

The Economics of Artificial Intelligence: An Agenda

Copyright 2019

National Bureau of Economics Research

DOI: 10.7208/chicago/9780226613475.003.0023

Edited by

**Ajay Agrawal, Joshua Gans,
and Avi Goldfarb**

The University of Chicago Press

Chicago and London

How Artificial Intelligence and Machine Learning Can Impact Market Design

Paul R. Milgrom and Steven Tadelis

23.1 Introduction

For millennia, markets have played a key role in providing individuals and businesses with the opportunity to gain from trade. More often than not, markets require structure and a variety of intuitional support to operate efficiently. For example, auctions have become a commonly used mechanism to generate gains from trade when price discovery is essential. Research in the area now commonly referred to as market design, going back to Vickrey (1961), demonstrated that it is critical to design auctions and market institutions more broadly in order to achieve efficient outcomes (see, e.g., Milgrom 2017; Roth 2015).

Any market designer needs to understand some fundamental details of the transactions that are expected to be consummated in order to design the most effective and efficient market structure to support these transactions. For example, the National Resident Matching Program, which matches doctors to hospital residencies, was originally designed in an era when nearly all doctors were men and wives followed them to their residencies. It needed to be redesigned in the 1990s to accommodate the needs of couples, when men and women doctors could no longer be assigned jobs in different cities. Even

Paul R. Milgrom is the Shirley and Leonard Ely Professor of Humanities and Sciences in the Department of Economics at Stanford University and professor, by courtesy, at both the Department of Management Science and Engineering and the Graduate School of Business. Steven Tadelis holds the James J. and Marianne B. Lowrey Chair in Business and is professor of economics, business, and public policy at the University of California, Berkeley, and a research associate of the National Bureau of Economic Research.

For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c14008.ack>.

something as mundane as the sale of a farm when a farmer dies requires knowledge of the structure and decisions about whether to sell the whole farm as a unit, or to separate the house for sale as a weekend retreat while selling the land to neighboring farmers, or selling the forest separately to a wildlife preservation fund.

In complex environments, it can be difficult to understand the underlying characteristics of transactions, and it is challenging to learn enough about them in order to design the best institutions to efficiently generate gains from trade. For example, consider the recent growth of online advertising exchanges that match advertisers with online ads. Many ads are allocated to advertisers using real-time auctions. But how should publishers design these auctions in order to make the best use of their advertising space, and how can they maximize the returns to their activities? Based on the early theoretical auction design work of Myerson (1981), Ostrovsky and Schwartz (2017) have shown that a little bit of market design in the form of setting better reserve prices can have a dramatic impact on the profits an online ad platform can earn.

But how can market designers learn the characteristics necessary to set optimal, or at least better, reserve prices? Or, more generally, how can market designers better learn the environment of their markets? In response to these challenges, artificial intelligence (AI) and machine learning are emerging as important tools for market design. Retailers and marketplaces such as eBay, TaoBao, Amazon, Uber, and many others are mining their vast amounts of data to identify patterns that help them create better experiences for their customers and increase the efficiency of their markets. By having better prediction tools, these and other companies can predict and better manage sophisticated and dynamic market environments. The improved forecasting that AI and machine-learning algorithms provide help marketplaces and retailers better anticipate consumer demand and producer supply as well as help target products and activities to finer segmented markets.

Turning back to markets for online advertising, two-sided markets such as Google, which match advertisers with consumers, are not only using AI to set reserve prices and segment consumers into finer categories for ad targeting, but they also develop AI-based tools to help advertisers bid on ads. In April 2017 Google introduced “Smart Bidding,” a product based on AI and machine learning that helps advertisers bid automatically on ads based on ad conversions so they can better determine their optimal bids. Google explained that the algorithms use vast amounts of data and continually refine models of users’ conversion to better spend an advertiser’s dollars to where they bring in the highest conversion.

Another important application of AI’s strength in improving forecasting to help markets operate more efficiently is in electricity markets. To operate efficiently, electricity market makers such as California’s Independent System Operator must engage in demand and supply forecasting. An inaccurate

forecast in the power grid can dramatically affect market outcomes causing high variance in prices, or worse, blackouts. By better predicting demand and supply, market makers can better allocate power generation to the most efficient power sources and maintain a more stable market.

As the examples above demonstrate, the applications of AI algorithms to market design are already widespread and diverse. Given the infancy of the technology, it is a safe bet that AI will play a growing role in the design and implementation of markets over a wide range of applications. In what follows, we describe several less obvious ways in which AI has played a key role in the operation of markets.

23.2 Machine Learning and the Incentive Auction

In the first part of the twentieth century, the most important infrastructure projects for the United States related to transportation and energy infrastructure. By the early twenty-first century, however, it was not just people and goods that needed to be transported in high volumes, but also information. The emergence of mobile devices, WiFi networks, video on demand, the Internet of Things, services supplied through the cloud, and much more has already created the need for major investments in the communication network, and with 5G technologies just around the corner, more is coming.

Wireless communications, however, depend on infrastructure and other resources. The wireless communication rate depends on the channel capacity, which in turn depends jointly on the communication technology used and the amount of radio spectrum bandwidth devoted to it. To encourage growth in bandwidth and the rapid develop of new uses, the Obama White House in 2010 issued its National Broadband Plan. That plan set a goal of freeing a huge amount of bandwidth from older, less productive uses to be used instead as part of the modern data highway system.

In 2016–2017, the US Federal Communications Commission (FCC) designed and ran an auction market to do part of that job. The radio spectrum licenses that it sold in that auction raised about \$20 billion in gross revenue. As part of the process of making room for those new licenses, the FCC purchased TV broadcast rights for about \$10 billion, and incurred nearly \$3 billion in costs to move other broadcasters to new TV channels. Some 84MHz of spectrum was made available in total, including 70MHz for wireless broadband and 14MHz for unlicensed uses. This section describes the processes that were used, and the role of AI and machine learning to improve the underlying algorithms that supported this market.

Reallocating spectrum from one use to another is, in general, neither easy nor straightforward, in either the planning or the implementation (Leyton-Brown, Milgrom, and Segal 2017). Planning such a change can involve surprisingly hard computational challenges, and the implementation requires high levels of coordination. In particular, the reallocation of a portion

of the spectrum band that had been used for UHF broadcast television required deciding how many channels to clear, which stations would cease broadcasting (to make room for the new uses), what TV channels would be assigned to the remaining stations that continued to broadcast, how to time the changes to avoid interference during the transition, and to assure that the TV tower teams, which would replace the old broadcast equipment, had sufficient capacity, and so on. Several of the computations involved are, in principle, nondeterministic polynomial time (NP)-hard, making this a particularly complex market-design problem. One of the most critical algorithms used for this process—the “feasibility checker”—was developed with the aid of machine-learning methods.

But why reallocate and reassign TV stations at all? Broadcast television changed enormously in the late twentieth century. In the early days of television, all viewing was of over-the-air broadcasts using an analog technology. Over the decades that followed, cable and satellite services expanded so much that, by 2010, more than 90 percent of the US population was reached by these alternative services. Standard definition TV signals were replaced by high definition and, eventually, 4K signals. Digital television and tuners reduced the importance of channel assignments, so that the channel used by consumers/viewers did not need to match the channel used by the broadcaster. Digital encoding made more efficient use of the band and it became possible to use multiplexing, so that what was once a single standard-definition broadcast channel could carry multiple high-definition broadcasts. Marginal spectrum had fallen in value compared to the alternative uses.

Still, the reallocation from television broadcasting would be daunting and beyond what an ordinary market mechanism could likely achieve. The signal from each of thousands of TV broadcast towers across the United States can interfere with potential uses for about 200 miles in every direction, so all of the broadcasts in any frequency needed to be cleared to make the frequencies available for new uses. Not only would it be necessary to coordinate among different areas of the United States, but coordination with Canada and Mexico would improve the allocation, too; most of the population of Canada lives, and most of its TV stations operate, within 200 miles of the US border. Because a frequency is not usable until virtually all of the relevant broadcasters have ceased operation, efficiency would demand that these changes would need to be coordinated in time, too; they should be roughly simultaneous. In addition, there needed to be coordination *across* frequencies. The reason is that we need to know in advance which channels will be cleared before the frequencies can be efficiently divided between uplink uses and downlink uses.

Among the many issues to be resolved, one would be how to determine which stations would continue to broadcast after the transition. If the goal were efficiency, then the problem can be formulated as maximizing the total value of the TV stations that continue to broadcast after the auction. Let N

be the set of all currently broadcasting TV stations and let $S \subseteq N$ be a subset of those TV stations. Let \mathcal{C} be the set of available channels to which to assign stations after the auction, and let \emptyset denote the null assignment for a station that does not continue to broadcast. A channel assignment is a mapping $A : N \rightarrow \mathcal{C} \cup \{\emptyset\}$. The constraints on the channels available for assignment are to ones that rule out interference between pairs of TV stations, taking the form: $A(n_1) = c_1 \Rightarrow A(n_2) \neq c_2$ for some $(c_1, c_2) \in \mathcal{C}^2$. Each such constraint is described by a fourtuple: (n_1, c_1, n_2, c_2) . There were more than a million such constraints in the FCC's problem. A channel assignment is feasible if it satisfies all the interference constraints; let \mathcal{A} denote the feasible set of assignments. A set of stations S' can be feasibly assigned to continue broadcasting, which we denote by $S' \in \mathcal{F}(\mathcal{C})$, if there exists some feasible channel assignment $A \in \mathcal{A}$ such that $\emptyset \notin A(S')$.

Most of the interference constraints took a special form. Those constraints assert that no two stations that are geographic neighbors can be assigned to the same channel. Let us call such stations "linked" and denote the relationship by $(n_1, n_2) \in L$. For such a pair of stations, the constraint can be written as: $A(n_1) = A(n_2) \Rightarrow A(n_1) = \emptyset$. These are the *cochannel interference constraints*. One can think of (N, L) as defining a graph with nodes N and arcs L . If the cochannel constraints were the only ones, then determining whether $S' \in \mathcal{F}$ would amount to deciding whether there exists a way to assign channels in \mathcal{C} to the stations in N so that no two linked nodes are on the same channel.

Figure 23.1 shows the graph of the cochannel interference constraints for the United States and Canada. The constraint graph is most dense in the eastern half of the United States and along the Pacific Coast.

In the special case of cochannel constraints, the problem of checking the feasibility of a set of stations is a standard *graph-coloring* problem. The problem is to decide whether it is possible to assign a color (channel) to each node (station) in the graph so that no two linked nodes are given the same color. Graph-coloring is in the class of NP-complete problems, for which there is no known algorithm that is guaranteed to be fast, and for which it is commonly hypothesized¹ that worst-case solution time grows exponentially in the problem size. Since the general station assignment problem includes the graph coloring problem, it, too, is NP-complete, and can become intractable at scales such as that of the FCC's problem.

The problem that the FCC would ideally like to solve using an auction is to maximize the value of the stations that remain on-air to broadcast, given the reduced set of channels \mathcal{C} . If the value of station j is v_j , the problem can be formulated as follows:

$$\max_{S \in \mathcal{F}(\mathcal{C})} \sum_{j \in S} v_j.$$

1. The standard computer science hypothesis that $P \neq NP$ implies that no fast algorithm exists for NP-complete problems.




