

Economists (and Economics) in Tech Companies

Susan Athey and Michael Luca

PhD economists have started to play an increasingly central role in tech companies, tackling problems such as platform design, pricing, and policy. Major companies, including Amazon, eBay, Google, Microsoft, Facebook, Airbnb, and Uber, have large teams of PhD economists working to engineer better design choices. For example, led by Pat Bajari, Amazon has hired more than 150 PhD economists in the past five years, making it the largest employer of tech economists. In fact, Amazon now has several times more full-time economists than the largest academic economics department, and continues to grow at a rapid pace. Companies such as Coursera, Expedia, Microsoft, Netflix, Pandora, Uber, Yelp, and Zillow have also hired economists. Table 1 shows a list of some technology companies that have hired PhD economists, although the list is not comprehensive.

Hiring of PhD economists has happened at all levels, from newly minted PhDs heading directly to the tech sector up through chief economists plucked from tenured positions at prestigious academic departments. The types of positions also vary greatly. Much of the recent growth has focused on economists working directly on business problems, with only a small fraction of the work resulting in academic papers. In contrast, some companies, such as Microsoft, have a chief economist who manages teams focused directly on business problems, but also have many economists working out of research centers, publishing self-guided research in academic

■ *Susan Athey is Economics of Technology Professor, Stanford Graduate School of Business, Stanford, California. Michael Luca is Lee J. Styslinger III Associate Professor of Business Administration, Harvard Business School, Boston, Massachusetts. Their email addresses are athey@susanathey.com and mluca@hbs.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at

<https://doi.org/10.1257/jep.33.1.209>

doi=10.1257/jep.33.1.209

Table 1

Examples of Tech Companies That Have Hired PhD Economists

Alibaba	Forkcast	LinkedIn	Redfin
AirBnB	Glassdoor	Lyft	Ripple
Amazon	Google	Microsoft	Rover
AppNexus	Granular	Netflix	Trulia
CoreLogic	Groupon	Nuna	Uber
Coursera	Houzz	Oath	Upwork
Dstillery	Huawei	OpenAI	Vericred
Didichuxing	IBM	Pandora	Visa
Digonex	Indeed	Pinterest	Walmart
eBay	ING	PoliticalSheepdog.com	Wealthfront
ECONorthwest	Intel	Prattle	Yahoo!
Expedia	Kensho	Quantco	Yelp
Facebook	Lending Club	Quora	Zillow

journals comparable to that of economists working in business schools or economics departments. These research centers, at their best, provide frontier insights, some of which will guide the future direction of the company.

Many tech companies now recruit directly through the American Economic Association's Job Openings for Economists (JOE) platform, which is where much of the recruiting for PhD economists begins. During the 2017–18 academic year, 21 tech companies were hiring through the JOE website. To put this into context, there are roughly half as many tech companies hiring through JOE as there are policy schools. Taking into account the fact that many of these companies had multiple positions, the number of positions available for economists in tech companies likely exceeded those at policy schools.

Moreover, Table 2 shows that the number of tech companies with job postings at JOE has generally risen in recent years. As technology platforms play an increasing role in the economy, topics relevant to them have become more important to the business school curriculum and to academic research in business schools. Business schools have seen increased demand for faculty specializing in online platforms and digitization, as well as in areas crucial to understanding data analysis, such as experimental methods and machine learning. For example, groups in business schools that historically focused on operations research or management of information systems have recently begun to focus more on economic problems such as marketplaces, pricing algorithms, and empirical studies of economic questions.

These shifts are partially driven by a growing need to prepare MBA students for a career in the technology sector. For example, Amazon was the largest employer of Harvard Business School's most recent graduating class of MBA students. Corresponding to the shifting career paths of MBA students, recent additions to the Harvard Business School curriculum in the past few years include courses on experimental methods, designing online marketplaces, digital marketing, technology strategy, and data science. Stanford's Graduate School of Business has seen similar

Table 2

The Number of Tech Companies, Policy Schools, Business Schools, and Economics Departments with Positions for PhD Economists

	<i>Tech companies</i>	<i>Policy school departments</i>	<i>Business school departments</i>	<i>Economics departments</i>
<i>August 2014—July 2015</i>	15	43	247	366
<i>August 2015—July 2016</i>	15	43	264	349
<i>August 2016—July 2017</i>	23	36	289	322
<i>August 2017—July 2018</i>	21	50	326	374

Note: Authors using data from Job Openings for Economists (JOE).

growth. More broadly, there has been a rapid expansion in courses directly related to the technology industry. Content related to the digital economy has increasingly been added to more traditional courses (such as core strategy and marketing courses) as well.

Within industry, there is little precedent for private companies recruiting academic economists as well as new PhDs with strong research skills so heavily for full-time positions. Organizations like the RAND Corporation and Mathematica Policy Research recruit economists on a large scale, but focus mainly on research and policy evaluations. Consulting firms like Cornerstone and the Analysis Group also recruit large numbers of economists, but primarily to support and serve as expert witnesses in legal matters in areas such as antitrust and intellectual property litigation.

Considering the tech firms that hire into research labs, such as Microsoft Research, perhaps the closest historical analog would be Bell Labs, which was operating as a division of AT&T when it created an economics team in 1968. The team grew to include about 30 economists, including high-profile economists such as Elizabeth Bailey, Roy Radner, and Robert Willig. In 1970, it launched the *Bell Journal of Economics and Management Science*, which lives on as the highly regarded *RAND Journal of Economics*. The team was phased out in 1983, coinciding with the breakup of AT&T. Some of its economists were folded into other parts of the company, while others left for other industry or academic jobs—including at Columbia University, Harvard Business School, New York University, Princeton University, and the University of Pennsylvania.

Although some tech companies hire economists using a lab model, the majority of economists in tech companies work on managerially relevant problems with data from the company, and many are in business roles. For example, outside of Microsoft Research, Microsoft has a business-focused chief economist whose team actively recruits PhD economists to work on problems ranging from cloud computing to search advertising. Amazon assigns economists to specific business problems across divisions, ranging from the e-commerce platform to digital content to the experimentation platform used to evaluate changes and innovations. Uber has teams

of economists focused on understanding policy issues in addition to pricing and incentive design—some of these teams produce outward-facing research published in academic journals while others are completely inward-facing. More broadly, many economists at tech companies do a combination of external research and internal work, continuing to attend conferences and publish in leading economics journals; they often hire summer interns from top PhD programs or collaborate with academic economists on such projects. Because many of the problems faced by tech companies are on the frontier of academic research, close ties to academics and rigorous, original thinking are highly valued in the tech sector.

Indeed, the interaction between tech companies and economists has given rise to new intellectual questions and a new field within economics—the “economics of digitization.” The field has explored a wide range of questions. For example, how does the advent of artificial intelligence and the use of large-scale consumer datasets affect industry structure and market power? How should tech companies be regulated? How should data from the tech sector inform policy? How do aggregators, search engines, reputation systems, and social media affect the decisions we make and the news we read? How should online marketplaces be designed to ensure safe and efficient transactions? Online platforms have also created novel datasets and testing grounds that have been used to inform virtually every field of economics, from market design to industrial organization to labor economics to behavioral economics.

