

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: Innovation Policy and the Economy, Volume 17

Volume Author/Editor: Shane Greenstein, Josh Lerner, and Scott Stern, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-48985-8 (cloth); 978-0-226-48999-5 (e-ISBN)

Volume URL: <http://www.nber.org/books/gree-2>

Conference Date: April 12, 2016

Publication Date: February 2017

Chapter Title: Designing Online Marketplaces: Trust and Reputation Mechanisms

Chapter Author(s): Michael Luca

Chapter URL: <http://www.nber.org/chapters/c13761>

Chapter pages in book: (p. 77–93)

Designing Online Marketplaces: Trust and Reputation Mechanisms

Michael Luca, *Harvard Business School*

Executive Summary

Online marketplaces have proliferated over the past decade, creating new markets where none existed. By reducing transaction costs, online marketplaces facilitate transactions that otherwise would not have occurred and enable easier entry of small sellers. One central challenge faced by designers of online marketplaces is how to build enough trust to facilitate transactions between strangers. This paper provides an economist's tool kit for designing online marketplaces, focusing on trust and reputation mechanisms.

I. Introduction

Over the past decade, there has been a proliferation of online marketplaces, ranging from eBay and Amazon to Uber and Airbnb. These online markets cannot exist without trust. Buyers need to trust that sellers are accurately describing products or services and fulfilling transactions as promised. Sellers need to trust that buyers will pay, and in the case of services and rentals, will abide by the agreed-upon terms of service. Trust between buyers and sellers is enabled by reputation systems and design choices made by online marketplaces.

This paper surveys the design choices and mechanisms that online marketplaces use to build trust and facilitate transactions. I focus on two central decisions faced by a marketplace. First, online marketplaces design review systems, which allow buyers and sellers to review each other and the product or service being transacted. Second, marketplaces choose what information buyers and sellers should have about each

other when deciding whether to transact, and how much flexibility market participants should have in choosing who they will transact with. I highlight challenges and trade-offs faced by online marketplaces in these design choices.

Historically, reviews have formed the backbone of reputation systems in the online marketplace. Sellers (and the products and services they are offering) on online platforms are rated and reviewed by buyers, and buyers can use this information to choose whom to interact with. Likewise, sellers on some platforms can review buyers. Reviews allow buyers and sellers to make sure they are transacting with someone deemed trustworthy enough to participate in the transaction. Reviews also create incentives for quality, as behavior of buyers and sellers is made more public. Despite the benefits of creating online reputation systems, several design challenges have been documented in the context of reviews. First, on platforms with reciprocal reviewing (i.e., where buyers and sellers review each other), users can have strategic incentives to manipulate reviews (Bolton, Greiner, and Ockenfels 2013; Fradkin et al. 2015). Second, reviews can suffer from selection bias (Hitt et al. 2008; Masterov, Meyer, and Tadelis 2015), as the people leaving reviews may differ from those who do not. Third, reviews may be distorted by promotional content in which businesses attempt to leave reviews for themselves (Mayzlin, Dover, and Chevalier 2014; Luca and Zervas 2015). Moreover, even if all reviews represent a user's true experience, some users may be more informative than others (Dai et al. 2014). Section III provides an overview of these issues, as well as potential design solutions for these challenges.

In first-generation platforms such as eBay, reviews—along with pictures and descriptions of the product—were the main information sources available to buyers at the time of purchase. Buyers and sellers typically did not see each other's name or picture—names were exchanged after the purchase was made. Over time, platforms have become considerably less anonymous. For example, Airbnb is an online platform for people to list and find short-term lodging; it attempts to build trust by allowing would-be renters to present personal profiles and post pictures of themselves. Hosts can then accept or reject guests based on their pictures and profiles. This design choice is now the norm on many platforms—indeed, it represents a broader trend of platforms' providing more information to market participants—not only about products and services, but about the people buying and selling them.

This also demonstrates the large degree of flexibility that online marketplaces have in their market-design choices (much more than is typically found in traditional markets).