Fig. 23.1 Cochannel interference graph for spectrum reallocation

This problem is very hard. Indeed, as we have just argued, even checking the condition $S \in \mathcal{F}(C)$ is NP-complete, and solving exactly the related optimization is even harder in practice. Computational experiments suggest that with weeks of computation approximate optimization is possible, but with an optimization shortfall that can be a few percent.

For a TV station owner, it would be daunting to formulate a bid in an auction in which even the auctioneer, with all the bids in hand, would find it challenging to determine the winners. Faced with such a problem, some station owners might choose not to participate. That concern led the FCC staff to prefer a strategy-proof design, in which the optimal bid for the owner of a single station is relatively simple, at least in concept: compute your station's value and bid that amount. As is well known, there is a unique strategy-proof auction that optimizes the allocation and pays zero to the losers: the Vickrey auction. According to the Vickrey rules, if the auctioneer purchases the broadcast rights from station j , it must pay the owner this price:

$$p_i = \left(\max_{S \in \mathcal{F}(C)} \sum_{j \in S} v_j \right) - \left(\max_{\substack{S \in \mathcal{F}(C) \\ i \notin S}} \sum_{j \in S} v_j \right).$$

For a winning station i , the Vickrey price p_i will be larger than the station value. With roughly 2,000 stations to include in the optimization, a 1 percent error in either of the two maximizations would result in a pricing error for p_i equal to about 2,000 percent of the value of an average station. Such huge potential pricing errors would likely raise hackles among some of the potential bidders.

One way to put the problem of the Vickrey auction into sharp relief is to imagine the letter that the FCC might write to broadcasters to encourage their participation:

Dear Mr. Broadcaster:

We have heard your concerns about the complexity of the spectrum reallocation process. You may even be unsure about whether to participate or how much to bid. To make things as easy as possible for you, we have adopted a Nobel Prize-winning auction procedure called the "Vickrey auction." In this auction, all you need to do is to tell us what your broadcast rights are worth to you. We'll figure out whether you are a winner and, if so, how much to pay to buy your rights. The rules will ensure that it is in your interest to report truthfully. That is the magic of the Vickrey auction!

The computations that we do will be very hard ones, and we cannot guarantee that they will be exactly correct.

Such a letter would leave many stations owners uncomfortable and unsure about whether to participate. The FCC decided to adopt a different design.

What we describe here is a simplified version of the design, in which the broadcasters' only choices are whether to sell their rights or to reject the

FCC's offer and continue to broadcast. Each individual broadcaster was comforted by the assurance that it could bid this way, even if it had additional options, too.²

In the simplified auction, each bidder i was quoted a price $p_i(t)$ at each round t of the auction that decreased from round-to-round. In each round, the bidder could “exit,” rejecting the current price and keeping its broadcast rights, or it could accept the current price. After a round t of bidding, stations were processed one at a time. When station i was processed, the auction software would use its *feasibility checker* to attempt to determine whether it could feasibly assign station i to continue broadcasting, given the other stations that had already exited and to which a channel must be assigned. This is the generalized graph-coloring problem, mentioned earlier. If the software timed out, or if it determined that it is impossible to assign the station, then the station would become a winner and be paid $p_i(t-1)$. Otherwise, its price would be reduced to $p_i(t)$ and it would exit or continue, according to the bidder's instructions. It would be obvious to a station owner that, regardless of the pricing formula and of how the software performed, its optimal choice when its value is v_i is to exit if $p_i(t) < v_i$ and otherwise to continue.³

The theory of clock auctions of this sort for problems with hard computations has been developed by Milgrom and Segal (2017), who also report simulations showing high performance in terms of efficiency and remarkably low costs of procuring TV broadcast rights.

The performance of this descending auction design depends deeply on the quality of the feasibility checker. Based on early simulations, our rough estimate was the each 1 percent of failures in feasibility checking would add about 1.5 percent—or about \$150 million—to the cost of procuring the broadcast rights. So, solving most of the problems very fast became a high priority for the auction-design team.

As a theoretical proposition, any known algorithm for feasibility checking in the spectrum-packing problem has worst-case performance that grows exponentially in the size of the problem. Nevertheless, if we know the distribution of likely problems, there can still be algorithms that are fast with

2. In the actual auction, some broadcasters also had the option to switch from a UHF TV channel to a channel in the high VHF band, or one in the low VHF band (the so-called HVHF and LVHF options).

3. The pricing formula that the FCC used for each station was $p_i(t) = (\text{Pop}_i \text{Links}_i)^{0.5} q(t)$. In this formula, $q(t)$ is the “base clock price” that scaled the price offers to all the bidders. This price began at a high level $q(0)$ to encourage participation, and it declined round-by-round during the auction; Pop_i denotes the population of the area served by the station, which stands in for the value of the station. By linking prices to population served, the auctioneer is able to offer higher prices to valuable stations in high-population areas that it might need to acquire for a successful auction; Links_i , measured the number of other stations to which station i was linked in the interference graph. It was hoped that, by including this term in the pricing formula, the auction would be able to offer higher prices to and buy the rights of stations that pose particularly difficult problems by interfering with many other stations.

high probability. But how can we know the distribution and how can such an algorithm be found?

The FCC auction used a feasibility checker developed by a team of Auctionomics researchers at the University of British Columbia, led by Professor Kevin Leyton-Brown. There were many steps in the development, as reported by Newman, Fréchet, and Leyton-Brown (forthcoming), but here we emphasize the role of machine learning. Auctionomics' goal was to be able to solve 99 percent of the problem instances in one minute or less.

The development effort began by simulating the planned auction to generate feasibility problems like those that might be encountered in a real auction. Running many simulations generated about 1.4 million problem instances that could be used for training and testing a feasibility-checking algorithm. The first step of the analysis was to formulate the problem as mixed integer programs and test standard commercial software—CPLEX and Gurobi—to see how close those could come to meeting the performance objectives. The answer was: not close. Using a 100-second cutoff, Gurobi could solve only about 10 percent of the problems and CPLEX only about 25 percent. These were not nearly good enough for decent performance in a real-time auction.

Next, the same problems were formulated as satisfiability problems and tested using seventeen research solvers that had participated in recent SAT-solving tournaments. These were better, but none could solve as many as two-thirds of the problems within the same 100-second cutoff. The goal remained 99 percent in sixty seconds.

The next step was to use automated algorithm configuration, a procedure developed by Hutter, Hoos, and Leyton-Brown (2011) and applied in this setting by Leyton-Brown and his students at the University of British Columbia. The idea is to start with a highly parameterized algorithm for solving satisfiability problems⁴ and to train a random forest model of the algorithm performance, given the parameters. To do that, we first ran simulated auctions with what we regarded as plausible behavior by the bidders to generate a large data set of representative problems. Then, we solved those problems using a variety of different parameter settings to determine the distribution of solution times for each vector of parameters. This generated a data set with parameters and performance measures. Two of the most interesting performance characteristics were the median run time and

4. There are no known algorithms for NP-complete problems that are guaranteed to be fast, so the best existing algorithms are all heuristics. These algorithms weight various characteristics of the problem to decide about such things as the order in which to check different branches of a search tree. These weights are among the parameters that can be set and adapted to work well for a particular class of problems, such as those that arise in the incentive auction application. The particular software algorithm that we used was CLASP, which had more than 100 exposed parameters that could be modified.

the fraction of instances solved within one minute. Then, using a Bayesian model, we incorporated uncertainty in which the experimenter “believes” that the actual performance is normally distributed with a mean determined by the random forest and a variance that depends on the distance of the parameter vector from the nearest points in the data set. Next, the system identifies the parameter vector that maximizes the expected improvement in performance, given the mean and variance of the prior and the performance of the best-known parameter vector. Finally, the system tests the actual performance for the identified parameters and adds that as an observation to the data set. Proceeding iteratively, the system identifies more parameters to test, investigates them, and adds them to the data to improve the model accuracy until the time budget is exhausted.

Eventually, this machine-learning method leads to diminishing returns to time invested. One can then create a new data set from the instances on which the parameterized algorithm was “slow,” for example, taking more than fifteen seconds to solve. By training a new algorithm on those instances, and running the two parameterized algorithms in parallel, the machine-learning techniques led to dramatic improvements in performance.

For the actual auction, several other problem-specific tricks were also applied to contribute to the speed-up. For example, to some extent it proved possible to decompose the full problem into smaller problems, to reuse old solutions as starting points for a search, to store partial solutions that might help guide solutions of further problems, and so on. In the end, the full set of techniques and tricks resulted in a very fast feasibility checker that solved all but a tiny fraction of the relevant problems within the allotted time.

23.3 Using AI to Promote Trust in Online Marketplaces

Online marketplaces such as eBay, Taobao, Airbnb, and many others have grown dramatically since their inception just over two decades ago, providing businesses and individuals with previously unavailable opportunities to purchase or profit from online trading. Wholesalers and retailers can market their goods or get rid of excess inventory; consumers can easily search marketplaces for whatever is on their mind, alleviating the need for businesses to invest in their own e-commerce website; individuals transform items they no longer use into cash; and more recently, the so called “gig economy” is comprised of marketplaces that allow individuals to share their time or assets across different productive activities and earn extra income.

The amazing success of online marketplaces was not fully anticipated, primarily because of the hazards of anonymous trade and asymmetric information. Namely, how can strangers who have never transacted with one another, and who may be thousands of miles apart, be willing to trust each other? Trust on both sides of the market is essential for parties to be willing to transact and for a marketplace to succeed. The early success of eBay is

often attributed to the innovation of introducing its famous feedback and reputation mechanism, which was adopted in one form or another by practically every other marketplace that came after eBay. These online feedback and reputation mechanisms provide a modern-day version of more ancient reputation mechanisms used in the physical marketplaces that were the medieval trade fairs of Europe (see Milgrom, North, and Weingast 1990).

Still, recent studies have shown that online reputation measures of marketplace sellers, which are based on buyer-generated feedback, don't accurately reflect their actual performance. Indeed, a growing literature has shown that user-generated feedback mechanisms are often biased, suffer from "grade inflation," and can be prone to manipulation by sellers.⁵ For example, the average percent positive for sellers on eBay is about 99.4 percent, with a median of 100 percent. This causes a challenge to interpret the true levels of satisfaction on online marketplaces.

A natural question emerges: Can online marketplaces use the treasure trove of data it collects to measure the quality of a transaction and predict which sellers will provide a better service to their buyers? It has become widely known that all online marketplaces, as well as other web-based services, collect vast amounts of data as part of the process of trade. Some refer to this as the "exhausts data" generated by the millions of transactions, searches, and browsing that occur on these marketplaces daily. By leveraging this data, marketplaces can create an environment that would promote trust, not unlike the ways in which institutions emerged in the medieval trade fairs of Europe that helped foster trust. The scope for market design goes far beyond the more mainstream application like setting rules of bidding and reserve prices for auctions or designing tiers of services, and in our view, includes the design of mechanisms that help foster trust in marketplaces. What follows are two examples from recent research that show some of the many ways that marketplaces can apply AI to the data they generate to help create more trust and better experiences for their customers.