We have had the opportunity to spend our careers thus far with one foot in academia, studying and teaching about online platforms, and the other in practice, helping to shape them. Outside of our academic roles, we work closely with tech companies. Susan previously served as consulting chief economist at Microsoft and currently sits on the boards of Expedia, Lending Club, Rover, and Ripple. While working with Microsoft, she also helped build the economics group at Microsoft’s research arm in New England. Mike works with a variety of tech companies, and created an economic research initiative at Yelp. As academics, we have taught hundreds of students and executives who now work in tech companies. Doctoral students have become interested in tech companies as well—our own students have worked at companies ranging from Facebook, Microsoft, and Amazon to Wealthfront, Uber, and Airbnb.

The core skills that economists use in tech companies have been important to economic research for decades prior to the tech era. The field of market design has been combining novel theoretical insights, empirical work, and experiments to solve real-world problems since Bob Wilson’s pioneering work on auctions in the 1960s. Assessing causal relationships and understanding incentives have been central themes in applied microeconomics and industrial organization for decades. With the advent of new technologies, the expertise developed by PhD economists has found new and influential uses in the tech sector. Furthermore, the frontiers of economic research in these areas has been advanced as the tech sector has simultaneously introduced new economic problems, provided new ways to bring ideas from economic theory into practice, and provided opportunities for new types of statistical analysis.

With the rise of economists in tech companies, we're frequently contacted by tech companies for recommendations about whom to hire and what types of roles economists should take on. We are also asked how undergraduates and PhD students can prepare for such careers, as well as what these careers will be like. Faculty are often interested in how they can get involved with tech companies, and what types of problems they might work on there. In this paper, we describe the skills that PhD economists apply in tech companies, the companies that hire them, the types of problems that economists are currently working on, and the areas of academic research that have emerged in relation to these problems.

What Tech-Relevant Skills Do PhD Economists Have?

To draw inspiration from Liam Neeson's line in the movie *Taken*, economists have "a very particular set of skills." Here, we focus on three broad skillsets that are part of the economics curriculum that allow economists to thrive in tech companies: the ability to assess and interpret empirical relationships and work with data; the ability to understand and design markets and incentives, taking into account the information environment and strategic interactions; and the ability to understand industry structure and equilibrium behavior by firms.

Assessing Empirical Relationships

Relative to other disciplines, economists have several strengths in thinking about data. First, economists are interested in understanding which relationships are *causal*—and which are not. Over the past 30 years, economics has developed a toolkit to identify causal relationships in real-world data. As the internet age has helped to usher in an era of unprecedented amounts of data, this has also contributed to the growing demand for economists.

For example, empirical applied microeconomics has developed tools for using "natural experiments" and for evaluating policies—tools such as instrumental variables, causal panel data models, and regression discontinuity (for a review of some of these, see Angrist and Pischke 2009). As we describe further in the next section, these tools are widely used in technology firms to answer questions about the effects of interventions such as price changes, the introduction of new products, changes to the user interface, and advertising effectiveness. Economists' attention to identifying causal effects, as well as to both the statistical and economic significance of findings, are important contributions to the practice of empirical analysis in tech firms. Industrial organization economists and market design economists have also developed methods for estimating the impact of counterfactual price changes or changes to market design. Perhaps surprisingly, these tools are less widely used in tech firms than the tools of empirical applied microeconomics, although there are notable exceptions.

Experiments are central to the decision-making process within the tech sector. Most large tech companies evaluate product changes through "A/B testing," or

randomized controlled trials, conducting thousands or tens of thousands of A/B tests per year. Experiments pose important managerial and technical questions, ranging from how to choose an appropriate sample, to how to design the intervention itself, to how to move from experimental results to a managerial decision.

With many experiments seeking to identify small effects over a massive number of users, changes to the methodology of A/B testing can be impactful. The science of experimental design has therefore become an important topic within tech companies, often pushing the research frontier. For example, Blake and Coey (2014) highlight challenges in running experiments in marketplaces, where equilibrium effects create interference between treatment and control groups in a paper motivated by challenges they faced at eBay and Facebook. Athey, Eckles, and Imbens (2018) examine issues that arise in evaluating experiments in a network setting, in a paper motivated by challenges they faced at Amazon and Facebook.

The widespread use of experiments in the tech sector has at times proved controversial, as when Facebook ran an experiment to test how users would react when Facebook varied whether users were shown more positive or negative posts in their newsfeeds (Kramer, Guillory, and Hancock 2014). Although the experiment ultimately found very small effects, it generated considerable public backlash against Facebook (Meyer 2014) and an expression of editorial concern from the journal that published the experiment (Verma 2014). In response to public pressure and broader concerns about the ethics of experimentation within companies, Facebook updated its internal procedure for deciding which experiments to run. Companies and policymakers are still exploring ways to establish best practices that allow for productive experimentation and uses of data, while protecting the privacy and safety of participants.

The widespread use of machine learning in tech firms has also created new challenges and opportunities. Initially, academic economists were slow to take up machine learning for reasons ranging from the lack of asymptotic results behind many approaches to machine learning to questions about whether prediction problems are important from an economics perspective. Thus, some economists came to tech firms unfamiliar with machine learning, requiring them to learn a new set of methods in order to communicate with the machine learning community. More recently, the interaction of economists with technology firms has contributed to an expansion of interest among economists in machine learning—focusing both on prediction problems and causal inference problems.

Motivated by the need to bring causal inference techniques to the large datasets of technology firms as well as the desire to make full use of these rich datasets, a recent literature has developed combining machine learning and causal inference (Athey, forthcoming), and this literature in turn has influenced the business practice of technology firms: for example, Hitsch and Misra (2018) apply Wager and Athey's (forthcoming) causal forest method in an application to targeted promotions, while the Athey and Imbens (2016) approach to recursive partitioning for causal effects has been applied in technology firms' A/B testing platforms. From a practical perspective, the intersection of machine learning and economics allows economists to understand what works, what doesn't, and why.

While experiments have played an important role within tech companies, they also have limitations. Economists have helped to bring a broader causal inference toolkit to supplement experiments within tech, using methods such as instrumental variables, causal panel data models, and regression discontinuity. This has allowed companies to obtain treatment effects in contexts where experiments might be difficult or costly to run.

In addition to their focus on causal relationships, economists are interested in understanding the tradeoffs involved in different outcome metrics. In many technology firms, decisions about product design, marketing, and even human resources are determined by empirical analysis (rather than subjective evaluation), and the choice of metrics will guide incentives throughout the companies. Economists have sought to understand the relationship between short-term metrics such as clicks on an advertisement (also called “surrogates,” as in Athey, Chetty, Imbens, and Kang 2016) which are easy to observe, and longer-term metrics (like revenue or the lifetime value of a customer), which are more difficult to observe, but better represent company goals.

For example, a large technology company made the following change in measurement for email marketing. The old measure, customer sales, was noisy. Consumers might take weeks before making a purchase. The new measure, opening the email, was immediately observable, and could be incorporated to adjust the content of the email very quickly. However the company found that within months, the number of sales *per email* declined precipitously, because the marketing emails evolved to maximize email opening rates without regard to final sales. For example, the successful emails (using the opening rate metric) had catchy subject lines and somewhat misleading promises. For economists, it is natural to think about a metric not only as a statistical measure, but also as implicitly governing worker incentives, and to suggest ways to provide incentives for long-term innovation as well as short-term metrics that better capture long-term effects. More broadly, economists are interested in the difference between short-term and long-term objectives, which can often lead to dramatically different conclusions in making product and market design choices, and in developing algorithms. Economists have focused on the link between experiments, algorithms, and managerial decisions.