The design choice of allowing hosts to reject guests on the basis of profile elements such as name, picture, and number of Facebook friends is a double-edged sword. On the one hand, it has the potential to make market participants feel more comfortable with each other, in part by reducing social distance. However, this design choice also facilitates discrimination—for example, African American guests (and hosts) are discriminated against on Airbnb (Edelman and Luca 2014; Edelman, Luca, and Svirsky 2016). This is part of a broader trend of discrimination that has been documented in online marketplaces including eBay (Nunley et al. 2011), Craigslist (Doleac and Stein 2013), and Prosper (Pope and Sydnor 2011). In labor market contexts, there is extensive evidence of discrimination against minority job applicants (Bertrand and Mullainathan 2004). Online labor markets such as Upwork are designed in a way that makes it easy for discrimination to manifest in online contexts as well. For example, Ghani, Kerr, and Stanton (2014) find that employers of Indian ethnicity are more likely to hire workers from India.

Overall, these findings highlight the promise and peril of online marketplaces. The existence of new online markets creates value for society, enabled by trust and reputation mechanisms. However, the same design choices allow for unintended consequences such as discrimination. In this paper, I document the evolution of new online markets, focusing on the design choices pertaining to trust and reputation. I highlight the key features of reputation systems and the problems that can arise. I then survey the evidence on discrimination in online markets and consider potential solutions that can be implemented by platforms.

II. The Rise of Online Marketplaces

The founding of Amazon and eBay in 1994 and 1995, respectively, ushered in the first generation of online marketplaces. These platforms were remarkable in their ability to facilitate transactions between strangers. Someone in upstate New York could order a used book from a stranger in Southern California and trust that it would arrive in a few days. These platforms facilitated transactions that would not otherwise occur, supported in part by reputation systems. At the same time, buyers and sellers generally did not provide pictures or names until after

a purchase was made. Both platforms had review systems to facilitate transactions.

In the 20 years since, a variety of more specialized platforms such as Airbnb, Uber, and Upwork have emerged, creating new markets and pushing a growing proportion of the economy onto the Internet. To give context for how quickly the industry is growing: Airbnb was founded in 2008 and is now valued at \$26 billion. It currently has more than two million listings, which is more than the largest hotel chain. Similarly remarkable is the statistic that there are now more Uber drivers-affiliated vehicles than traditional taxis in New York City. As of 2015, Uber is valued at roughly \$50 billion. Relative to first-generation platforms, these marketplaces have less anonymity—pictures and names are fast becoming the norm—but still have the reputation systems that were staples of earlier platforms. Table 1 provides a sample of marketplaces that are active as of 2016. Table 2 provides data on a subsample of those.

III. Reputation Systems

In this section, I highlight the main challenges—and potential solutions—to building a robust reputation in an online marketplace, focusing on the role of user reviews and drawing on Luca (forthcoming). Broadly speaking, there are two main mechanisms that a platform can use to improve the quality of its review system, taking into account the types of biases that commonly arise. First, it can improve the incentives to leave high-quality reviews. Second, taking the reviews as given,

Table 1
Sample Markets with Online Marketplaces

Market	Sample Platforms
Short-term housing	Airbnb, HomeAway, Couchsurfing
Labor	TaskRabbit, Upwork, MechanicalTurk
Pet sitting	DogVacay, Rover, Fetch!
Home services	HouseCall, Handy
Senior and child care	Care.com
Food delivery	Grubhub, Postmates, Eat24
Funding	Prosper, Lending Club, Kickstarter, Indiegogo
Dating	Tinder, OkCupid, CoffeeMeetsBagel, Match
Handmade goods	Etsy, CustomMade
Fashion	Renttherunway, Rocksbox, LeTote
Car rentals	Getaround, car2go
Car rides	Uber, Lyft, PickupPal, BlaBlaCar

Table 2
Sample Online Marketplaces

Marketplace	Year Founded	Valuation \$ Billions (as of 2015)	Revenue \$ Millions (as of 2015)	User Base Millions
Amazon	1994	250	107,010	304 (active customer accounts)
eBay	1995	69	6,103	162 (active buyers)
Craigslist	1995	3	381	60 (monthly users in United States)
Priceline	1997	62	9,220	13 (unique monthly users)
Etsy	2005	2	273	24 (active buyers)
Airbnb	2008	26	900	60 (cumulative guests)
Uber	2009	50	2,000	8 (cumulative riders)

it can aggregate the reviews in a way to increase their informational content (Dai et al. 2014). I then look beyond reviews for other tools to facilitate trust.