23.3.1 Using AI to Assess the Quality of Sellers

One of the ways that online marketplaces help participants build trust is by letting them communicate through online messaging platforms. For example, on eBay buyers can contact sellers to ask them questions about their products, which may be particularly useful for used or unique products for which buyers may want to get more refined information than is listed. Similarly, Airbnb allows potential renters to send messages to hosts and ask questions about the property that may not be answered in the original listing.

Using Natural Language Processing (NLP), a mature area in AI, market-

5. On bias and grade inflation see, for example, Nosko and Tadelis (2015), Zervas, Proserpio, and Byers (2015), and Filippas, Horton, and Golden (2017). On seller manipulation of feedback scores see, for example, Mayzlin, Dover, and Chevalier (2014) and Xu et al. (2015).

places can mine the data generated by these messages in order to better predict the kind of features that customers value. However, there may also be subtler ways to apply AI to manage the quality of marketplaces. The messaging platforms are not restricted to pretransaction inquiries, but also offer the parties to send messages to each other *after* the transaction has been completed. An obvious question then emerges: How could a marketplace analyze the messages sent between buyers and sellers post the transaction to infer something about the quality of the transaction that feedback doesn't seem to capture?

This question was posed and answered in a recent paper by Masterov, Mayer, and Tadelis (2015) using internal data from eBay's marketplace. The analysis they performed was divided into two stages. In the first stage, the goal was to see if NLP can identify transactions that went bad when there was an independent indication that the buyer was unhappy. To do this, they collected internal data from transactions in which messages were sent from the buyer to the seller after the transaction was completed, and matched it with another internal data source that recorded actions by buyers indicating that the buyer had a poor experience with the transactions. Actions that indicate an unhappy buyer include a buyer claiming that the item was not received, or that the item was significantly not as described, or leaves negative or neutral feedback, to name a few.

The simple NLP approach they use creates a "poor-experience" indicator as the target (dependent variable) that the machine-learning model will try to predict, and uses the messages' content as the independent variables. In its simplest form and as a proof of concept, a regular expression search was used that included a standard list of negative words such as "annoyed," "dissatisfied," "damaged," or "negative feedback" to identify a message as negative. If none of the designated terms appeared, then the message was considered neutral. Using this classification, they grouped transactions into three distinct types: (a) no posttransaction messages from buyer to seller, (b) one or more negative messages, or (c) one or more neutral messages with no negative messages.

Figure 23.2, which appears in Masterov, Mayer, and Tadelis (2015), describes the distribution of transactions with the different message classifications together with their association with poor experiences. The x-axis of figure 23.1 shows that approximately 85 percent of transactions fall into the benign first category of no posttransaction messages. Buyers sent at least one message in the remaining 15 percent of all transactions, evenly split between negative and neutral messages. The top of the y-axis shows the poor experience rate for each message type. When no messages are exchanged, only 4 percent of buyers report a poor experience. Whenever a neutral message is sent, the rate of poor experiences jumps to 13 percent, and if the message's content was negative, over one-third of buyers express a poor experience.

In the second stage of the analysis, Masterov, Mayer, and Tadelis (2015)




Fig. 23.2 Message content and poor experiences on eBay

Source: Masterov et al. 2015. ©2015 Association for Computing Machinery, Inc. Reprinted by permission. <https://doi.org/10.1145/2764468.2764499>.

used the fact that negative messages are associated with poor experiences to construct a novel measure of seller quality based on the idea that sellers who receive a higher frequency of negative messages are worse sellers. For example, imagine that seller A and seller B both sold 100 items and that seller A had five transactions with at least one negative message, while seller B had eight such transactions. The implied quality score of seller A is then 0.05 while that of seller B is 0.08, and the premise is that seller B is a worse seller than seller A. Masterov, Mayer, and Tadelis (2015) show that the relationship between this ratio, which is calculated for every seller at any point in time using aggregated negative messages from past sales, and the likelihood that a current transaction will result in a poor experience, is monotonically increasing.

This simple exercise is a proof of concept that shows that by using the message data and a simple natural language processing AI procedure, they were able to better predict which sellers will create poor experiences than one can infer from the very inflated feedback data. eBay is not unique in allowing the parties to exchange messages and the lessons from this research are easily generalizable to other marketplaces. The key is that there is information in

communication between market participants, and past communication can help identify and predict the sellers or products that will cause buyers poor experiences and negatively impact the overall trust in the marketplace.

23.2.2 Using AI to Create a Market for Feedback

Aside from the fact that feedback is often inflated as described earlier, another problem with feedback is that not all buyers choose not to leave feedback at all. In fact, through the lens of mainstream economic theory, it is surprising that a significant fraction of online consumers leave feedback. After all, it is a selfless act that requires time, and it creates a classic free-rider problem. Furthermore, because potential buyers are attracted to buy from sellers or products that already have an established good track record, this creates a “cold-start” problem: new sellers (or products) with no feedback will face a barrier-to-entry in that buyers will be hesitant to give them a fair shot. How could we solve these free-rider and cold-start problems?

These questions were analyzed in a recent paper by Li, Tadelis, and Zhou (2016) using a unique and novel implementation of a market for feedback on the huge Chinese marketplace Taobao where they let sellers pay buyers to leave them feedback. Naturally, one may be concerned about allowing sellers to pay for feedback as it seems like a practice in which they will only pay for good feedback and suppress any bad feedback, which would not add any value in promoting trust. However, Taobao implemented a clever use of NLP to solve this problem: it is the platform, using an NLP AI model, that decides whether feedback is relevant and not the seller who pays for the feedback. Hence, the reward to the buyer for leaving feedback was actually managed by the marketplace, and was handed out for informative feedback rather than for positive feedback.

Specifically, in March 2012, Taobao launched a “Rebate-for-Feedback” (RFF) feature through which sellers can set a rebate value for any item they sell (cash back or store coupon) as a reward for a buyer’s feedback. If a seller chooses this option, then Taobao guarantees that the rebate will be transferred from the seller’s account to a buyer who leaves high-quality feedback. Importantly, feedback quality only depends on how informative it is, rather than whether the feedback is positive or negative. Taobao measures the quality of feedback with a NLP algorithm that examines the comment’s content and length and finds out whether key features of the item are mentioned. Hence, the marketplace manages the market for feedback by forcing the seller to deposit at Taobao a certain amount for a chosen period, so that funds are guaranteed for buyers who meet the rebate criterion, which itself is determined by Taobao.⁶

6. According to a Taobao survey (published in March 2012), 64.8 percent of buyers believed that they will be more willing to buy items that have the RFF feature, and 84.2 percent of buyers believed that the RFF option will make them more likely to write detailed comments.

Taobao's motivation behind the RFF mechanism was to promote more informative feedback, but as Li, Tadelis, and Zhou (2016) noted, economic theory offers some insights into how the RFF feature can act as a potent signaling mechanism that will further separate higher- from lower-quality sellers and products. To see this, recall the literature launched by Nelson (1970) who suggested that advertising acts as a signal of quality. According to the theory, advertising—which is a form of burning money—acts as a signal that attracts buyers who correctly believe that only high-quality sellers will choose to advertise. Incentive compatibility is achieved through repeat purchases: buyers who purchase and experience the products of advertisers will return in the future only if the goods sold are of high enough quality. The cost of advertising can be high enough to deter low-quality sellers from being willing to spend the money and sell only once because those sellers will not attract repeat customers, and still low enough to leave profits for higher-quality sellers. Hence, ads act as signals that separate high-quality sellers, and in turn attract buyers to their products.

As Li, Tadelis, and Zhou (2016) argue, the RFF mechanism plays a similar signaling role as ads do. Assuming that consumers express their experiences truthfully in written feedback, any consumer who buys a product and is given incentives to leave feedback will leave positive feedback only if the buying experience was satisfactory. Hence, a seller will offer RFF incentives to buyers only if the seller expects to receive positive feedback, and this will happen only if the seller will provide high quality. If a seller knows that their goods and services are unsatisfactory, then paying for feedback will generate negative feedback that will harm the low-quality seller. Equilibrium behavior then implies that RFF, as a signal of high quality, will attract more buyers and result in more sales. The role of AI was precisely to reward buyers for information, not for positive feedback.

Li, Tadelis, and Zhou (2016) proceeded to analyze data from the period where the RFF mechanism was featured and confirmed that first, as expected, more feedback was left in response to the incentives provided by the RFF feature. More important, the additional feedback did not exhibit any biases, suggesting that the NLP algorithms used were able to create the kind of screening needed to select informative feedback. Also, the predictions of the simple signaling story were borne out in the data, suggesting that using NLP to support a novel market for feedback did indeed solve both the free-rider problem and the cold-start problem that can hamper the growth of online marketplaces.

23.4 Using AI to Reduce Search Frictions

An important application of AI and machine learning in online marketplaces is the way in which potential buyers engage with the site and proceed to search for products or services. Search engines that power the search of

products online are based on a variety of AI algorithms that are trained to maximize what the provider believes to be the right objective. Often this boils down to conversion, under the belief that the sooner a consumer converts a search to a purchase, the happier the consumer is both in the short and the long run. The rationale is simply that search itself is a friction, and hence, maximizing the successful conversion of search activity to a purchase reduces this friction.

This is not inconsistent with economic theory that has modeled search as an inevitable costly process that separates consumers from the products they want. The canonical search models in economics either build on the seminal work of Stigler (1961), who assumes that consumers sample a fixed number of stores and choose to buy the lowest priced item, or more often, on the models of McCall (1970) and Mortensen (1970), who posit that a model of sequential search is a better description of consumer search behavior. In both modeling approaches consumers know exactly what they wish to buy.

However, it turns out that unlike the simplistic models of search employed in economic theory, where consumers know what they are looking for and the activity of search is just a costly friction, in reality, people's search behavior is rich and varied. A recent paper by Blake, Nosko, and Tadelis (2016) uses comprehensive data from eBay to shed light on the search process with minimal modeling assumptions. Their data show that consumers search significantly more than other studies—which had limited access to search behavior over time—have suggested.

Furthermore, search often proceeds from the vague to the specific. For example, early in a search a user may use the query “watch,” then refine it to “men's watch,” and later add further qualifying words such as color, shape, strap type, and more. This suggests that consumers often learn about their own tastes, and what product characteristics exist, as part of the search process. Indeed, Blake et al. (2016) show that the average number of terms in the query rises over time, and the propensity to use the default-ranking algorithm declines over time as users move to more focused searches like price sorting.

These observations suggest that marketplaces and retailers alike could design their online search algorithms to understand search intent so as to better serve their consumers. If a consumer is in the earlier, exploratory phases of the search process, then offering some breadth will help the consumer better learn their tastes and the options available in the market. But when the consumer is driven to purchase something particular, offering a narrower set of products that match the consumer's preferences would be better. Hence, machine learning and AI can play an instrumental role in recognizing customer intent.