Finally, the theoretical and empirical training of economists prepares them to think about both intended and unintended consequences of different decisions. For example, Airbnb originally made it very easy for landlords to reject guests after seeing their name and picture. While this extent of flexibility may have led to short-term user growth (the metric that Airbnb had been focusing on), an experiment run by Edelman, Luca, and Svirsky (2017) showed that it also led to widespread racial discrimination against African-Americans on the platform. Thus, Airbnb’s market design choices raised the possibility of reintroducing discrimination to a market that had worked hard to reduce it. Fisman and Luca (2016) proposed a series of market design choices that might reduce discrimination in online markets more generally—such as further automating transactions on platforms. As a result of this work, the company created a task force that weighed the different options, which

led to a full-time team of data scientists to explore discrimination on an ongoing basis. Regulators also became involved, prompting Airbnb to continue these efforts. Ultimately, Airbnb implemented a variety of changes that balanced their desire for short-term growth against the company's goal of reducing discrimination on the platform, objectives which were not always aligned.

Designing Markets and Incentives

The rise of economists in tech companies has coincided with the rise of market design, a field that was pioneered by Stanford economist Bob Wilson and extended into a variety of application areas by economists such as Paul Milgrom and Al Roth (who won the Nobel Prize for his pioneering work in this field). Market design has shifted economists away from using a primarily descriptive lens to a more prescriptive one, using the tools of economics to engineer better-functioning markets. These economists—and Roth in particular—have promoted the idea of the “economist as engineer,” whereby the economist becomes deeply involved in the implementation of economic ideas and tailors recommendations to the fine details of the problem. While market design research initially focused on offline marketplaces such as spectrum auctions, residency matching programs, and kidney exchange, economists have more recently taken the market design mindset into the tech sector. For example, the lens that Roth has long used in offline markets—exploring issues around market thickness, congestion, and safety of participants—has gained further prominence in online marketplaces, where design choices are front and center.

Applications of market design in tech firms range from Google, Yahoo!, and Microsoft's marketplaces for selling advertisements (Varian 2007; Edelman, Ostrovsky, and Schwarz 2007; Athey and Ellison 2011; Agarwal, Athey, and Yang 2009; Athey and Nekipelov 2013) to Uber's market for rides (Cohen, Hahn, Hall, Levitt, and Metcalfe 2016). Much of this literature has examined pricing and allocation mechanisms, as well as reputation systems. Other work has focused on search costs (Athey and Ellison 2011; Fradkin 2017; Cullen and Farronato, 2018). Multisided platforms are especially ripe for an economist's skills, since these are exactly the kinds of settings in which it is critical to think through strategic behavior, interactions, and equilibrium effects.

Bringing together their unique perspectives on assessing empirical relationships with their expertise in market design, economists offer particular value to technology firms by bringing together theory and data to predict not just the immediate effect of a decision, but how a decision affects equilibrium behavior in a market.

Analyzing Equilibrium Market Structure

Tech companies think a lot about which markets to enter, taking into account the current and potential competitive landscape. For example, questions about market structure have arisen in the battle between Uber and Lyft, and helped to shape expansion and acquisition strategies. Economic theory, including the theory of platforms and market design, speaks to the forces that might lead a market to

be highly competitive, as well as the forces that make monopoly more likely. This is helpful for platforms deciding a strategy about which markets to enter, and also for policymakers and regulators. Currently, the question of market power is hotly debated in the technology industry, and economists can help by putting structure on the debate, even if they cannot perfectly predict the future.

Applications of Economics in Technology Firms

Economists now work on a variety of issues pertaining to tech companies. In this section, we highlight several exemplars of economics in tech companies: designing advertising auctions, estimating the returns to advertising, designing review and reputation systems, and studying the effects of reviews on firms.

Design of Online Advertising

Advertising has changed dramatically with the advent of online technology, and with the involvement of economists. This involvement has been concentrated in two areas: the design of advertising auctions and estimating the returns to advertising.

The involvement of economists in online advertising auctions dates back to the late 1990s, when Simon Wilkie, an economics professor at Cal Tech, started advising GoTo, a company that later became Overture and eventually powered Yahoo!'s search advertising auctions. In 2002, Hal Varian received a call from Eric Schmidt, the chairman of a young company called Google. Schmidt was intrigued by *Information Rules*, a book Varian had coauthored with Carl Shapiro, his fellow economist and colleague at the University of California, Berkeley. After speaking with Schmidt, Varian became a consultant for Google, and ultimately, the company's chief economist, the first academic microeconomist to become chief economist of a major technology firm. Preston McAfee, another market design economist, joined Yahoo! Research from Cal Tech a few years later, while Susan became consulting chief economist at Microsoft while on leave from Harvard in 2008. Susan and Preston McAfee also initially focused on market design and strategy questions surrounding online advertising.

To understand some of the issues involved with search advertising, consider the way it works. Search engines, ranging from general engines like Google and Bing to more specialized search engines like Yelp, generally sell advertising through auctions for specific terms. Bids are expressed in terms of a willingness to pay per outcome, such as a click, and advertisers with higher bids are rewarded with more favorable ad placement. Thus, firms must make choices about auction formats and parameters.

One complexity arises because in a traditional second-price auction with a single winner, the winner pays the bid of the second-highest bidder, which in turn means that the best strategy is to bid one's true value (and not to worry about being an outlying high bidder). However, in ad sales, the result is a ranking of bidders, not a single winner. The auction used by Google, Microsoft, and Yahoo! is a generalized

second-price auction, where each advertiser pays the price bid by the next lowest bidder. Work by Edelman, Ostrovsky, and Schwarz (of Yahoo! Research) (2007) shows that the generalized second-price auctions do not have the same properties as a second-price auction with a single winner (for instance there is not a dominant strategy equilibrium), but they remain useful in search engine advertising applications.

Athey and Ellison (2011) incorporate rational consumer search into the market design of auctions, motivating the use of reserve price not only as an instrument for raising revenue, but also as a tool for managing advertising quality and thus increasing users' incentive to search. One of us, Susan, used this as a framework for advising Microsoft to improve the ad quality on Microsoft's search engine. Later she took the theoretical models to the data and built an econometric model (Athey and Nekipelov 2012) that could be used to infer advertiser valuations and profits from their bidding behavior. This type of model can be used to understand how changes in algorithms affect advertiser well-being and thus forecast the future engagement of advertisers on the platform.

At Yahoo!, Ostrovsky and Schwarz (2016) observed that the reserve prices the company was setting were lower than what auction theory predicted would be revenue-maximizing for the seller. The pair assigned search keywords to a treatment and a control group. Keywords in the treatment group received a theoretically optimal reserve price calculated by the authors, while keywords in the control group used a default reserve price of \$0.10 per click. The treatment group increased ad revenue by several percentage points, leading Yahoo! to change its reserve price policies for all of its search advertising—and making the company millions of dollars in additional revenue.