A. Reciprocal Reviewing

The process of reciprocal reviewing (i.e., of buyers and sellers reviewing each other) is valuable because it builds trust on both sides of the market. However, this can also create incentives for upward-biased reporting if reviewers fear retaliation. For example, when Airbnb's policy allowed renters' reviews to be posted before hosts' (as was previously the case), guests might have been hesitant to leave bad reviews out of concern that hosts would retaliate.

To circumvent this problem, platforms such as Airbnb have moved to a system sometimes referred to as "simultaneous reveal": they do not display reviews until both sides have left a review (or until the time to review has expired). Simultaneous revelation of reviews reduces the strategic problems associated with reciprocal reviewing (Bolton et al. 2013; Fradkin et al. 2015). However, even in a simultaneous-reveal system, strategic incentives persist. For example, buyers may be reluctant to provide negative feedback if they suspect that it would discourage other sellers from transacting with them.

One potential solution to the problems created by reciprocal reviewing would be to allow users to leave anonymous ratings (in situations where this is feasible), which could allow users to be more honest without fearing direct or indirect retaliation. To increase anonymity, platforms might consider showing only aggregated feedback. While this reduces the total amount of information being provided to users, it may increase the quality of that information. A second solution would be to have private ratings that are provided from a user to the platform, but not shared publicly. This would again reduce fear of retaliation on the part of users, while allowing the platform to use the information. The platform could use this private feedback in a variety of ways, such as updating its sorting algorithms, or using it to address concerns directly with the user being reviewed.

B. Self-Selection

Because reviews in online marketplaces are voluntary, they can suffer from selection bias, in that reviews are left by users who chose both to purchase the product or service and to leave a review online. In particular, users may be more likely to leave a review after an especially positive or negative experience. Indeed, Hu et al. (2009) find that reviews on Amazon tend to exhibit an asymmetric bimodal (J-shaped) distribution, with more positive than negative reviews. They argue that experiences for many products are more likely to resemble a normal distribution, and hence the J-shape suggests that people are more likely to leave reviews after extreme experiences. Masterov et al. (2015) find consistent evidence from eBay, where buyers are more likely to leave a review after a good experience.

There are a variety of tools that online marketplaces can use to improve this selection problem, ranging from sending repeated e-mail notifications encouraging people to report their experiences to paying people who leave reviews. Alternatively, to the extent that platforms know who is leaving a review, they can incorporate this information into the reputation score—for example, penalizing sellers who receive low rates of feedback.

C. Promotional Content

Another potential bias in online reviews occurs when people or businesses surreptitiously leave reviews about themselves or competitors.

Promotional content is driven, at least in part, due to changing economic incentives for a business (Mayzlin, Dover, and Chevalier 2014; Luca and Zervas 2015). It is more prevalent among independent businesses, when there is nearby competition, and when there is a negative shock to a business's reputation. One mechanism for reducing promotional reviews is to verify whether a transaction has occurred. While this may help to reduce fake or promotional reviews, it may also prevent legitimate reviews by increasing the barriers to contributing content.

Beyond verification of transactions, there are several other potential approaches to reducing promotional content: spam can be identified through algorithms that mine review text and characteristics (e.g., Ott et al. 2011; Akoglu et al. 2013). Platforms can also give more weight to reviewers who are less likely to be contributing promotional or fake content—for example, reviewers with longer transaction histories could receive more weight.