Artificial intelligence and machine learning cannot only help predict a customer's intent, but given the large heterogeneity on consumer tastes, AI

can help a marketplace or retailer better segment the many customers into groups that can be better served with tailored information. Of course, the idea of using AI for more refined customer segmentation, or even personalized experiences, also raises concerns about price discrimination. For example, in 2012 the *Wall Street Journal* reported that “Orbitz Worldwide Inc. has found that people who use . . . Mac computers spend as much as 30% more a night on hotels, so the online travel agency is starting to show them different, and sometimes costlier, travel options than Windows visitors see. The Orbitz effort, which is in its early stages, demonstrates how tracking people’s online activities can use even seemingly innocuous information—in this case, the fact that customers are visiting Orbitz.com from a Mac—to start predicting their tastes and spending habits.”⁷

Whether these practices of employing consumer data and AI will help or harm consumers is not obvious, as it is well known from economic theory that price discrimination can either increase or reduce consumer welfare. If, on average, Mac users prefer staying at fancier and more expensive hotels because owning a Mac is correlated with higher income and tastes for luxury, then the Orbitz practice is beneficial because it shows people what they want to see and reduces search frictions. However, if this is just a way to extract more surplus from consumers who are less price sensitive, but do not necessarily care for the snazzier hotel rooms, then it harms these consumers.

There is currently a lot of interest in policy circles regarding the potential harms to consumers from AI-based price discrimination and market segmentation. McSweeney and O’Dea (2017) suggest that once AI is used to create more targeted market segments, this may not only have implications only for consumer welfare, but for antitrust policy and market definitions for mergers. But, as Gal and Elkin-Koren (2017) suggest, the same AI-targeting tools used by retailers and marketplaces to better segment consumers may be developed into tools for consumers that will help them shop for better deals and limit the ways in which marketplaces and retailers can engage in price discrimination.

23.5 Concluding Remarks

In its early years, classical economic theory paid little attention to market frictions and treated information and computation as free. That theory led to conclusions about efficiency, competitive prices for most goods, and full employment of valuable resources. To address the failures of that theory, economists began to study models with search frictions, which predicted that price competition would be attenuated, that some workers and resources

7. See “On Orbitz, Mac Users Steered to Pricier Hotels,” Dana Mattioli, *The Wall Street Journal*, Aug. 23, 2012. <https://www.wsj.com/articles/SB10001424052702304458604577488822667325882>.

could remain unemployed, and that it could be costly to distinguish reliable trading partners from others. They also built markets for complex resource-allocation problems in which computations and some communications were centralized, lifting the burden of coordination from individual market participants.

With these as the key frictions in the traditional economy, AI holds enormous potential to improve efficiency. In this chapter, we have described some of the ways that AI can overcome computational barriers, reduce search frictions, and distinguish reliable partners. These are among the most important causes of inefficiency in traditional economies, and there is no longer any question that AI is helping to overcome them, with the promise of widespread benefits for all of us. As Roth (2002) noted, market designers “cannot work only with the simple conceptual models used for theoretical insights into the general working of markets. Instead, market design calls for an engineering approach.” Artificial intelligence has already proven to be a valuable tool in the economist-as-engineer tool box.

References

- Blake, T., C. Nosko, and S. Tadelis. 2016. “Returns to Consumer Search: Evidence from eBay.” *17th ACM Conference on Electronic Commerce (EC 2016)*, 531–45.
- Filippas, A., J. J. Horton, and J. M. Golden. 2017. “Reputation in the Long-Run.” Working paper, Stern School of Business, New York University.
- Gal, M. S., and N. Elkin-Koren. 2017. “Algorithmic Consumers.” *Harvard Journal of Law & Technology* 30:1–45.
- Hutter, F., H. Hoos, and K. Leyton-Brown. 2011. “Sequential Model-Based Optimization for General Algorithm Configuration.” Conference paper, International Conference on Learning and Intelligent Optimization. https://link.springer.com/chapter/10.1007/978-3-642-25566-3_40.
- Leyton-Brown, K., P. R. Milgrom, and I. Segal. 2017. “Economics and Computer Science of a Radio Spectrum Reallocation.” *Proceedings of the National Academy of Sciences* 114 (28): 7202–09. www.pnas.org/cgi/doi/10.1073/pnas.1701997114.
- Li, L. I., S. Tadelis, and X. Zhou. 2016. “Buying Reputation as a Signal of Quality: Evidence from an Online Marketplace.” NBER Working Paper no. 22584, Cambridge, MA.
- Masterov, D. V., U. F. Mayer, and S. Tadelis. 2015. “Canary in the E-commerce Coal Mine: Detecting and Predicting Poor Experiences Using Buyer-to-Seller Messages.” In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC ‘15, 81–93.
- Mayzlin, D., Y. Dover, and J. Chevalier. 2014. “Promotional Reviews: An Empirical Investigation of Online Review Manipulation.” *American Economic Review* 104 (8): 2421–55.
- McCall, J. J. 1970. “Economics of Information and Job Search.” *Quarterly Journal of Economics* 84 (1): 113–26.
- McSweeney, T., and B. O’Dea. 2017. “The Implications of Algorithmic Pricing for

- Coordinated Effects, Analysis and Price Discrimination Markets in Antitrust Enforcement.” *Antitrust* 32 (1): 75–81.
- Milgrom, P. R. 2017. *Discovering Prices: Auction Design in Markets with Complex Constraints*. New York: Columbia University Press.
- Milgrom, P. R., D. C. North, and B. R. Weingast. 1990. “The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs.” *Economics and Politics* 2 (1): 1–23.
- Milgrom, P. R., and I. Segal. 2017. “Deferred Acceptance Auctions and Radio Spectrum Reallocation.” Working paper.
- Mortensen, D. T. 1970. “Job Search, the Duration of Unemployment and the Phillips Curve.” *American Economic Review* 60 (5): 847–62.
- Myerson, R. B. 1981. “Optimal Auction Design.” *Mathematics of Operations Research* 6 (1): 58–73.
- Nelson, P. 1970. “Information and Consumer Behavior.” *Journal of Political Economy* 78 (2): 311–29.
- Newman, N., A. Fréchet, and K. Leyton-Brown. Forthcoming. “Deep Optimization for Spectrum Repacking.” *Communications of the ACM (CACM)*.
- Nosko, C., and S. Tadelis. 2015. “The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment.” NBER Working Paper no. 20830, Cambridge, MA.
- Ostrovsky, M., and M. Schwartz. 2017. “Reserve Prices in Internet Advertising: A Field Experiment.” Working paper.
- Stigler, G. J. 1961. “The Economics of Information.” *Journal of Political Economy* 69 (3): 213–25.
- Roth, A. E. 2002. “The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics.” *Econometrica* 70 (4): 1341–78.
- Roth, A. E. 2015. *Who Gets What—and Why: The New Economics of Matchmaking and Market Design*. New York: Houghton Mifflin Harcourt.
- Vickrey, W. 1961. “Counterspeculation, Auctions, and Competitive Sealed Tenders.” *Journal of Finance, American Finance Association* 16 (1): 8–37.
- Xu, H., D. Liu, H. Wang, and A. Stavrou. 2015. “E-commerce Reputation Manipulation: The Emergence of Reputation-Escalation-as-a-Service.” *Proceedings of 24th World Wide Web Conference (WWW 2015)*:1296–306.
- Zervas, G., D. Proserpio, and J. W. Byers. 2015. “A First Look at Online Reputation on Airbnb, Where Every Stay Is Above Average.” Working paper, Boston University.

Contributors

Daron Acemoglu
Department of Economics
Massachusetts Institute of
Technology
50 Memorial Drive
Cambridge, MA 02142-1347

Philippe Aghion
Collège de France
3 Rue D'Ulm
75005 Paris, France

Ajay Agrawal
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Susan Athey
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305

James Bessen
Technology & Policy Research
Initiative
Boston University School of Law
765 Commonwealth Avenue
Boston, MA 02215

Erik Brynjolfsson
MIT Sloan School of Management
100 Main Street, E62-414
Cambridge, MA 02142

Colin F. Camerer
Department of Economics
California Institute of Technology
1200 East California Boulevard
Pasadena, CA 91125

Judith Chevalier
Yale School of Management
135 Prospect Street
New Haven, CT 06520

Iain M. Cockburn
School of Management
Boston University
595 Commonwealth Avenue
Boston, MA 02215

Tyler Cowen
Department of Economics
George Mason University
4400 University Drive
Fairfax, VA 22030

Patrick Francois
Vancouver School of Economics
University of British Columbia
IONA Building, 6000 Iona Drive
Vancouver, BC V6T 2E8, Canada

Jason Furman
Harvard Kennedy School
79 John F. Kennedy Street
Cambridge, MA 02138

Alberto Galasso
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Joshua Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Avi Goldfarb
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Austan Goolsbee
University of Chicago Booth School
of Business
5807 S. Woodlawn Avenue
Chicago, IL 60637

Rebecca Henderson
Harvard Business School
Morgan Hall 445
Soldiers Field Road
Boston, MA 02163

Ginger Zhe Jin
Department of Economics
University of Maryland
3115F Tydings Hall
College Park, MD 20742-7211

Benjamin F. Jones
Department of Management and
Strategy
Kellogg School of Management
Northwestern University
2211 Campus Drive
Evanston, IL 60208

Charles I. Jones
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305-4800

Daniel Kahneman
Woodrow Wilson School
Princeton University
Princeton, NJ 08544-1013

Anton Korinek
University of Virginia
Monroe Hall 246
248 McCormick Road
Charlottesville, VA 22904

Mara Lederman
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, Ontario M5S 3E6, Canada

Hong Luo
Harvard Business School
Morgan Hall 241
Soldiers Field Road
Boston, MA 02163

John McHale
108 Cairnes Building
School of Business and Economics
National University of Ireland
Galway H91 TK33, Ireland

Paul R. Milgrom
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305-6072

Matthew Mitchell
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Alexander Oettl
Scheller College of Business
Georgia Institute of Technology
800 West Peachtree Street NW
Atlanta, GA 30308

Andrea Prat
Columbia Business School
Uris Hall 624
3022 Broadway
New York, NY 10027-6902

Manav Raj
Stern School of Business
New York University
44 West Fourth Street
New York, NY 10012

Pascual Restrepo
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215

Daniel Rock
MIT Sloan School of Management
100 Main Street, E62-365
Cambridge, MA 02142

Jeffrey D. Sachs
Center for Sustainable Development,
Earth Institute
Columbia University
535 West 116th Street, MC 4327
New York, NY 10027

Robert Seamans
Stern School of Business
New York University
44 West 4th Street, KMC 7-58
New York, NY 10012

Scott Stern
MIT Sloan School of Management
100 Main Street, E62-476
Cambridge, MA 02142

Betsey Stevenson
Gerald R. Ford School of Public Policy
University of Michigan
5224 Weill Hall
735 South State Street
Ann Arbor, MI 48109-3091

Joseph E. Stiglitz
Columbia University
Uris Hall 212
3022 Broadway
New York, NY 10027

Chad Syverson
University of Chicago Booth School of
Business
5807 S. Woodlawn Avenue
Chicago, IL 60637

Matt Taddy
University of Chicago Booth School of
Business
5807 S. Woodlawn Avenue
Chicago, IL 60637

Steven Tadelis
Haas School of Business
University of California, Berkeley
545 Student Services Building
Berkeley, CA 94720

Manuel Trajtenberg
Eitan Berglas School of Economics
Tel Aviv University
Tel Aviv 69978, Israel

Daniel Trefler
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Catherine Tucker
MIT Sloan School of Management
100 Main Street, E62-536
Cambridge, MA 02142

Hal Varian
School of Information
University of California, Berkeley
102 South Hall
Berkeley, CA 94720-4600