Tech firms have also hired economists to solve challenges relating to the choice of outcome of advertising, such as pay-per-click versus alternatives. Agarwal, Athey, and Yang (2009) explore the benefits and drawbacks of pay-per-click compared to pay-per-action, in which advertisers only pay each time an individual performs an action after clicking the ad link—such as buying a product. Pay-per-action also allows advertisers to better express their benefits from different ad placements; some ad placements may attract consumers who are likely to purchase expensive items, but other ad placements may attract consumers who purchase cheap items, or do not purchase at all. Pay-per-action allows the advertising platform to optimize on behalf of the advertiser, obtaining more placements in scenarios where more profitable consumer behavior is likely. However, if the pay-per-action system allows bidders to bid on several different types of actions, the bidder may have opportunities to game the system, potentially making the revenue to the search engine arbitrarily low.

Finally, although not much academic work has analyzed Facebook's online advertising auctions, Facebook's early decision to adopt a Vickrey auction to sell its advertising space was heavily influenced by the training of a Facebook employee, John Hegeman, in the graduate program at the economics department at Stanford, which has considerable expertise in auctions (Amit, Grefl, and Hegeman 2013).

The Role of Ranking and Incentives in Marketplaces

Equilibrium effects can be especially challenging to understand in the platforms and marketplaces that are common in the tech industry. For example, a change to the user interface at eBay that made it easier for consumers to find the products they want, and thus to do price comparisons, affected consumer choice behavior (Dinerstein, Einav, Levin, and Sundaresan 2018), but that in turn can affect the prices charged by sellers. Over the long term, changes in pricing behavior by sellers affect consumers' desire to shop on eBay at all, which in turn influences seller behavior. Similar issues arise in many marketplaces. In general, the way marketplaces and intermediaries rank the offers from sellers or service-providers can be thought of as an incentive system. Marketplaces like Airbnb incentivize owners to maintain their calendars accurately and accept booking requests from travelers by prioritizing owners who behave as desired, and demoting those that do not. Economists are well positioned to analyze issues that arise in ranking offers from sellers, not just on short-term user behavior, but also the equilibrium impact on the marketplace as a whole.

Estimating the Returns to Advertising

Estimating the returns to advertising has traditionally been difficult. Older media, such as print and television, do not allow for showing different advertisements or tracking behavior at the individual consumer level, which makes designing randomized experiments difficult. Nonrandomized observational studies are biased due to selection issues. Thus, most traditional studies of advertising were plagued by poor identification strategies, limited data on outcomes, and small sample sizes.

The digital age has allowed a better understanding of the returns to advertising. Platforms such as Facebook, Google, and Microsoft collect vast amounts of data on user behavior, and regularly run experiments to test the effectiveness of their online advertising systems—allowing them to make progress on understanding the conditions under which advertising is most effective. Economists at such firms can thus draw on existing theories of market design, generate new ideas, and rapidly test and evaluate those ideas.

Economists at companies that advertise online have also made significant progress in understanding the effectiveness of advertising. For example, while working at eBay Research Labs, Blake, Nosko, and Tadelis (2015) conducted field experiments that allowed them to understand the impact of eBay's advertising campaigns on Google and Bing. They found that search engine marketing—purchasing ads to be displayed on search engines when certain search terms are entered—was only effective when ads are viewed by new or infrequent eBay customers and when the search terms already contain the firm being searched for. Since frequent customers drive most of their sales, the overall returns were negative, a significant result given that eBay's yearly US search engine marketing budget was over \$50 million at the time of the experiment.

In other contexts, advertising appears to be a positive investment. Johnson, Lewis, and Reiley (2016) report a 3.6 percent increase in sales among consumers

shown advertisements for a large retailer on Yahoo!, with a point estimate, though not statistically significant, of positive returns. Their experiment used a sample size in the millions, a control group shown an irrelevant ad (in addition to a group shown no ads), and a large set of individual covariates. Dai, Kim, and Luca (2016) collaborated with Yelp to display ads randomly for a set of previously non-advertising restaurants—a design that allowed them to include many small businesses rather than a small number of well-known businesses. Restaurants for which ads were shown had 25 percent more page views and 5 percent more reviews (which can be viewed as a proxy for actual visits to the restaurant)—and a back-of-the-envelope calculation suggests a positive return on investment.

Economists have also designed long-term experiments that examine the impact of ads on the propensity of users to buy or use the advertised product; Huang, Reilly, and Riabov (2018) study Pandora Internet Radio consumers over a 21-month period, estimating a fairly linear negative relationship between the quantity of ads shown to each consumer while listening to the internet radio and usage of the service, and further showing that increasing the ad load increased purchases of paid, ad-free subscriptions.

But it remains challenging to measure the returns to advertising. Lewis and Rao (2015), two economists formerly at Yahoo!, discuss the challenges in a meta-analysis spanning 25 online advertising field experiments. They argue that even studies of returns to advertising that can use online data are still held back by the signal-to-noise ratio in individual sales data, where standard deviations are often an order of magnitude higher than means. Even studies with hundreds of thousands of users often produce confidence intervals too wide even to distinguish highly profitable ads from wholly ineffective ones.

Designing Review and Reputation Systems

Online reviews and reputation systems have become increasingly prevalent in the past decade. Platforms like Yelp and TripAdvisor contain hundreds of millions of reviews for businesses ranging from plumbers to hotels. Uber, Airbnb, and other online marketplaces also rely heavily on reputation systems to facilitate trust between strangers, and traditional retailers ranging from Home Depot to Gap have developed review systems of their own.

Economists have been involved in the design of reputation systems—focusing on understanding the systematic biases that can occur in review ecosystems and the design choices that might mitigate these biases. One bias that has been documented in review systems in online marketplaces arises from reciprocal reviewing, in marketplaces where buyers and sellers review each other. While reciprocal reviewing can be a valuable way to build trust on both sides of the market, it has the potential to create incentives for upward-biased reporting. When Airbnb allowed the reviews of renters to be posted before those of the hosts, guests might have been hesitant to leave bad reviews out of concern that hosts would reciprocate. Bolton, Greiner, and Ockenfels (2013) propose a fix to this dilemma in the context of eBay, which offered reciprocal reviewing where both buyer and seller

reviews were immediately posted. The solution eBay (and Airbnb) explored is to postpone displaying reviews until both sides have left a review, or until a certain amount of time has expired. Under this system, however, buyers may still be reluctant to provide negative feedback if they suspect that it would discourage future sellers from transacting with them. Therefore, eBay added an anonymous, one-way review component called a “detailed seller rating,” where buyers assign sellers several numerical scores and the results are only viewable in aggregate form. Fradkin, Grewal, and Holtz (2018) study this issue using a randomized experiment at Airbnb (working within the company), and find results consistent with the hypothesis that reducing the possibility of retribution increases the informativeness of reviews.

A second bias can arise because reviews in online marketplaces are voluntary and so may suffer from selection bias. In particular, users may be more likely to leave a review after an especially positive or negative experience. For example, a group from eBay’s team of economists found evidence that eBay users were more likely to leave a review after a positive experience, relative to a negative one (Masterov, Meyer, and Tadelis 2015). Review platforms have a variety of tools to tackle the selection problem, such as sending emails to encourage consumers to leave reviews and even paying reviewers. Alternatively, platforms can incorporate information about buyer and seller review frequency into reputation scores—for example, penalizing sellers who receive low rates of feedback. Upon the recommendation of an in-house economist, a large online labor market developed a system that allowed for both private and public feedback, finding that private feedback was less inflated than public-facing reviews.