D. Social Distance

Social distance has the potential to affect a variety of other behaviors, such as generosity (Hoffman, McCabe, and Smith 1996) and reciprocity (Charness, Haruvy, and Sonsino 2007)—even in one-shot games. With pictures and profiles becoming an increasingly common design choice, online marketplaces are beginning to shrink the social distance between buyers and sellers. One might expect this to lead to higher ratings. This could be good if it makes users more comfortable with each other, but may distort reviews to the extent that higher ratings reflect a reluctance to leave negative feedback after a bad experience. The level of social distance can be a choice variable for platforms.

E. Beyond Reviews

Online marketplaces can supplement reviews through other trust-building mechanisms. The marketplace itself can do more to screen or authenticate information about buyers and sellers. For example, Airbnb could conduct interviews of renters and hosts, or they could run background checks. For example, Care.com, a marketplace for child-care providers, has incorporated a greater degree of platform-driven screening. Platforms can also provide insurance—and clear liability rules—for situations where something does go wrong.

IV. Discrimination in Online Marketplaces

Another area in which online platforms can fall short is in their potential to allow discrimination. In the early days of electronic commerce, economists hypothesized that online platforms might decrease the amount of discrimination in commercial transactions. For example, consider a customer looking to purchase a car. This is a market where prices are negotiated separately for each buyer, and the final price paid is opaque—and this is also a market where racial discrimination is prevalent (Ayres and Siegelman 1995). However, when a purchase is initiated through an online platform, Scott Morton et al. (2003) find no difference in outcomes on the basis of race. Similarly, eBay has less scope for discrimination because indicators of race and gender are generally not very salient on the platform. On other platforms, such as Amazon and Expedia, sellers essentially precommit to accepting all buyers regardless of race or ethnicity. Marketplaces have the potential to facilitate transactions while reducing discrimination.

Over time, though, the design of online platforms has changed, moving toward systems with less anonymity and fuller user profiles, coupled with more flexibility on the part of sellers and buyers to do business based on these attributes. For example, Expedia effectively prevents a hotel from rejecting a guest based on perceived race, ethnicity, or almost any other factor. But if the same hotel lists a room on Airbnb, it could reject a guest based on these or other factors. This highlights the fact that, while the Internet has the potential to reduce discrimination, this benefit depends on the design choices made by platforms.

In part because online platforms have evolved in this way, the conditions that made online markets potential havens free from discrimination are not prevalent on all platforms. To see the contrast between the prominent norms of online and offline markets, suppose that a senior executive at a hotel chain were to propose the following change to its reservation policy:

Let's start encouraging guests to upload their pictures, and let's allow branch managers to reject guests if they don't like the way they look. Potential guests would continue to make reservations through the website as before, but they would be nudged to upload pictures and links to Facebook pages or LinkedIn accounts. Branch managers would then have 24 hours to decide whether the guest looked sufficiently trustworthy. If the guest looked trustworthy, he or she would be accepted for a reservation. If not, the guest would be rejected. Managers would like the policy because they could accept the guests that they are most comfortable with. We'd also save money because the most costly guests would go elsewhere.

Of course, this idea sounds unrealistic. While it is conceivable that a branch manager could distinguish some of the trustworthy guests from the bad on the basis of looks or name, the adoption of screening on this basis is fraught with risk. In particular, what sorts of pictures and names should a manager screen for? The Civil Rights Act of 1964 prohibits discrimination on the basis of race, color, religion, or national origin. While the manager could certainly look for subtler cues, incorporating appearance in the reservations process would likely be too risky to justify any benefit the hotel might receive. Yet this design choice is the emerging norm in online marketplaces for short-term housing—and in many other marketplaces.

Discrimination has now been documented in a variety of online marketplaces. Table 3 summarizes the evidence on discrimination across different online marketplaces. This section documents the evidence, as well as design features that might facilitate discrimination.

A. *Housing Rental Markets*

Most searches for accommodations—both short term and long term—now begin online. For long-term apartment leases, platforms including Craigslist, Zillow, and many of their competitors provide rental information. In these markets, the platform typically provides the initial connection between potential landlords and lessees (and often charges a fee for advertising), but does not facilitate actual payment or charge fees based on transacted leases.