Author Index

- Abadie, A., 531
Ablon, L., 448
Abrahamson, Z., 420
Abramovitz, M., 32
Acemoglu, D., 23n1, 43, 89, 90, 105, 127, 141, 152, 197, 198, 201, 202, 203, 203n4, 204, 204n5, 205, 206, 208, 210, 211, 212n7, 219, 220, 220n11, 223, 224, 225n14, 238, 240, 243, 271, 283, 293n3, 376n27, 376n28, 554
Acquisti, A., 410, 416, 424, 440n2, 444, 448, 451, 457, 457n42, 459, 483n19
Adee, S., 426
Agarwal, A., 82n12, 83
Aghion, P., 122, 172, 262, 262n19, 263, 265, 267, 268, 373n23, 465, 477, 479, 495
Agrawal, A., 5, 39, 90, 97n8, 150, 161, 167, 241, 425, 464, 501
Aguiar, M., 386n36
Airoldi, E., 80
Akerlof, G., 378
Akerman, A., 7, 303
Alexopoulos, M., 555
Allen, R. C., 204, 209
Alloway, T., 29
Alon, T., 43
Alpaydin, E., 92
Altman, S., 325
Alvarez-Cuadrado, F., 241n4
Anderson, R., 406, 416
Andrade, E. B., 592n3
Andrews, D., 30
Angermueller, C., 168
Aral, S., 43
Arntz, M., 321
Aronoff, M., 497
Arrow, K., 9, 43, 110, 118, 148, 364n11, 366, 412
Arthur, B. W., 150
Asher, S., 525
Athey, S., 68, 425, 448, 449, 510, 514, 515, 516, 517, 519, 523, 524, 525, 527, 528, 529, 530, 531, 532, 533, 534, 536, 538
Atkeson, A., 42n19
Autor, D. H., 7, 23n1, 30, 89, 198, 202n3, 203, 208, 220n11, 238, 239n3, 240, 271, 309, 322, 475, 555
Axelrod, R., 413
Ayres, R., 201
Azoulay, P., 475, 477
Babcock, L., 592n5
Bai, J., 532
Baker, D., 366n14
Baker, G., 107
Banbura, M., 537
Barkai, S., 271
Barrat, J., 350
Baslandze, S., 263, 264
Bastani, H., 529
Baumol, W., 38n14, 238
Bayati, M., 529, 531, 532, 536
Belloni, A., 93, 522
Bengio, S., 401

- Bengio, Y., 71, 74, 75, 75n6, 149, 168, 168n6, 172
- Benzell, S. G., 90, 336
- Berg, A., 373n24
- Beron, K. J., 483
- Bertocchi, G., 430
- Bessen, J. E., 23n1, 203, 300, 413, 555
- Bhatt, M. A., 592n4
- Bickel, P. J., 522
- Binmore, K., 589, 589n1, 593
- Bjorkegren, D., 516, 519
- Blake, T., 582
- Blattberg, R. C., 598
- Blei, D. M., 507, 510, 515, 532, 533, 534
- Bloom, N., 27, 150, 223, 259n15, 268, 560
- Bolton, P., 92, 95, 95n5, 96n7, 100, 112
- Boppart, T., 241n4, 295
- Borenstein, S., 413
- Bornmann, L., 153
- Bostrom, N., 3, 286, 381, 382, 382n31
- Bottou, L., 79
- Bousquet, O., 79
- Bowen, W., 38n14
- Brandeis, L. D., 431, 459
- Brander, J. A., 473
- Brandimarte, L., 448
- Brandt, L., 471
- Bresnahan, T. F., 4, 39, 42, 116, 119, 120, 169, 176n2, 310
- Bridgman, B., 38
- Brooks, R., 124
- Broseta, B., 593n9, 594n9
- Brown, J., 314
- Brunskill, E., 528
- Brynjolfsson, E., 23, 23n1, 28n7, 30, 39, 40, 42, 43, 47, 50, 89, 119, 120, 150, 197, 201, 204, 309, 555, 557, 560, 563
- Brzeski, C., 556
- Buera, F. J., 293, 293n2
- Buffie, E. F., 373n24
- Bughin, J., 408
- Burk, I., 556
- Busch, M. L., 474, 475n13
- Byers, J. W., 577n5
- Byrne, D. M., 29, 319
- Camerer, C. F., 589, 589n1, 590, 593, 594n10, 595, 598
- Campbell, K., 450
- Cardarelli, R., 29
- Carley, M., 130, 132
- Case, A., 26
- Catalini, C., 150, 425, 448, 449, 454
- Cavusoglu, H., 450
- Cette, G., 27
- Chandler, A. D., 206
- Chapelle, O., 529
- Chapman, J. P., 599
- Chapman, L. J., 599
- Chernozhukov, V., 93, 522, 523, 527
- Chetty, R., 270
- Chevalier, J., 577n5
- Chiou, L., 428, 448
- Chong, J.-K., 593, 595
- Christie, W. G., 414
- Chui, M., 331
- Clark, C., 295
- Cockburn, I., 150
- Coey, D., 514
- Cohen, J., 555
- Cohen, L., 501
- Comin, D., 241n4, 295
- Corrigan, B., 597
- Costa-Gomes, M. A., 593, 593n9, 594n9
- Courville, A., 71, 75n6
- Cowen, T., 27, 150
- Cowgill, B., 562
- Cranor, L. F., 424, 458
- Crawford, V. P., 593, 593n9, 594n9
- Crisuolo, C., 30
- Dana, J., 599
- Danaylov, N., 253
- Dasgupta, P., 366n15
- Datta, A., 433
- Dauth, W., 556
- David, P. A., 4, 41, 42n19, 119
- Davies, R. B., 483
- Dávila, E., 354
- Dawes, R. M., 597, 598, 599, 599n15
- Dawsey, K., 29
- Deaton, A., 26, 85
- Della Vigna, S., 411
- Delli Gatti, D., 361, 362n9, 380
- De Loecker, J., 30
- Dennis, B. N., 293n2
- Dewatripoint, M., 112
- Diamond, A., 531
- Dietvorst, B. J., 426
- Dimakopoulou, M., 529
- Dimico, A., 430
- Dobson, W., 485n23
- Dogan, M., 107n9
- Doi, E., 599
- Domingos, P., 93
- Doms, M., 562

- Dorantes, C., 450
 Dorn, D., 208
 Dosi, G., 366n14
 Doudchenko, N., 531
 Dover, Y., 577n5
 Drandakis, E., 376n27
 Dubé, J.-P., 410
 Dudik, M., 528
 Duffo, E., 539
 Dunne, T., 562
 Dupuit, J., 296, 298
 Duranton, G., 480, 567
 Dwork, C., 453
 Dzamba, M., 116

 Eckles, D., 538
 Edwards, D. D., 598
 Edwards, J. S., 598
 Edwards, L., 294
 Edwards, W., 595
 Eeckhout, J., 30
 Egami, N., 510
 Einhorn, H. J., 598, 599
 Elkin-Koren, N., 583
 Elsby, M. W. L., 270, 329
 Engel, E., 295
 Engerman, S. L., 430
 Engstrom, R., 537
 Erlingsson, U., 453
 Ethier, W. J., 476
 Etzioni, O., 410
 Evenson, R. E., 140
 Ezrachi, A., 414

 Fajgelbaum, P., 480
 Farrell, J., 424
 Faure-Grimaud, A., 92, 95, 95n5, 96n7,
 100, 112
 Faust, D., 597
 Fehr, E., 358n8
 Feldman, M. P., 560
 Feldstein, M., 29
 Feraud, R., 529
 Fernald, J. G., 27, 29, 32, 319
 Feurer, M., 70
 Filippas, A., 577n5
 Finklestein, A., 503
 Fleming, L., 161, 479
 Florêncio, D., 452
 Foellmi, R., 295
 Forbes, S., 108
 Ford, M., 201, 204, 301
 Forsythe, R., 590n2

 Fortunato, M., 25
 Fredriksson, P. G., 483
 Frey, C. B., 291, 321, 331, 350n1, 553, 555,
 556
 Friedman, J., 93
 Frosst, N., 76n7
 Fudenberg, D., 413, 594
 Fujii, H., 465
 Fung, A., 459
 Furman, J. L., 30, 122, 140, 150, 553

 Gaarder, I., 7, 303
 Gabaix, X., 604
 Gaggl, P., 303
 Gal, M. S., 583
 Gal, P., 30
 Galasso, A., 495, 496, 497, 498, 501
 Gans, J., 5, 39, 90, 97n8, 171, 425, 454, 464,
 550
 Garicano, L., 267
 Geanakoplos, J., 354, 368n19
 Geank, 369
 Gehring, J., 25
 Gentszkow, M., 411
 Gibbons, R., 107
 Glaeser, E. L., 516, 537, 548
 Goel, S., 516
 Goeree, J., 593
 Goh, G., 151n1, 168
 Goldberg, L. R., 597, 601n19
 Golden, J. M., 577n5
 Goldenshluger, A., 529
 Goldfarb, A., 5, 39, 90, 97n8, 148, 150,
 161, 167, 425, 427, 448, 464, 482, 483,
 483n19, 550, 593n8
 Goldin, C., 209, 322
 Goldman, M., 536
 Gomez-Uribe, C., 603
 Good, I. J., 238, 253
 Goodfellow, I., 66, 71, 75n6, 401
 Goolsbee, A. D., 310, 312, 314
 Gopalan, P., 510
 Gordon, R. D., 305, 310, 349
 Gordon, R. J., 27, 150, 175, 210, 223,
 259n15, 264
 Graepel, T., 429
 Graetz, G., 201, 274, 319, 554
 Graff Zivin, J. S., 475, 477
 Graham, M., 459
 Green, J., 502
 Greenstein, S., 119
 Greenwald, B., 354, 364n11, 366n15,
 368n19, 369, 370, 412