A third bias in online reviews occurs when businesses, or individuals hired by businesses, surreptitiously leave reviews about themselves or their competitors. Luca and Zervas (2016) explore the role of economic incentives in a business’s decision to commit review fraud, finding that independent restaurants and restaurants with a declining reputation are more likely to commit review fraud, and restaurants with high competition are more often targeted with fake negative reviews. One mechanism for reducing fraudulent reviews is to verify whether a transaction has occurred before allowing a review, as is policy on Airbnb, for example; other sites, such as Amazon, label reviews that come from a verified purchase. While this precaution may reduce fake reviews, it may also prevent legitimate reviews on some platforms by increasing the barriers to contributing content. Mayzlin, Dover, and Chevalier (2014) find evidence of promotional reviews in the context of TripAdvisor (which does not verify that reviewers have stayed at a property) and Expedia (which does). They find that relative to chains, independent hotels tend to have more five-star reviews on TripAdvisor (relative to Expedia). Moreover, the competitors of independent hotels tend to have more one-star reviews on TripAdvisor relative to Expedia.

In addition to creating incentives for people to leave high-quality reviews, platforms face a problem of how to aggregate reviews once the reviews are in place (Dai, Jin, Lee, and Luca 2018). In practice, review platforms such as Yelp and TripAdvisor

use algorithms to identify and remove content that is thought to be fake or of low quality. Platforms can also adjust and weight ratings to account for the informational content of each review, increasing the overall informational content of average ratings being presented to users. In practice, platforms also have to consider the incentive effects that different approaches to filtering and aggregating content might have.

Another perspective on reviews that is natural from an economist's training is to consider the cost of a user's time in writing a review as balanced against the value of information from a review. For example, Uber makes a decision about whether to require all riders to leave a review, or whether to request reviews only in some situations. It may not be worthwhile to request a review from every user who interacts with a highly experienced and well-rated seller on the marketplace. On the other hand, it is important to continue to collect some reviews to provide continued incentives for the seller to provide quality. In addition, there may be aspects of the user experience that can be directly measured by a marketplace (for example, time it took for the seller to ship, whether an Uber rider exceeded the speed limit, or how much a rider tips the driver). In such cases, it may be more efficient to ask the buyer to review aspects of the service that are more difficult to observe or infer directly.

The Effects of Reviews

The effects of online reviews on demand for products can be hard to identify. For example, hotels with higher TripAdvisor ratings may have higher demand either because ratings drive demand or simply because better hotels have higher ratings. However, economists have used a variety of methods to identify the causal impact of online reviews.

As one example, consider a book that is sold both on Amazon and on the Barnes & Noble website. The book would almost certainly have different ratings on the two platforms. Moreover, if an Amazon user left a review, the rating would change on Amazon, but not on Barnes & Noble, leading to variation in ratings across platforms and over time. Arguing that the exact timing of incoming reviews is plausibly exogenous, Chevalier and Mayzlin (2006) use this variation to estimate the impact of reviews on online book purchases. Specifically, they look for increases in sales on Amazon (relative to Barnes & Noble) after a review was left on Amazon (but not on Barnes & Noble)—implementing a difference-in-differences strategy. Using a regression discontinuity approach, Luca (2016) finds that higher ratings lead to higher sales for independent restaurants, but finds no evidence of this for chains. Anderson and Magruder (2012) find similar effects of Yelp ratings on restaurant reservations. Ghose, Ipeiritos, and Li (2012) uses a similar approach to understand the impact of TripAdvisor reviews. Beyond the average rating, other aspects of reviews are potentially important. For example, Sun (2012) explores the impact of the variance of product reviews, and highlights that if the variation in reviews of a product is driven by heterogeneous preferences, then, holding fixed the average review of a seller, it may be better to match some customers with sellers who have more variable reviews—as the variation may reflect the fact that the product is a good match for some customers but not for others.

Consumer reviews also have important implications for market structure and consumer welfare. Clemons, Gao, and Hitt (2006) argue that information provided in reviews can help to grow demand for products that are more differentiated by increasing the quality of the match, and find evidence generally consistent with this argument when looking at reviews for beer and growth in demand. Bar-Isaac, Caruana, and Cuñat (2012) theoretically show that introducing new information into a market can lead to a higher degree of product differentiation in markets. This finding suggests that the existence of online reviews may lead to a greater variety of products and services. Lewis and Zervas (2018) estimate the welfare effects of TripAdvisor reviews, focusing on the reduced search costs in markets with more review content.

Acquisitions, Exclusive Deals, and Strategy

The first question Susan was asked at Microsoft was whether internet search with search advertising was an industry that could sustain two or three players, or whether it was destined to be a monopoly. Her analysis of scale economies and indirect network effects in search suggested that sufficient scale was necessary for a second search engine to be viable; this analysis was used to value Microsoft's acquisition of Yahoo!'s search business, as well as other large business deals involving search. Later, the question arose of whether the smartphone market could sustain three platforms, something that has proved difficult to achieve. Questions about vertical integration also arise in these markets; for example, Google acquired the ITA travel search engine in 2010. Prior to that, ITA was providing the technology powering the travel search results for Microsoft's competing search engine, setting the stage for Google to increase dramatically its share of travel searches. This acquisition was closely reviewed by the US Department of Justice and was eventually approved with certain conditions (Miller 2011). Later, the European Commission imposed large fines on Google for biasing search results in favor of its own vertically integrated specialized search services (Scott 2017), and later for tying its search engine and mapping services to the applications store for Android (Satarino and Nicas 2018). Banks around the world have complained that Apple gives the Apple Wallet exclusive access to the NFC radio, a crucial component of mobile payments, in the iPhone. Apple then takes a fee for every credit card transaction that takes place on the Apple Wallet, a fee that is large (up to 0.15 percent) relative to the profits of the credit card networks (Zhu, Athey, and Lane 2018). Banks faced difficult strategic questions about whether to enable Apple Wallet in light of these fees as well as the control they would give up to Apple. Tech economists have been involved in analyzing all of these issues from both a business and regulatory perspective.

Economic theory and empirical approaches can also be critical in analyzing exclusive deals in the tech industry. For example, when gaming platforms such as Microsoft's Xbox and Sony's Playstation release new generations, they typically sign exclusive deals for games. Economic theory and empirical methods (for example, Lee 2013) can be used to value these exclusive deals, incorporating the direct impact

of those games on the sales of consoles at the time of launch, but also the indirect effect of those additional consoles on the subsequent incentives of game developers to develop for a platform, which in turn affects consumers, and so on.

Positions for Economists at Tech Companies

Economists have had mixed reception in tech companies. While some companies like Amazon have been quick to bring economists into the highest levels of decision-making, others have been slower, with economists sitting within data science teams or policy teams with less influence over the direction of the firms. In practice, economists within technology companies take on a number of roles ranging from Chief Economist to Product Manager. Economists often work within inward-facing groups at companies, including forecasting and planning, pricing, testing, and data science teams as well as outward-facing groups including policy, public relations, and marketing teams. We outline some examples of these roles.