In short-term rentals on marketplaces such as Airbnb and Home-Away, the market looks very different. Transactions are often agreed upon from afar, and the platform does facilitate payment. On Airbnb, prospective hosts and guests provide information not just about their listings and preferences, but also about themselves—often posting photos, providing their name, and so forth. All of these factors are provided before a host decides whether to accept or reject a guest.

Table 3
Discrimination in Online Marketplaces

Airbnb	African American guests 15% less likely to be accepted.
eBay	Bids 20% lower for identifiably African American sellers.
Craigslist	Buyers less likely to contact identifiably African American sellers.
Upwork	Employers of Indian ethnicity more likely to hire workers from India.
Prosper	Loan listings for black borrowers are 25–35% less likely to get funded than loan listings for white borrowers with similar credit scores.

My collaborators and I find that these design choices matter—African American guests (and hosts) are discriminated against on Airbnb in a way that would be infeasible under different design choices (Edelman and Luca 2014; Edelman, Luca, and Svirsky 2016). Looking at listings in New York City, Edelman and Luca (2014) find that African American hosts earn about 12% less than white hosts for similar listings. Edelman, Luca, and Svirsky (2016) implement a field experiment across five cities, varying only the race (as statistically signaled by the name of the guest, following the methodological approach of Bertrand and Mullainathan) of the person requesting to stay with a host. African American guests are roughly 16% less likely to be accepted relative to white guests, holding all else constant.

B. Labor Markets

Looking at the online labor market Upwork (which was called oDesk at the time), Ghani et al. (2014) find that relative to other employers, employers of Indian descent are more likely to hire Indian nationals.¹ Given the design of Upwork, which provides pictures of employees during the recruiting process, one might expect other forms of discrimination to be prevalent as well. Figure 1 presents results from a search for potential employees on Upwork. A variety of more specialized labor markets have arisen in the digital age as well, ranging from Topcoder (specializing in programmers) to Uber and Lyft (specializing in drivers). As of 2015, roughly 0.5% of all US workers provided contract services through online marketplaces (Katz and Krueger 2016).

The rapid growth of online labor markets has the potential to increase or decrease discrimination. For example, consider the taxi industry, an industry in which men are historically overrepresented. As of 2014, there were more females on Uber relative to the traditional taxi industry (Hall and Krueger 2015). New data from a survey done by the Benenson Strategy Group shows that these numbers are far from static—in fact, there has been a 30% increase in female Uber drivers in the past year alone.²

Online labor markets have a unique opportunity to prevent discrimination using various tools, some of which we discuss in the next section.

C. Credit Markets

Credit and funding markets are rapidly changing as peer-to-peer lending and crowdfunding platforms are on the rise. On peer-to-peer lend-

Data Entry Specialists

Post A Job. It's Free!



Dana Evans
Data Entry

United States - Tests: 4 - Portfolio: 2

My objective is to find work I can do from home. My strongest skills are data entry and web research. I've taken college classes on Excel and the Microsoft Office ... more

Microsoft Excel

100% Job Success

\$15.00 /hr

623 hours



Fred Gotterson
Data Entry professional

South Africa - Tests: 5

Willingness to learn new skills and communicating effectively has broadened my scope of capabilities to not only data entry and internet research but also the

TOP RATED

100% Job Success

\$11.11 /hr

8,589 hours

Fig. 1. Applicants on Upwork

ing platforms such as Prosper.com, would-be borrowers post profiles and an amount they are looking to borrow. Would-be lenders select among borrowers that they would like to fund. Looking at Prosper, Pope and Snyder (2011) find that loan listings for black borrowers are 25–35% less likely to get funded than loan listings for white borrowers with similar credit scores. However, black borrowers are also more likely to default on a loan through Prosper, leading the authors to conclude that this is statistical discrimination. In their context, reducing discrimination would also reduce efficiency, unless the platform were to supplement listings with further information to predict default rates.