- Gregory, T., 321
Griliches, Z., 116, 120, 170
Grissen, D., 516, 519
Groover, M., 201, 205
Groshen, E. L., 350
Gross, R., 451
Grossman, G. M., 464, 476, 477, 478, 479, 480, 482
Grove, W. M., 597
Guerrieri, V., 293n3
Gurun, U., 501
Gutiérrez, G., 30
Guvenen, F., 29
- Ha, Y., 599
Hahn, J., 522
Haile, P. A., 515
Hainmueller, J., 531
Hall, B. H., 132
Hall, R. E., 47
Hansen, C., 93, 522
Hanson, G. H., 208
Hanson, R., 386n36
Harari, Y., 383
Hardt, M., 80
Harrell, E., 443, 444
Hart, O., 269
Hartford, J., 68, 75, 79n9, 526, 594
Haskel, J., 18
Hastie, T., 93, 526
Hatzius, J., 29
Haugeland, J., 69
Hawkins, J., 93
Hay, B., 494n2, 499
Hazan, E., 408
He, K., 75
Heckman, J. J., 85, 182
Heifels, A., 116
Helpman, E., 169, 464, 476, 477, 478, 479, 480
Hemous, D., 241
Hendel, I., 412
Henderson, R., 42, 43, 150
Herley, C., 452
Herrendorf, B., 204, 241n4
Hersh, J., 537
Hicks, J., 350, 368
Hillis, A., 516
Hinton, G. E., 68, 74, 76n7, 124, 125, 149, 168, 168n6, 172
Hirschman, A., 178
Hitt, L., 42, 43, 120
Ho, T.-H., 592n3, 593, 595
Hobijn, B., 270, 329
Hoch, S. J., 598
Hochreiter, S., 73, 76
Hodas, N., 151n1, 168
Hoffman, D., 444, 457
Hoffman, M., 447n17
Hofman, J. M., 510
Hogarth, R. M., 598
Holmes, T. J., 43
Holmstrom, B., 112, 269
Holt, C., 593
Hoos, H., 575
Hornik, K., 74, 525
Hortaçsu, A., 43, 593n8
Horton, J. J., 577n5
Hotborn, T., 525
Howitt, P., 122, 262n19, 263, 477
Hubbard, F., 494n2, 500
Huber, P., 494
Hunt, N., 603
Hutter, F., 575
Hutter, M., 237n1
- Imbens, G. W., 68, 517, 522, 523, 524, 525, 529, 530, 536, 538
Iriberry, N., 593, 594n10
Irwin, D. A., 477
İşçan, T. B., 293n2
- Jaffe, A. B., 132, 150
Jaravel, X., 312
Jarrell, G., 502
Jayadev, A., 366n14
Jean, N., 537
Jewitt, I., 112
Jha, S., 94, 95
Jiang, N., 528
Jin, G. Z., 442n4
Joachims, T., 528
Johnson, E. J., 589, 593, 603
Jones, B. F., 127, 150, 155, 161, 259, 259n15, 373n23
Jones, C. I., 46, 150, 152, 159, 160, 170n9, 171n9, 240, 252, 259n15, 373n23
Jones, R. W., 472
Jordan, M. I., 509, 532
Jovanovic, B., 41
- Kaboski, J. P., 293, 293n2
Kahn, L. B., 447n17
Kahneman, D., 596, 600, 604
Kaldor, N., 240
Kallus, N., 528

- Kaplan, J., 3
 Kaplow, L., 500
 Kapur, D., 150
 Karabarounis, L., 270, 329
 Karagözoğlu, E., 589
 Katz, L. F., 7, 209, 322
 Kehoe, P. J., 42
 Kehrig, M., 270
 Kendrick, J. W., 43
 Kennan, J., 590n2
 Kennedy, C., 376
 Keynes, J. M., 379
 Kislev, Y., 140
 Kitagawa, T., 527
 Klayman, J., 599
 Kleinberg, J., 15, 507, 516, 518, 548, 588, 597
 Klenow, P. J., 310, 312, 477
 Klette, T. J., 479
 Kleven, H. J., 270
 Knetsch, L., 604
 Ko, M., 450
 Kogler, D. F., 560
 Koh, D., 329
 Kohno, T., 446
 Kolb, D. A., 207
 Kollmeyer, C., 293n2
 Komarova, T., 538
 Kominers, D., 501, 516
 Kongsamut, P., 241n4, 295
 Korinek, A., 354, 365, 366n14, 371, 373, 383, 384
 Korolova, A., 453
 Kortum, S. S., 259, 259n15, 479
 Kosinski, M., 429
 Kotlikoff, L. J., 336, 344
 Krajbich, I., 592n6
 Kravitz, L., 325
 Krisiloff, M., 325
 Krizhevsky, A., 68, 74, 125
 Krueger, A. B., 7, 482
 Krugman, P. R., 464, 472, 479, 480
 Krussell, P., 265
 Künzel, S., 527
 Kurakin, A., 401
 Kurakin, P. R., 464
 Kurzweil, R., 238, 253, 253n12, 350, 373n23, 381, 382
 Kuznets, S., 205, 206

 Lada, A., 540
 Laffont, J.-J., 514
 LaGarda, G., 336
 Laibson, D., 606
 Lambrecht, A., 425, 426
 Lancaster, K., 362
 Landes, D., 206
 Lane, J., 554, 562
 Langford, J., 528
 Lanier, J., 84
 Lashkari, D., 241n4, 295
 Lawrence, R. Z., 294
 LeCun, Y., 75, 149, 168, 168n6, 172
 Lederman, M., 108, 550
 Legg, S., 237n1
 Leung, M. K. K., 121
 Levin, J., 514, 515
 Levine, D. K., 43
 Levine, S., 38
 Levinthal, D., 161
 Levitt, S. D., 43
 Levy, F., 37, 202n3, 238, 239n3
 Lewis, E., 562
 Lewis, G., 526, 527
 Lewis, M., 596
 Lewis, R. A., 541n2
 Lewis, W. A., 294
 Leyton-Brown, K., 569, 575, 594
 Li, D., 447n17, 529
 Li, L., 528, 529, 580, 581
 Liang, A., 588, 594
 Lim, K., 465, 479
 Lindor, R., 494
 Lipsey, R., 362
 List, J. A., 43
 Litan, R. E., 441n2, 494
 Liu, Y., 38
 Long, N., 241n4
 Lovallo, D., 600
 Lowenstein, G., 448, 592n5
 Lu, Y., 529
 Luo, H., 495, 496, 497, 498, 501
 Lusinyan, L., 29
 Lusted, L. B., 94

 Malthus, T. R., 384
 Mamoshina, P., 168n7
 Managi, S., 465
 Mandel, M., 555
 Manning, A., 379
 Mantoux, P., 200, 203
 Manuelli, R. E., 205, 243n6
 Manyika, J., 331
 Marchant, G., 494n2
 Marco, A., 130, 132
 Markov, J., 93

- Markusen, J. R., 476
Marthews, A., 425
Marx, M., 479
Massey, C., 426
Masterov, D. V., 578, 579
Matsuyama, K., 294, 295
Mayer, U. F., 578, 579
Mayer, W., 471
Mayzlin, D., 577n5
McAfee, A., 23, 23n1, 30, 40, 50, 89, 150, 201, 204, 309, 555
McCall, J. J., 582
McClelland, J. L., 602
McDonald, A. M., 424, 458
McElheran, K., 560
McFadden, D., 85, 514
Mcguckin, R. H., 562
McHale, J., 150, 241
McLaren, J., 469
McSweeney, T., 583
Meade, J. E., 362
Meadows, M., 420
Meehl, P. E., 596, 597
Meltz, M. J., 488
Mestieri, M., 241n4, 295
Mian, A., 208
Michaels, G., 201, 274, 319, 554
Mikolov, T., 76
Milgrom, P. R., 43, 111, 567, 569, 574, 577
Miller, A., 428, 429, 448, 483
Miller, S. M., 201
Milliment, D. L., 483
Minsky, M., 28, 208
Miremadi, M., 331
Mishel, L., 322
Misra, S., 410
Mitchell, T., 39, 557, 563
Mnih, V., 84
Mobius, M. M., 510
Mogstad, M., 7, 303
Mojon, B., 27
Mokyr, J., 29, 121, 150, 175, 201, 206, 209
Monro, S., 78
Moore, G. E., 384n35, 498
Moore, M., 496
Morris, D. Z., 37
Mortensen, D. T., 379, 582
Muellbauer, J., 85
Mullainathan, S., 169n8, 511, 512, 518, 588
Murdoch, J. C., 483
Murnane, R. J., 202n3, 238, 239n3
Murphy, K., 107
Murray, C., 325
Murray, F., 141
Mussa, M. L., 471
Mutz, R., 153
Myers, A., 130
Myerson, R. B., 568
Naecker, J., 592
Naik, N., 516, 537
Nakamura, L., 29
Nave, G., 590
Neal, R. M., 74
Neelin, J., 589n1
Neiman, B., 270, 329
Nekipelov, D., 538
Nelson, P., 581
Nelson, R., 118, 161
Newbery, D., 362
Newell, A., 123
Newhouse, D., 537
Ng, A., 93, 509, 532
Ng, S., 532
Ngai, L. R., 241n4, 294
Nickell, S., 293n2
Nielsen, M., 152, 161
Nilsson, N., 122, 207, 208
Nordhaus, W. D., 28, 152, 172, 238
North, D. C., 577
Nosko, C., 577n5, 582
O'Dea, B., 583
Odlyzko, A., 441n2
Oettle, A., 241
Olmstead, A. L., 204, 205, 206
Olsen, M., 241
O'Mahony, S., 141
O'Neil, C., 433
O'Reilly, B., 499
Orlikowski, W. J., 43
Orszag, P., 30
Osborne, M. A., 291, 321, 331, 350n1, 553, 555, 556
Osindero, S., 74
Oskamp, S., 601
Ossard, H., 514
Östling, R., 593n8
Ostrovsky, M., 568
Pajarinen, M., 556
Pál, J., 510
Palfrey, T., 593
Parchomovsky, G., 494

- Pate, R. H., 420
 Pelzman, S., 502
 Peretto, P. F., 241
 Peterson, N., 599
 Peysakhovich, A., 540, 592
 Phelps, E., 376n27
 Philippon, T., 30
 Pierce, D. G., 413
 Pihur, V., 453
 Piketty, T., 360
 Pissarides, C. A., 241n4, 294, 379
 Polemarchakis, H., 354, 368n19, 369
 Polinsky, M., 494n2
 Porter, M. E., 481, 494
 Poschke, M., 241n4
 Posner, R. A., 439
 Prantl, S., 263
 Pratt, G. A., 40
 Proserpio, D., 577n5
 Puga, D., 480

 Raghaven, M., 518
 Ramaswamy, R., 291, 293n2
 Ramlogan, R., 480
 Rao, J. M., 516, 536, 541n2
 Rasmussen, W. D., 200, 203, 205n6, 206
 Rauch, J. E., 469
 Rebelo, S., 241n4, 295
 Recht, B., 80
 Redding, S., 293n2
 Reinsdorf, M. B., 29
 Rennie, J., 80
 Restrepo, P., 23n1, 43, 90, 105, 127, 152, 197,
 201, 202, 203, 203n4, 204, 204n5, 206,
 208, 210, 211, 212n7, 219, 220, 220n11,
 223, 224, 225n14, 238, 240, 241, 243,
 271, 283, 376n27, 379, 554
 Rhode, P. W., 204, 205, 206
 Rhodes, E., 325
 Rivera-Batiz, L. A., 477, 478
 Rivkin, J., 161
 Robbins, H., 78
 Roberts, J., 43, 111
 Robinson, P. M., 524, 527
 Rock, D., 557
 Rodrik, D., 293
 Roesner, F., 446
 Rogerson, R., 204, 241n4
 Romanosky, S., 444, 450, 457, 457n42
 Romer, P. M., 46, 122, 149, 153, 155n2, 159,
 169, 171, 171n9, 172, 255, 477, 478
 Rosenberg, N., 150, 169

 Rosenblatt, F., 73, 124
 Rosenfeld, J., 322
 Rossi-Hansberg, E., 476
 Roth, A. E., 539, 567, 584, 589
 Rousseau, P. L., 41
 Rouvinen, P., 556
 Rowthorn, R., 291, 293n2
 Rubin, D. B., 517, 522, 527, 529
 Rubinstein, A., 427
 Ruffin, R. J., 472
 Ruiz, F. J., 515, 533, 534
 Rumelhart, D. E., 73, 77, 124, 602