Data science/analytics is one of the fastest-growing job areas as tech companies become more data-driven. Economists use observational and experimental data to answer business questions, such as whether to introduce new products, how to understand the effectiveness of large initiatives, and how to evaluate the impact of competitors. Because this work directly informs the decisions of many other departments, some firms have embedded data scientists in product teams while others have centralized data science teams. For example, Amazon currently embeds data scientists within product teams, while Yelp has a centralized data science team. Economists often help to manage teams of data scientists as well—for example at Coursera, or for a period of time at HomeAway.

Tech companies are increasingly using *experimentation or A/B testing* to answer product or platform design questions, such as the launch of a new product or advertising campaign. Economists can help to manage the design, process, and analytics around randomized controlled experiments. Some firms have embedded A/B testing specialists within their functional teams (for example, in marketing teams) while others have a separate team to manage a larger testing platform. For example, Uber and Facebook have economists involved in managing experimentation platforms and process in a context with strong network effects and many experiments. Other economists have developed and applied techniques for estimating heterogeneous treatment effects in A/B testing platforms (for example, Athey and Imbens 2016; Wager and Athey forthcoming).

Some tech firms have embedded experimentation or data scientists into their *advertising/marketing analytics*. These teams typically evaluate the effectiveness of advertising, design experiments around advertising, optimize advertising spending, and predict the success of advertising campaigns. For example, Netflix has a team working on these issues.

Economists working as *product managers* can also design experiments and surveys that answer questions that guide product designs and other strategic decisions,

including ranking algorithms in search platforms or presentation of information in stores. These tasks often involve drawing causal inferences from observational data—for example, using difference-in-differences methods to evaluate the impact of a new product or feature.

In *regulation/litigation* settings, the role of an economist includes writing policy white papers that translate theory and empirical work for a legal or policy audience, contributing knowledge of specific subject areas such as telecommunications policy, intellectual property, and antitrust from an economic perspective. Chief economists often spend a share of time on these issues as well. Airbnb has economists trying to understand housing markets and policy. Uber has economists investigating the impact of Uber on the taxi industry and quality of rides. Google (and previously, Yahoo! and Microsoft) has had economists studying antitrust issues related to Google's dominant position in the search industry.

Tech companies also have economists in a *public policy role*, helping to partner with policymakers, often through data-sharing and analysis. For example, Yelp partnered with the City of Boston to develop an algorithm that allowed the city to help target inspections for restaurant health code violations (Glaeser, Hillis, Kominers, and Luca 2016). Yelp data has been used to forecast government statistics (Glaeser, Kim, and Luca 2017), to understand how neighborhoods change during gentrification (Glaeser, Kim, and Luca 2018), and to estimate the impact of the minimum wage on restaurant exit and prices (Luca and Luca 2018). Yelp also partnered with cities (and a third party data provider) to display hygiene violations online, providing a modern, digital implementation of the hygiene disclosure policies analyzed by Jin and Leslie (2003), where regulation forced restaurants to prominently display their hygiene ratings in their stores. This initiative helped to steer customers away from restaurants with the most violations of health code policy (Dai and Luca 2018). Zillow, the online real estate company, creates reports of local housing markets. Search data from Yahoo! and Google has been used to help forecast economic activity (Goel, Hofman, Lahaie, Pennock, and Watts 2010; Choi and Varian 2012; Wu and Brynjolfsson 2015). LinkedIn is exploring the ways in which its data can help to shed light on labor markets. Uber's public policy team examines issues such as the impact of driving for Uber on driver welfare (Chen, Chevalier, Rossi, and Oehlsen 2017), the impact of Uber on labor markets and local economies (Hall, Horton, and Knoepfle 2017), and the role of gender in labor markets (Cook, Diamond, Hall, List, and Oyer 2018).

Several leading technology companies, including Zillow and Houzz, employ economists to do research designed for *public and media relations*, intended to inform potential customers and to create awareness for the company. For example, a primary mechanism for Zillow to attract consumers in its early years was that its chief economist created analyses of real estate markets to be covered by local and national news media. As another example, Houzz employs PhD economists who analyze and publish trends and data relevant to home remodeling.

Members of the *chief economist team* conduct and oversee many of the roles outlined above and also may make strategic decisions for the company. These decisions might include acquisitions and partnerships (one of us, Susan, worked on

strategy and empirical analysis for Microsoft's investment in Facebook, the acquisition of Yahoo!'s search business, and the company's strategy for cloud computing), as well as pricing and market entry.

Depending on the size of the company, economists have also gone into forecasting and planning teams (using time-series econometrics and modeling), pricing teams (using market design and supply-and-demand modeling), and academic relations (recruiting academics to fill the economic roles and build academic awareness surrounding policy and public relations issues).

Discussion

While we have focused mainly on economists working directly in tech firms, the rise of tech companies and emergence of the economics of digitization has important implications for academia as well. The shifting field leads not only to new research questions, but also to new academic positions, opportunities for collaborations, and potential career shifts. In this section, we address these opportunities.

Partnerships with Academics

While a growing number of economists now work within tech companies, collaborations with academics remain central to the strategy of tech firms and to the diffusion of economics within companies. For example, Airbnb, Amazon, eBay, Facebook, Indeed, LinkedIn, Microsoft, Rover, TaskRabbit, Uber, Upwork, Yelp, and Zillow have all collaborated with academic economists. These collaborations have several advantages for companies.

First, academics often have deep expertise in focused areas, including the key areas highlighted in this paper and many others; for example, a behavioral economist might shed light on the role of habit formation in user behavior. A market design economist might have unique insight into mechanisms driving market thickness. An econometrician might offer new ways to run experiments in a market with complicated network effects. Academics are also well-positioned to draw on insight from different contexts, since their work is less concentrated on a single platform.

Second, economists working full time within companies are often under pressure to deal with immediate issues (such as whether to change prices in a given quarter, or whether a specific advertising campaign was productive). Academics are insulated from these pressures, and so can explore longer-term strategic issues such as whether a company is even tracking the right metrics, or whether it makes sense to shift product composition.

Third, the hiring of economists by tech companies has brought forth a related risk—little research is being conducted internally on the shortcomings of tech companies and the negative implications of their models. For example, Airbnb did not examine racial discrimination on the platforms until academics documented it in academic research, thus bringing it to the attention of policymakers, Airbnb users, and ultimately, Airbnb managers. Working with academics and allowing a

broad degree of autonomy can help to get more credible and objective assessments of issues with which the companies are dealing.

At the same time, challenges do arise with academic partnerships. For example, academics often sign agreements with firms guaranteeing the ability to publish their results regardless of the result. In principle, this helps to reduce publication bias. However, firms may choose not to sign agreements around research topics where they are concerned about what the answers might be, potentially creating a bias towards papers favorable towards firms and creating an incomplete snapshot of an issue. This issue is not new, since economists have obtained data from firms and government agencies at their discretion for many years. However, as collaborations become more standard, this issue becomes more important.