D. Other Markets

Other markets face discrimination challenges as well. For example, Doleac and Stein (2013) implement a field experiment on Craigslist in which they sell used iPods. In the posted pictures, they vary the hand that is holding the iPod. They find that the demand for the iPod is lower when the hand holding it is African American. Looking at sales of baseball cards on eBay using a similar design, Ayres et al. (2015) find that cards held by African American sellers earn roughly 20% less. Nunley et al. (2011) also find discrimination on eBay, but note that the extent of discrimination also depends on the amount of competition. In their design, the name of the seller is varied (as opposed to the picture). Relative to platforms such as Airbnb and Upwork, both eBay and Craigslist

have less of a norm for sellers to post personal pictures of themselves. Hence, while discrimination can exist on eBay and Craigslist, these platforms are doing less to facilitate it than are platforms that encourage (and in some cases require) users to publicly post information that signals their race, gender, or other personal information.

E. Debiasing Marketplaces

There are many market design solutions that online platforms could implement in order to reduce discrimination. Moreover, there is variation in the choices that platforms make in this area that can influence the extent of discrimination. In this section, I lay out potential ways to reduce discrimination in an online marketplace. This is not intended to be a comprehensive list, but rather to highlight the fact that there are a variety of market design solutions that could be implemented by a well-intentioned marketplace.

Optimal Information Provision

Until recently, most musicians within major symphony orchestras were male. Over time, gender equality has increased. Goldin and Rouse (2001) find that roughly 25% of the change in composition was driven by one simple change—blind auditions. Historically, when a musician would audition for the orchestra, he or she would be both heard and watched by the evaluating committee. Over time, there has been a shift in norms toward blind auditions; now, there is typically a screen between the musician and the evaluating committee. This makes factors such as looks, race, and gender less salient, allowing the committee to focus on the quality of the music.

This highlights the fact that while information is necessary to make markets function, more information is not always better. Providing information about factors such as race and gender can facilitate discrimination; removing this information can moderate discrimination effects.

In offline settings, it is often difficult to blind transactions. For example, it would be challenging for an interviewer to put a screen up in front of a candidate that she is interviewing. One of the unique features about *online* marketplaces is that they provide an opportunity to decide when and where virtual screens should be placed. In the context of Airbnb, it would be a trivial exercise to remove identifying information about guests and hosts. For example, they could limit pictures,

names, and racial identifiers of hosts until after a transaction is completed. Similarly, they could remove elements of guest profiles that provided race information until after the host had accepted a reservation request.

In practice, there is considerable variation in the amount and timing of information that is provided to buyers and sellers in online marketplaces. For example, Airbnb and HomeAway (a main competitor of Airbnb) take different approaches to information provision. Complete guest profiles (many of which have pictures—in our sample, we saw roughly 40% had pictures) are more the norm on Airbnb. Figures 2 and 3 display typical guest profiles on Airbnb, highlighting the differential prevalence of pictures and social profiles across the two platforms. Differences also become apparent when looking at the search process. When searching on Airbnb, results show the pictures of hosts alongside pictures of listings on the main results page. In contrast, HomeAway shows only the pictures of listings; users need to click on a listing before seeing host information. Figures 4 and 5 provide sample results for each platform.

An online platform could choose not to provide any identifying characteristics of buyers and sellers until after the transaction is completed. For example, both eBay and Amazon function successfully with little information to identify personal characteristics such as race and gender and rely mostly on user ratings.



Catherine was very attentive and kind host! Staying with her was a very peaceful and warm experience. She made me feel very comfortable and made sure I had everything I needed. I would stay with her again, and recommend her room- not only for her hosting skills, but also for the ideal location. I could walk everywhere from her apt, and had a wonderful time exploring town on foot!



Catherine is lovely, very sweet and accommodating. She gave lots of suggestions on where to go and what to visit and was very flexible with hours which I thought was fantastic. The apartment is centrally located and it's easy to get almost anywhere. Would stay again on a next visit.

