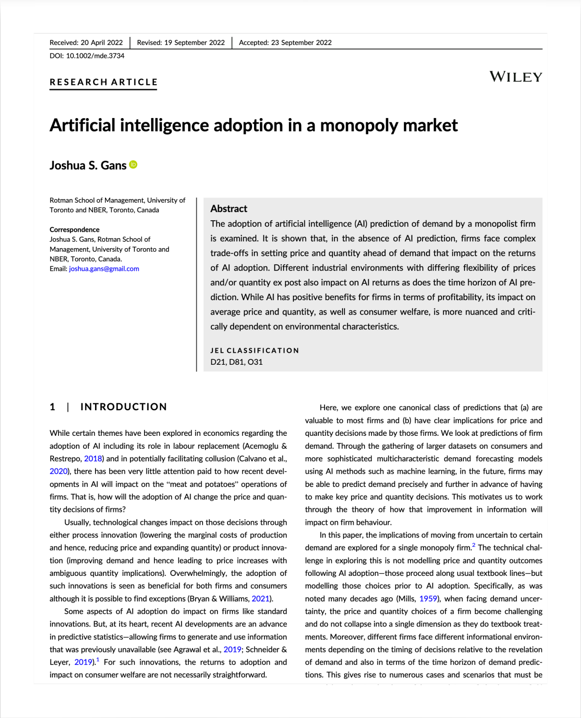

Received: 20 April 2022 | Revised: 19 September 2022 | Accepted: 23 September 2022
DOI: 10.1002/mde.3734

RESEARCH ARTICLE

WILEY

Artificial intelligence adoption in a monopoly market

Joshua S. Gans

Rotman School of Management, University of Toronto and NBER, Toronto, Canada

Correspondence
Joshua S. Gans, Rotman School of Management, University of Toronto and NBER, Toronto, Canada.
Email: joshua.gans@gmail.com

Abstract
The adoption of artificial intelligence (AI) prediction of demand by a monopolist firm is examined. It is shown that, in the absence of AI prediction, firms face complex trade-offs in setting price and quantity ahead of demand that impact on the returns of AI adoption. Different industrial environments with differing flexibility of prices and/or quantity ex post also impact on AI returns as does the time horizon of AI prediction. While AI has positive benefits for firms in terms of profitability, its impact on average price and quantity, as well as consumer welfare, is more nuanced and critically dependent on environmental characteristics.

JEL CLASSIFICATION
D21, D81, O31

1 | INTRODUCTION

While certain themes have been explored in economics regarding the adoption of AI including its role in labour replacement (Acemoglu & Restrepo, 2018) and in potentially facilitating collusion (Calvano et al., 2020), there has been very little attention paid to how recent developments in AI will impact on the "meat and potatoes" operations of firms. That is, how will the adoption of AI change the price and quantity decisions of firms?

Usually, technological changes impact on those decisions through either process innovation (lowering the marginal costs of production and hence, reducing price and expanding quantity) or product innovation (improving demand and hence leading to price increases with ambiguous quantity implications). Overwhelmingly, the adoption of such innovations is seen as beneficial for both firms and consumers although it is possible to find exceptions (Bryan & Williams, 2021).

Some aspects of AI adoption do impact on firms like standard innovations. But, at its heart, recent AI developments are an advance in predictive statistics—allowing firms to generate and use information that was previously unavailable (see Agrawal et al., 2019; Schneider & Leyer, 2019).[1] For such innovations, the returns to adoption and impact on consumer welfare are not necessarily straightforward.

Here, we explore one canonical class of predictions that (a) are valuable to most firms and (b) have clear implications for price and quantity decisions made by those firms. We look at predictions of firm demand. Through the gathering of larger datasets on consumers and more sophisticated multicharacteristic demand forecasting models using AI methods such as machine learning, in the future, firms may be able to predict demand precisely and further in advance of having to make key price and quantity decisions. This motivates us to work through the theory of how that improvement in information will impact on firm behaviour.

In this paper, the implications of moving from uncertain to certain demand are explored for a single monopoly firm.[2] The technical challenge in exploring this is not modelling price and quantity outcomes following AI adoption—those proceed along usual textbook lines—but modelling those choices prior to AI adoption. Specifically, as was noted many decades ago (Mills, 1959), when facing demand uncertainty, the price and quantity choices of a firm become challenging and do not collapse into a single dimension as they do textbook treatments. Moreover, different firms face different informational environments depending on the timing of decisions relative to the revelation of demand and also in terms of the time horizon of demand predictions. This gives rise to numerous cases and scenarios that must be

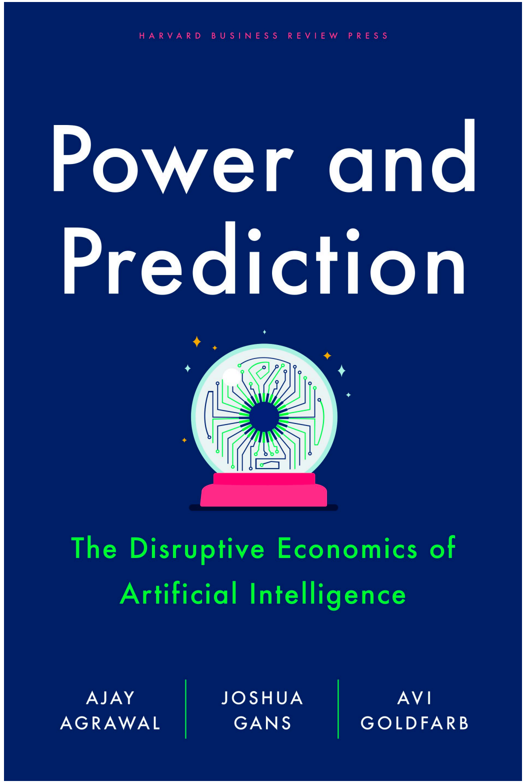

Conclusion

Conclusion

- Work on the competition effects of AI and AI’s effect on competition is still in its infancy

Conclusion

- Work on the competition effects of AI and AI's effect on competition is still in its infancy
- Current results and those inferred from the past literature indicate that the competition arguments are if anything, more nuanced when it comes to AI

Conclusion

- Work on the competition effects of AI and AI’s effect on competition is still in its infancy
- Current results and those inferred from the past literature indicate that the competition arguments are if anything, more nuanced when it comes to AI
- Even without this, beyond the standard antitrust tools, new instruments or policy approaches have not been developed