Academic Jobs for Digitization Economists

The number of academic positions for digitization economists is growing. While some of these are in economics departments, digitization economists now also teach in business schools in strategy, marketing, information systems, entrepreneurship, and other departments. Doctoral students with interests in these areas should be aware that while recruiting for some of these positions takes place through the American Economic Association, recruiting for other positions, such as in marketing, operations, and information systems, often takes place on other timelines and outside of AEA mechanisms.

Tech companies have also created strong demand for undergraduate economics majors, who take roles ranging from product management to policy. Leading universities including Dartmouth, Harvard, MIT, Princeton, Stanford, and Yale teach about online platforms in their introductory microeconomics courses, or have created entire courses related to the “economics of digitization” (including courses on e-commerce, online platforms, and related areas). MIT’s economics and computer science departments have partnered to create a new major in computer science, economics, and data science. Harvard, MIT, and other universities have developed data science initiatives, drawing in computer scientists, economists, and other social scientists. We see opportunities to expand these course offerings, and to combine them with additional course material for students looking for a career at tech companies. Courses about marketplaces and platforms, taught from an economics perspective, have also proliferated among business schools, such as Boston University, Harvard, New York University, and Stanford.

While PhD economists are well suited to tech careers in many ways, we also see areas for the field to improve the preparation of PhD economists for working with or in tech companies. First, with the importance of prediction, targeting, and precise estimates in tech companies, machine learning plays an important role in tech companies. While the field of economics has long been a leader in causal inference, the field is still in the process of incorporating machine learning into its standard toolkit. Second, economists have historically received less training, relative to computer scientists, at coding and at optimizing code to run statistical algorithms at large scale. Investing in these skills (and incorporating them into the

PhD curriculum) can help to prepare economists to work in this area. At the same time, it remains important that economists have a strong conceptual understanding of economic issues like incentives and equilibrium effects, as well as strong empirical skills in the areas such as causal inference that we have described in this paper.

Shifts Between Academia and Practice

Economists in this area have growing opportunities to shift between academia and practice. Microsoft, Google, Yahoo!, Facebook, Amazon, eBay, Yelp, Uber, and other companies have all hosted faculty during sabbaticals. Tenured faculty members have left academia for positions at Amazon, Google, and elsewhere. Practitioners have also transitioned into academia—for example, leaving Facebook and Microsoft for MIT and Stanford. We believe this is the beginning of a larger movement in which a greater share of academic economists spend time in practice, acquiring a deeper understanding of what issues are most important for efficiency and profitability in technology firms, as well as getting exposure to unsolved business problems that may highlight fruitful academic research questions. As more PhD economists accept positions at tech companies, clearer paths for spending time (or re-entering) academia will likely appear, for those who are interested in this option. Firms that allow their economists to continue to publish will likely have an advantage in recruiting and retaining economists who want to retain flexibility in their career paths.

■ *We thank Duncan Gilchrist and Guido Imbens for valuable feedback. Stephanie Chan and Byron Perpetua provided excellent research assistance.*

References

- Agarwal, Nikhil, Susan Athey, and David Yang. 2009. "Skewed Bidding in Pay-per-Action Auctions for Online Advertising." *American Economic Review* 99(2): 441–47.
- Amit, Alon, Yaron Grelf, and John Hegeman. 2013. "Advertisement Selection and Pricing Using Discounts based on Placement." U.S. Patent 2013/0325585 A1 file June 4 2012, and published December 5, 2013. <https://patents.google.com/patent/US20130325585A1/en>.
- Anderson, Michael, and Jeremy Magruder. 2012. "Learning from the Crowd: Regression Discontinuity Estimates of the Effects of an Online Review Database." *Economic Journal* 122(563): 957–89.
- Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics*. Princeton University Press.
- Athey, Susan. Forthcoming. "The Impact of Machine Learning on Economics." Forthcoming in *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay K. Agrawal, Joshua Gans, and Avi Goldfarb. University of Chicago Press.
- Athey, Susan, Raj Chetty, Guido Imbens, and Hyunseung Kang. 2016. "Estimating Treatment Effects Using Multiple Surrogates: The Role of the Surrogate Score and the Surrogate Index." Unpublished paper.
- Athey, Susan, Dean Eckles, and Guido Imbens. 2018. "Exact p -values for Network Interference."

Journal of the American Statistical Association 113(521): 230–40.

Athey, Susan, and Glenn Ellison. 2011. “Position Auctions with Consumer Search.” *Quarterly Journal of Economics* 126(3): 1213–70.

Athey, Susan, and Guido Imbens. 2016. “Recursive Partitioning for Heterogeneous Causal Effects.” *PNAS* 113(27): 7353–60.

Athey, Susan, and Denis Nekipelov. 2012. “A Structural Model of Sponsored Search Advertising.” https://people.stanford.edu/athey/sites/default/files/structural_sponsored_search.pdf.

Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat. 2012. “Search, Design, and Market Structure.” *American Economic Review* 102(2): 1140–60.

Blake, Thomas, and Dominic Coey. 2014. “Why Marketplace Experimentation is Harder than It Seems: The Role of Test-Control Interference.” In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, 567–582.

Blake, Thomas, Chris Nosko, and Steven Tadelis. 2015. “Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment.” *Econometrica* 83(1): 155–74.

Bolton, Gary, Ben Greiner, and Axel Ockenfels. 2013. “Engineering Trust: Reciprocity in the Production of Reputation Information.” *Management Science* 59(2): 265–85.

Chen, M. Keith, Judith A. Chevalier, Peter E. Rossi, and Emily Oehlsen. 2017. “The Value of Flexible Work: Evidence from Uber Drivers.” NBER Working Paper 23296.

Chevalier, Judith A., and Dina Mayzlin. 2006. “The Effect of Word of Mouth on Sales: Online Book Reviews.” *Journal of Marketing Research* 43(3): 345–54.

Choi, Hyunyoung, and Hal Varian. 2012. “Predicting the Present with Google Trends.” *Economic Record* 88(1): 2–9.

Clemons, Eric K., Guodong Gao, and Lorin Hitt. 2006. “When Online Reviews Meet Hyperdifferentiation: A Study of the Craft Beer Industry.” *Journal of Management Information Systems* 23(2): 149–71.

Cohen, Peter, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe. 2016. “Using Big Data to Estimate Consumer Surplus: The Case of Uber.” NBER Working Paper 22627.

Cook, Cody, Rebecca Diamond, Jonathan Hall, John A. List, and Paul Oyer. 2018. “The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Rideshare Drivers.” NBER Working Paper 24732.

Cullen, Zoë, and Chiara Farronato. 2018. “Outsourcing Tasks Online: Matching Supply and Demand on Peer-to-Peer Internet Platforms.”

<https://www.hbs.edu/faculty/Pages/item.aspx?num=50051>.

Dai, Weijia, Ginger Jin, Jungmin Lee, and Michael Luca. 2018. “Aggregation of Consumer Ratings: An Application to Yelp.com.” *Quantitative Marketing and Economics* 16(3): 289–339.

Dai, Weijia (Daisy), Hyunjin Kim, and Michael Luca. 2016. “Effectiveness of Paid Search Advertising: Experimental Evidence.” Harvard Business School Working Paper 17-025.

Dai, Weijia (Daisy), and Michael Luca. 2018. “Digitizing Disclosure: The Case of Restaurant Hygiene Scores.” Unpublished paper.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2018. “Consumer Price Search and Platform Design in Internet Commerce.” *American Economic Review* 108(7): 1820–59.