Fig. 2. Sample guests on Airbnb

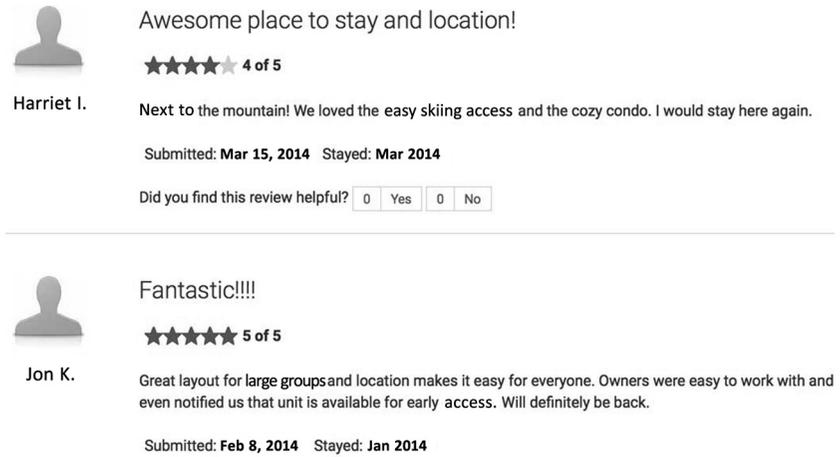


Fig. 3. Sample guests on HomeAway

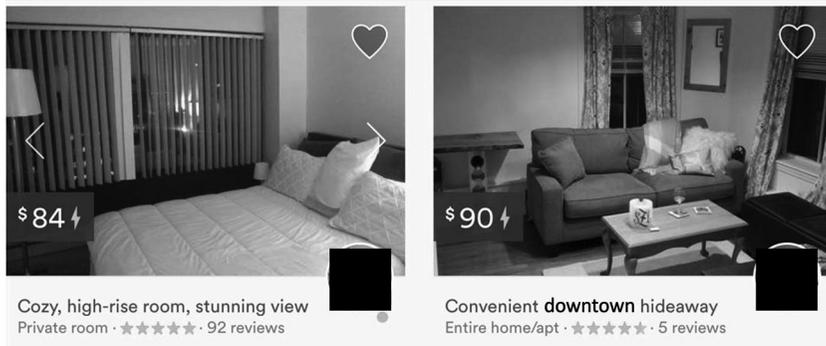


Fig. 4. Sample search result on Airbnb

Increasing Salience of Diversity Goals

An alternative (and potentially complementary) approach would be to make the objective of increasing diversity more top-of-mind for users. For example, Airbnb could require hosts to check an acknowledgment box that says, “Airbnb prohibits discriminatory behavior” each time that a host chooses whether to accept or reject a guest. Similar to the way that priming ethics can induce ethical behavior (Shu et al. 2014), this might prime users to think more carefully about the possibility that bias is creeping into their decisions.

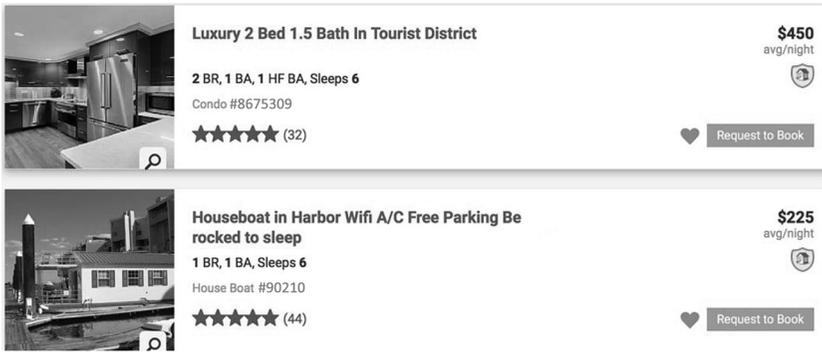


Fig. 5. Sample search result on HomeAway

Automating the Transaction Process

Another approach to reducing discrimination on online platforms is to further automate the transaction process. For example, Uber riders do not see the picture or name of the driver until after the ride is booked. While in principle, riders could cancel a ride and look for a new driver, this design choice can greatly reduce the scope for discrimination. Similarly, Airbnb has a feature called “instant book,” in which hosts can opt to automatically accept qualified guests. By removing the ability to reject guests on the basis of features such as name and picture, this virtually eliminates the potential to discriminate. Yet instant booking is currently only used by a small fraction of Airbnb hosts. If online platforms were to move toward further automation of booking, this could help to reduce discrimination in a variety of online contexts.