 Sabour, S., 76n7
 Sachs, J. D., 336, 344
 Saez, E., 270, 360
 Şahin, E., 270, 329
 Salakhutdinov, R. R., 125
 Salomons, A., 23n1, 309, 555
 Samuelson, P., 376
 Santaaulalia-Llopolis, R., 329
 Saon, G., 25
 Sawyer, J., 597
 Saxenian, A.L., 4, 478
 Schierholz, H., 322
 Schmidhuber, J., 73, 76
 Schmidt, K. M., 358n8
 Schmitt, J., 322, 379
 Schmitz, J. A., 43
 Schmucki, R., 497
 Schultz, P. H., 414
 Schumpeter, J., 148
 Schwartz, M., 568
 Scotchmer, S., 121, 496, 502
 Scott, S. L., 528
 Seater, J. J., 241
 Segal, I., 414, 569, 574
 Seira, E., 514
 Seshadri, A., 205, 243n6
 Shah, A. K., 517
 Shaked, A., 589n1
 Shavell, S., 494n2
 Shaw, J. C., 123
 Shiller, B. R., 410
 Shroff, R., 516
 Silver, D., 63, 66, 453
 Simcoe, T., 487
 Simmons, J. P., 426
 Simon, H. A., 107, 123, 207–8
 Sims, C., 113
 Simsek, A., 362n10
 Singer, Y., 80

- Slovic, P., 596
Smith, A., 590
Smith, M. D., 43
Smith, N., 29
Sokoloff, K. L., 430
Solomonoff, R. J., 253n12
Solove, D., 447, 455
Soloveichik, R., 29
Solow, R. M., 24, 46, 350
Somanchi, S., 444
Song, J., 30
Sonnenschein, H., 589n1
Sopher, B., 590n2
Sorensen, O., 161
Spencer, B. J., 473
Spezio, M., 594n10
Spiegel, M., 589n1
Spiegel, Y., 412
Spier, K., 494, 499, 502
Spiess, J., 169n8, 511, 512
Spindler, M., 522
Srivastava, N., 80
Stahl, D. O., 593
Stantcheva, S., 360
Stein, A., 494
Stern, A., 325
Stern, S., 122, 141
Stevenson, B., 194n5
Stigler, G. J., 439, 582
Stiglitz, J. E., 30n9, 354, 360, 362, 364,
364n11, 366n14, 366n15, 368n19, 369,
370, 370n21, 371, 373, 376n27, 376n28,
378, 412
Stillwell, D., 429
Stinchcombe, M., 74
Stivers, A., 442n4
Stole, L., 268
Strehl, A., 527
Streitwieser, M. L., 562
Strotz, R. H., 427
Stucke, M. E., 414
Stutzman, F., 451
Sufi, A., 208
Summers, L. H., 175
Sutskever, I., 68, 74, 125
Sutton, J., 467, 589n1
Swaffield, J., 293n2
Swaminathan, A., 528
Sweeney, L., 433
Swire, P. P., 441n2
Syverson, C., 23, 29, 43, 210, 319, 320
Taddy, M., 83, 526, 527
Tadelis, S., 107, 577n5, 578, 579, 580, 581,
582
Tang, J., 453
Taylor, C. R., 416, 424, 441n2, 448, 459,
483n19
Tegmark, M., 382, 384n34
Teh, Y.-W., 74
Telang, R., 444, 450, 457, 457n42
Teodoridis, F., 150, 161
Tetenov, A., 527
Thaler, R., 604
Thomas, K., 443
Thomas, P., 528
Thompson, W. R., 82
Thrnton, B., 598
Tibshirani, R., 93, 525, 527
Tirole, J., 106, 112, 264, 268, 269
Topol, E. J., 94, 95
Tory, J., 485n23
Toulis, P., 80
Trajtenberg, M., 4, 39, 116, 119, 132, 150,
169, 176n2
Trefler, D., 465, 467, 479, 485n23
Troske, K. R., 562
Tschantz, M. C., 433
Tucker, C. E., 148, 425, 426, 427, 428, 429,
448, 449, 482, 483, 483n19
Turing, A. M., 123, 385
Tuzel, S., 210
Tversky, A., 596
Ugander, J., 538
Uyarra, E., 480
Vadlamannati, K. C., 483
Valentinyi, Á., 204, 241n4
Van der Laan, M. J., 522, 527
Van Seijen, H., 62, 84
Varian, H. R., 93, 310, 410, 413, 424, 425,
440, 511
Venables, A. J., 480
Vesteger, M., 420
Vickrey, W., 567
Vijverberg, W. P. M., 483
Vincent, N., 270
Vines, P., 446
Vinge, V., 238, 253, 253n12, 373n23
Viscusi, K., 496, 498
Vishnu, A., 151n1, 168
Von Hippel, E., 499

- Von Weizacker, C. C., 376
Von Winterfeldt, D., 595
Vuong, Q., 514
- Wager, S., 523, 525, 527, 528
Wagman, L., 416, 424, 441n2, 448, 459, 483n19
Wallach, I., 116
Wan, M., 534
Wang, J., 475, 477, 594n10
Warren, S. D., 431, 459
Waseem, M., 270
Wattal, S., 450
Weil, D., 459
Weingast, B. R., 577
Weiss, A., 354
Weitzman, M., 149, 151, 157, 171, 171n9, 172, 241
Western, B., 322
Westlake, S., 18
Whinston, M. D., 148
White, H., 74, 511
Williams, H., 122, 364n11
Williams, R., 124
Wilson, P. W., 593
Winter, S., 161
Wolfers, J., 195n5
Wooldridge, J. M., 522
Wright, G., 42n19
Wright, G. C., 303
- Wright, J. R., 594
Wu, L., 43
- Xiao, M., 593n8
Xie, D., 241n4, 295
Xu, H., 577
- Yakovlev, E., 538
Yang, S., 42, 47
Yellen, J., 378
Yeomans, M., 517
Yildirim, P., 107n9
Yu, M., 465, 479
Yudkowsky, E., 253n12
- Zanna, L.-F., 373n24
Zeevi, A., 529
Zeileis, A., 525
Zeira, J., 152, 198, 212n7, 238, 239, 241
Zervas, G., 577n5
Zhang, M. B., 210
Zheng, Y., 329
Zhou, D., 529, 581
Zhou, X., 580
Zhu, X., 471
Zierahn, U., 321
Zubizarreta, J. R., 523
Zweimüller, J., 295
Zwiebel, J., 268

Subject Index

Note: Page numbers followed by “f” or “t” refer to figures or tables, respectively.

- adoption, technological: implications of speed of, for job market and inequality, 310–12
- adversarial artificial intelligence, 401
- aggregate productivity statistics, technologies and, 26–28
- AI. *See* artificial intelligence (AI)
- AlphaGo (Google), 63
- Amazon Go concept, 67
- applied artificial intelligence, 208
- artificial intelligence (AI), 1; and automation of production, 239–50; as basis for learning, 120–21; benefit of more, 318–20; bibliometric data on evolution of, 128–32; capital shares and, 270–74, 272–73f; in context, 84–85; defined, 3–4, 62–67, 93, 122, 237, 468; economies of scale from data and, 468–69; economies of scale from overhead of developing AI capabilities, 469–70; evolution of, 122–25; expected productivity effects of, 45–46; firm organization and, 264–70; future of research on economics of, 17; as general purpose technology, 4–7, 39–41; in idea production function, 250–52; impact of, on innovation, 125–28; impact of long-term decline in labor force participation rate and, 323–25; implications of, 349–53; income distribution and, 351; inequality and, 320–23; internal agreements and, 463; international macroeconomics and, 488; knowledge externalities and, 470–71; likely productivity effects of, and acceleration of, 45–46; longer-term prospects of, 381–86; market structure and, 262–63; as “next big thing,” 175; political economy of, 11, 394–95; prediction costs and, 92–93; privacy and, 425–26; privacy concerns and, 423–24; for promoting trust in online marketplaces, 576–81; recent approach to, 93; for reducing search frictions, 581–83; return of Malthus and, 381–86; revolution, international effects of, 393–94; in Schumpeterian model with creative destruction, 276–79; sectoral reallocation and, 263–64; statistics on, 465–66, 466t; studies on economic effect of, 556–58; theory of privacy in economics and, 424–26; as tool, 16–17; world’s largest companies and exposure to, 465–67, 467t. *See also* machine learning (ML)
- artificial intelligence capital, measuring, 46–50
- artificial intelligence–general purpose technology (GPT) era: education strategies for, 179–82; human-enhancing innovations vs. human-replacing innovations

- artificial intelligence—general purpose technology (GPT) era (*continued*)
 for, 184–85; professionalization of personal services strategies for, 182–84; top skills required for employment in, 180–81, 181t
- artificial intelligence revolution, international effects of, 393–94
- Atomwise, case of, 115–16, 120, 154
- automatic teller machines (ATMs), security policy and, 416
- automation, 3–4, 105–6; basic model, 336–41; Baumol's cost disease and, 241–50; to date, and capital shares, 270–74; decline in labor share and, 329–31; deepening of, 198, 204–5, 216–17; economic adjustment and, 208–9; employment and, 190–91; excessive, 224–26; model of, 211–14; of production, and artificial intelligence, 239–50; productivity and, 210–11; sector of economy affected by, 330–33; studies on employment on, 555–58; wages and, 200–211; winners, 190; work and, 200–211; Zeira model of growth and, 239–41
- average treatment effects, 522–24
- bandits (algorithms), problem of, 528–29
- Baumol's cost disease. *See* cost disease, Baumol's
- BenchSci search technology, 153
- buy/make decisions (firm boundaries), 107–8
- capital accumulation, 198, 204–5, 216
- capital shares, and automation to date, 270–74
- causal inference, new literature on, 519–34
- Children's Online Privacy Protection Act of 1998 (COPPA), 454
- cloud-computing facilities, 402–3
- cluster policies, 480–81
- clusters, regional, theory of, 479
- CNNs (convolutional neural networks), 75–76, 75n6
- collusion, strategies for facilitating, 413–14
- combinatorial-based knowledge production function, 154–61; potential uses of new, 170–71; with team production, 161–67
- Communications Act (1986), 456
- competition policy, innovations and, 141–43
- complexity, 103–4
- consumer attitude, 448–49
- consumer privacy: challenging issues in, 457–59; consumer attitude and, 448–49; consumer risk and, 443–48; data risk and, 442–43; nature of problem of, 442; policy landscape in United States, 454–57; supply side actions and, 450–54. *See also* privacy
- consumer surplus, 11; distribution of, 391–93
- contracting, 106–7
- convolutional neural networks (CNNs), 75–76, 75n6
- cooperation, evolution of, 414
- cost disease, Baumol's, 8–9, 238–39; automation and, 241–50
- creative destruction, 260–61
- data, 61; acquisition methods, 403–4; decreasing marginal returns of, 406, 407f; economics of, 14; importance of, 13–14; important characteristics of, 404–6; localization rules, trade policies and, 485–86; persistence of predictive power of, 427–28; privileged access to government, trade policies and, 486–87; types of, and generation of spillovers, 431–34
- data access, 405–6
- data generation, as pillar of artificial intelligence, 62f, 65–66
- data ownership, 405–6
- data persistence, 426–27; predictive power and, 427–28
- data pipeline, 402
- data pyramid, 404, 405f
- data repurposing, 428–31
- data security: challenging issues in, 457–59; policy landscape in United States, 454–57
- data spillovers, 431–34
- data warehouses, 402–3
- decision-making, baseline model for, 95–103; complexity and, 103–8
- deepening of automation, 198, 204–5, 216–17
- Deep Genomics, 154
- deep learning, 3, 71–77, 400; as general purpose invention in method of invention, 139–43; as general purpose technology, 133–39; as new discovery tool, 167–69; patent systems and, 142