Edelman, Benjamin, Michael Luca, and Dan Svirsky. 2017. “Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment.” *American Economic Journal: Applied Economics* 9(2): 1–22.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. “Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords.” *American Economic Review* 97(1): 242–59.

Fisman, Raymond, and Michael Luca. 2016. “Fixing Discrimination in Online Marketplaces.” *Harvard Business Review* 94(12): 88–95.

Fradkin, Andrey. 2018. “Search, Matching, and the Role of Digital Marketplace Design in Enabling Trade: Evidence from Airbnb.” Available at SSRN: <https://ssrn.com/abstract=2939084>.

Fradkin, Andrey, Elena Grewal, and David Holtz. 2017. “The Determinants of Online Review Informativeness: Evidence from Field Experiments on Airbnb.” Available at SSRN: <https://ssrn.com/abstract=2939064>.

Ghose, Anindya, Panagiotis G. Ipeirotis, and Beibei Li. 2012. “Designing Ranking Systems for Hotels on Travel Search Engines by Mining User-Generated and Crowdsourced Content.” *Marketing Science* 31(3): 493–520.

Glaeser, Edward, Andrew Hillis, Scott Duke Kominers, and Michael Luca. 2016. “Crowdsourcing City Government: Using Tournaments to Improve Inspection Accuracy.” *American Economic Review* 106(5): 114–18.

Glaeser, Edward L., Hyunjin Kim, and Michael Luca. 2017. “Nowcasting the Local Economy: Using Yelp Data to Measure Economic Activity.” NBER Working Paper 24010.

Glaeser, Edward L., Hyunjin Kim, and Michael Luca. 2018. “Nowcasting Gentrification: Using Yelp Data to Quantify Gentrification.” *American*

Economic Association: Papers and Proceedings 108: 77–82.

Goel, Sharad, Jake M. Hofman, Sébastien Lahaie, David M. Pennock, and Duncan J. Watts. 2010. “Predicting Consumer Behavior with Web Search.” *PNAS* 107(41): 17486–90.

Hall, Jonathan V., John J. Horton, and Daniel T. Knoepfle. 2018. “Pricing Efficiently in Designed Markets: Evidence from Ride-Sharing.” http://john-joseph-horton.com/papers/uber_price.pdf.

Hitsch, Günter J., and Sanjog Misra. 2018. “Heterogeneous Treatment Effects and Optimal Targeting Policy Evaluation. Available at SSRN: <https://ssrn.com/abstract=3111957>.

Hu, Nan, Jie Zhang, and Paul A. Pavlou. 2009. “Overcoming the J-shaped Distribution of Product Reviews.” *Communications of the ACM* 52(10): 144–47.

Huang, Jason, David Reilly, and Nikolai Riabov. 2018. “Measuring Consumer Sensitivity to Audio Advertising: A Field Experiment on Pandora Internet Radio.” Available at SSRN: <https://ssrn.com/abstract=3166676>.

Jin, Ginger, and Phillip Leslie. 2003. “The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards.” *Quarterly Journal of Economics* 118(2): 409–51.

Johnson, Garrett A., Randall A. Lewis, and Elmar I. Nubbemeyer. 2017. “Ghost Ads: Improving the Economics of Measuring Online Ad Effectiveness.” *Journal of Marketing Research* 54(6): 867–84.

Johnson, Garrett A., Randall A. Lewis, and David H. Reiley. 2016. “When Less is More: Data and Power in Advertising Experiments.” *Marketing Science* 36(1): 43–53.

Kramer, Adam D. I., Jamie E. Guillory, and Jeffrey T. Hancock. 2014. “Experimental Evidence of Massive-Scale Emotional Contagion through Social Networks.” *PNAS* 111(24): 8788–90.

Lee, Robin S. 2013. “Vertical Integration and Exclusivity in Platform and Two-Sided Markets.” *American Economic Review* 103(7): 2960–3000.

Lewis, Gregory, and Giorgos Zervas. 2018. “Welfare Impact of Consumer Reviews: A Case Study of the Hotel Industry.” Unpublished paper.

Lewis, Randall A., and Justin M. Rao. 2015. “The Unfavorable Economics of Measuring the Returns to Advertising.” *Quarterly Journal of Economics* 130(4): 1941–73.

Luca, Dara Lee, and Michael Luca. 2018. “Survival of the Fittest: The Impact of the Minimum Wage on Firm Exit.” <https://www.hbs.edu/faculty/Pages/item.aspx?num=52552>.

Luca, Michael. 2016. “Reviews, Reputation, and Revenue: The Case of Yelp.com.” https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1928601.

Luca, Michael, and Georgios Zervas. 2016. “Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud.” *Management Science* 62(12): 3412–27.

Masterov, Dimitriy V., Uwe F. Mayer, and Steven Tadelis. 2015. “Canary in the e-Commerce Coal Mine: Detecting and Predicting Poor Experiences using Buyer-to-Seller Messages.” In *EC ’15 Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pp. 81–93.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier. 2014. “Promotional Reviews: An Empirical Investigation of Online Review Manipulation.” *American Economic Review* 104(8): 2421–55.

Meyer, Robinson. 2014. “Everything We Know About Facebook’s Secret Mood Manipulation Experiment.” *The Atlantic*, June 28.

Miller, Claire Cain. 2011. “U.S. Clears Google Acquisition of Travel Software.” *New York Times*, April 8.

Ostrovsky, Michael, and Michael Schwarz. 2016. “Reserve Prices in Internet Advertising Auctions: A Field Experiment.” <http://web.stanford.edu/~ost/papers/rp.pdf>.

Satariano, Adam, and Jack Nicas. 2018. “E.U. Fines Google \$5.1 Billion in Android Antitrust Case.” *New York Times*, July 18, 2018.

Scott, Mark. 2017. “Google Fined Record \$2.7 Billion in E.U. Antitrust Ruling.” *New York Times*, June 27.

Shapiro Carl, and Hal R. Varian. 1999. “Information Rules: A Strategic Guide to the Network Economy.” *Harvard Business School Press*.

Sun, Monic. 2012. “How Does the Variance of Product Ratings Matter?” *Management Science* 58(4): 696–707.

Varian, Hal R. 2007. “Position Auctions.” *International Journal of Industrial Organization* 25(6): 1163–78.

Verma, Inder. 2014. “Editorial Expression of Concern: Experimental Evidence of Massivescale Emotional Contagion through Social Networks.” *PNAS* 111(29).

Wager, Stefan, and Susan Athey. Forthcoming. “Estimation and Inference of Heterogeneous Treatment Effects using Random Forests.” *Journal of the American Statistical Association* 109(508).

Wu, Lynn, and Erik Brynjolfsson. 2015. “The Future of Prediction: How Google Searches Fore-shadow Housing Prices and Sales.” In *Economic Analysis of the Digital Economy*, edited by Avi Goldfarb, Shane M. Greenstein, and Catherine E. Tucker, 89–118. University of Chicago Press.

Zhu, Feng, Susan Athey and David Lane. 2018. “Apple Pay and Mobile Payments in Australia (A).” *Harvard Business School Case no. 619-010*.