- deep learning networks, 94
 deep learning techniques, 25
 deep neural networks (DNNs), 25–26, 61,
 63; structure in, 76–77
 demand, importance of, 301–2
 destruction, creative, 260–61
 difference-in-difference models, 530–31
 digital information, 334–35
 direct network effects, 412
 displacement effect, 8, 198, 208, 214
 DNNs. *See* deep neural networks (DNNs)
 domain structure, as pillar of artificial intel-
 ligence, 62f, 63–65
 double machine learning, 523–24
- economic growth: artificial intelligence and,
 262–70; prospects for technology-
 driven, 149–53; Zeira model of auto-
 mation and, 239–41. *See also* growth
 economics, impact of machine learning on
 practice of, 15–16
 education, factory model of, 180
 Electronic Communications Privacy Act of
 1986: (ECPA), 456
 employment: automation and, 190–91;
 levels of, and new technologies, 220–21;
 long-run vs. short run, 192–94; studies
 on automation and, 555–57; work out-
 side of, 194–95
 evolution of cooperation, 414
- factory model of education, 180
 Federal Trade Commission (FTC), 454–55
 firm boundaries (make/buy decisions),
 107–8
 firm-level data: need for, 558–59; strategies
 for collecting, 561–62
 firm-level research questions, 560–61
 firms: artificial intelligence and, 262–70;
 impact of machine learning on, 12
- general purpose machine learning (GPML),
 67–71
 general purpose technology (GPT), 2, 65,
 119–20, 169–70; artificial intelligence
 as, 4–7, 39–41; deep learning as, 133–
 39; viewing today's Solow paradox
 through previous history of, 44–45
 generative adversarial networks (GANs),
 66–67
 GPML (general purpose machine learning),
 67–71
- GPT. *See* general purpose technology (GPT)
 Gramm-Leach-Bliley Act (GLBA), 454
 growth: impact of artificial intelligence on,
 7–9. *See also* economic growth
- Health Insurance Portability and Account-
 ability Act of 1996 (HIPAA), 454
 HEI (human-enhancing innovations),
 184–85
 heterogeneous treatment effects, 524–28
 hierarchical Poisson factorization, 510
 HIPAA. *See* Health Insurance Portabil-
 ity and Accountability Act of 1996
 (HIPAA)
 HRI (human-replacing innovations), 184–85
 human-enhancing innovations (HEI), 184–85
 human-replacing innovations (HRI), 184–85
- idea production function, artificial intelli-
 gence in, 250–52
 implementation/restructuring lags, as expla-
 nation for Solow paradox, 31–36
 incentive auctions, machine learning and,
 569–76
 income, artificial intelligence and, 189–90
 income distribution: artificial intelligence
 and, 351; impact of AI on, 11
 income inequality: artificial intelligence and,
 320–23; impact of AI on, 7–8, 11–12;
 speed of technological adoption and,
 310–12
 income redistribution, political economy of,
 394–95
 indirect network effects, 412
 industrial regulation, trade policies and, 487
 inequality. *See* income inequality
 information technology (IT), 24
 innovation, 115–18; competition policy and,
 141–43; early stage, 121–22; impact
 of artificial intelligence on, 125–28;
 institutions and, 141–43; management
 and organization of, 140–41; product
 liability and, 494
 institutions, innovations and, 141–43
 intelligence-assisting innovation (IA),
 350–51
 International Federation of Robotics
 (IFR), 16
 international macroeconomics, artificial
 intelligence and, 488
 international trade, economics of data
 and, 14

- invention of a method of inventing (IMI), 120–21, 124
- inverted-U pattern, 293; simple model of, 297–301
- JDM (judgment and decision-making) research, 596–98
- job displacement, 310–12
- job losses, 291
- job markets, speed of technological adoption and, 310–12
- jobs, impact of artificial intelligence on, 7–8, 9–11
- judgment: in absence of prediction, 96–101; as complements/substitutes to prediction, 102–3; creating role for, 91; prediction and, 91–92
- judgment and decision-making (JDM) research, 596–98
- knowledge creation, 477–79
- knowledge externalities, and artificial intelligence, 470–71
- knowledge spillovers, 479
- labor: comparative advantage of, and new tasks, 217–18; model of demand for, 211–14
- labor demand, technology and, 214–21
- labor productivity growth, technologies and, 26–28
- learning by doing, 412–13
- liability, innovation and, 494; empirical evidence on, 496–98; theoretical model of, 494–96
- liability, tort, development of artificial intelligence technologies and, 498–502
- machine learning (ML), 3, 24–25; applications of, 401–2; defined, 509–10; double, 523–24; early use cases of, 510–15; for finding behavioral variables, 587–96; general purpose, 67–71; human prediction as imperfect, 496–603; impact of, 12–13, 15–17, 507–9; incentive auctions and, 569–76; new literature on, 519–34; overview, 399–406; predictions about impact of on economics, 534–42; regulation and, 12–15; supervised, 511–12; unsupervised, 510–11; vertical integration and, 408–9. *See also* artificial intelligence (AI)
- machine learning–provision industries, 414–15
- machine learning–using industries, 408f; boundaries and, 409–10; firm size and, 409–10; price differentiation and, 410–11; pricing and, 410; returns to scale and, 411–14; structure of, 406–8
- macroeconomics, international, artificial intelligence and, 488
- make/buy decisions (firm boundaries), 107–8
- Maluuba, 62–63
- market design, introduction to, 567–69
- massive open online courses (MOOCs), 181
- matrix completion problem, 531–32
- matrix factorization, 51, 511
- meta technologies, 153
- ML. *See* machine learning (ML)
- model averaging, 511
- MOOCs (massive open online courses), 181
- neural networks, 72–74, 123, 124–25, 510, 511
- new economic geography (NEG), 479–80
- online marketplaces, using artificial intelligence to promote trust in, 576–81
- optical lenses, invention of, 121
- optimism, sources of technological, 24–26
- Pareto improvement, 363
- patent systems, deep learning and, 142
- Pen Register Act (1986), 456
- policy analysis, using methods of prediction for, 516–19
- political economy: of artificial intelligence, 11; of income redistribution, 394–95; of technological disruptions, 176–79
- prediction: in absence of judgment, 101–2; artificial intelligence as tool for, 16–17; as complements/substitutes to judgment, 102–3; costs of, and AI, 92–93; judgment and, 91–92; using methods of, in policy analysis, 516–19
- premature deindustrialization, 393–94
- price discrimination, artificial intelligence and, 604
- principal components analysis, 74, 92, 510
- privacy, 13–14; artificial intelligence and,

- 425–26; current models of economics and, 424–25; data spillovers and, 431.
See also consumer privacy
- privacy policy, 416–17
- privacy regulation, trade policies and, 482–85
- privileged access to government data, trade policies and, 486–87
- production, automation of, and artificial intelligence, 239–50
- productivity: automation and, 210–11; missing growth of, and new technologies, 223–26
- productivity effects, 198, 203–4, 214–16
- productivity growth: low current, reasons why it is consistent with future technological growth, 41–44; rates of, technologies and, 26–28; slow, and future productivity growth, 31–36
- productivity optimism, technology-driven case for, 36–39
- radiology, case of, 94–95
- random forest, 511
- R&D. *See* research and development (R&D)
- recommender systems, 603
- regional clusters, theory of, 479
- regression trees, 511
- regularization on norm of matrix, 510
- regularized regression, 511
- regulation, 12–15; machine learning and, 12–15
- reinforcement learning, 25, 66, 81–84, 400–401
- reinstatement effect, 8, 198, 206
- research and development (R&D), 336; productivity, effects of rise in, 341–43
- research tools, economics of new, 118–22
- robotics, 123–24; tort law and, 493–94.
See also robots
- robots, studies on, 554–55. *See also* robotics
- Romer/Jones knowledge production function, 151–52
- scale, economies of, and artificial intelligence, 470
- Schumpeterian model with artificial intelligence, 276–79
- scientific discovery, rate of, 6
- scientists, role of, 472–73; superstar, 474–76
- scope, economies of, and artificial intelligence, 470
- search frictions, artificial intelligence for reducing, 581–83
- security policy, 416
- singularities, 253–61; examples of technological, 254–58; objections to, 258–61
- skills: mismatch of technologies and, 221–23; technologies and, 209
- Solow paradox, 24; potential explanations for, 28–31
- source code, trade policies and, 487–88
- spectrum reallocation, 569–76, 572f
- spillovers, data, 431
- spreadsheet software, invention and impact of, 90
- stochastic gradient descent optimization, 77–81
- strategic trade policy, 473–74
- structural change, 293–96
- structural models, 532–34
- superstar scientists, role of, 474–76
- supervised machine learning, 511; methods for, 511–12
- supplementary analysis, 530
- support vector machines, 511
- symbolic processing hypothesis, 123
- symbolic systems, 123
- tasks, new, 205–7; comparative advantage of labor and, 217–18; creation of, 198; model of, 211–14
- technological changes: factor-biased, 376–77; and levels of employment, 220–21; types of, 212–13
- technological disruptions, political economy of, 176–79
- technological growth, reasons why it is consistent with low current productivity growth, 41–44
- technological optimism, sources of, 24–26
- technological progress: channels of inequality and, 365–70; determining scenarios that best describe economy, 363–64; endogenous, 364–65; first-best scenario, 353–56; imperfect markets scenario, 361–62; perfect markets but costly redistribution scenario, 358–61; perfect markets ex post and no costs of redistribution scenario, 356–58; welfare and, 353–65; worker-replacing, redistribution and, 370–77
- technological singularities, examples of, 254–58

- technological unemployment, 377–81
- technologies: future progress of, and low current productivity growth, 41–44; labor demand and, 214–21; mismatch of skills and, 221–23; skills and, 209
- technology-driven economic growth, prospects for, 149–53
- tort law, robotics and, 493
- tort liability, 14
- total factor productivity growth, 32
- trade models, basic, 476–77
- trade policies: data localization rules and, 485–86; industrial and strategic, case for, 471–81; industrial regulation and, 487; privacy regulation and, 482–85; privileged access to government data and, 486–87; role of university-related talent, 472–76; source code and, 487–88; strategic, 473–74
- UBI. *See* universal basic income (UBI)
- unemployment, technological, 377–81
- universal basic income (UBI), 312–14; cost of replacing current safety net with, 325–26
- unsupervised machine learning, 510–11
- USA Patriot Act (2001), 456
- vertical integration, machine learning and, 408–9
- vertical research spillovers, 122
- wages, automation and, 200–211
- welfare, technological progress and, 353–65
- work, automation and, 200–211
- worker-replacing technological progress: dynamic implications of, 373–74; redistributing innovators' surplus and, 374–76; redistribution and, 370–77; static pecuniary externalities of, 371–72
- Zeira model of automation and growth, 239–41
- “zero-shot” learning systems, 66