V. Conclusion

The evolution of trust on the Internet and development of principles for designing reputation systems have allowed online marketplaces to thrive. Early work suggested that the arms-length nature of online transactions has the potential to reduce the amount of discrimination we see. And in some cases, platforms likely are bringing us closer to the bias-free ideal. For example, Ayres, Vars, and Zakariya (2005) find evidence of racial bias in the context of tipping taxi drivers. Because it disallows tipping and instead provides market prices, Uber’s design choice eliminates this inequality among driver wages.

But the amount of discrimination in an online marketplace is a choice variable, determined by the design features that the marketplace selects. Ultimately, the decisions made by platforms will depend on ethical considerations, legal considerations, operational considerations, competition considerations, and public relations considerations. This paper has provided insight into the design principles underlying the design of reputation systems, and a tool kit for helping to reduce discrimination in the digital age.

Endnotes

1. There is extensive evidence of discrimination in offline labor markets, which is beyond the scope of this chapter. See, for example, Bertrand and Mullainathan (2004) and Lang and Lehman (2011).
2. Benenson Strategy Group: “Uber: The Driver Roadmap 2.0.”

References

- Akoglu, L., R. Chandy, and C. Faloutsos. 2013. “Opinion Fraud Detection in Online Reviews by Network Effects.” The 7th International AAAI Conference on Weblogs and Social Media.
- Ayres, Ian, and Peter Siegelman. 1995. “Race and Gender Discrimination in Bargaining for a New Car.” *American Economic Review* 85 (3): 304–21.
- Ayres, Ian, Frederick Vars, and Nasser Zakariya. 2005. “To Insure Prejudice: Racial Disparities in Taxicab Tipping.” *Yale Law Journal* 114 (7): 1613–74.
- Ayers, Ian, Mahzarin Banaji, and Christine Jolls. 2015. “Race effects on eBay.” *The RAND Journal of Economics* 46 (4): 891–917.
- Bertrand, Marianne, and Sendhil Mullainathan. 2004. “Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination.” *American Economic Review* 94 (4): 991–1013.
- Bolton, Gary, Ben Greiner, and Axel Ockenfels. 2013. “Engineering Trust: Reciprocity in the Production of Reputation Information.” *Management Science* 59 (2): 265–85.
- Charness, Gary, Eman Haruvy, and Doron Sonsino. 2007. “Social Distance and Reciprocity: An Internet Experiment.” *Journal of Economic Behavior and Organization* 63 (1): 88–103.
- Dai, Weijia, Ginger Jin, Jungmin Lee, and Michael Luca. 2014. “Optimal Aggregation of Consumer Ratings: An Application to Yelp.com.” Working Paper no. 18567, NBER, Cambridge, MA.
- Doleac, Jennifer, and Luke Stein. 2013. “The Visible Hand: Race and Online Market Outcomes.” *Economic Journal* 123 (572): F469–92.
- Edelman, Benjamin, and Michael Luca. 2014. “Digital Discrimination: The Case of Airbnb.com.” HBS Working Paper no. 14-054, Harvard Business School.
- Edelman, Benjamin, Michael Luca, and Daniel Svirsky. 2016. “Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment.” *American Economic Journal: Applied Economics* (forthcoming).
- Fradkin, Andrey, Elena Grewal, David Holtz, and Matthew Pearson. 2015.

- "Bias and Reciprocity in Online Reviews: Evidence from Field Experiments on Airbnb.com." In *Proceedings of the Sixteenth ACM Conference on Economics and Computation* (EC '15), 641. New York: Association for Computing Machinery. doi: <http://dx.doi.org/10.1145/2764468.2764528>.
- Ghani, Ejaz, William Kerr, and Christopher Stanton. 2014. "Diasporas and Outsourcing: Evidence from oDesk and India." *Management Science* 60 (7): 1677–97.
- Goldin, Claudia, and Cecilia Rouse. 2001. "Orchestrating Impartiality: The Impact of 'Blind' Auditions on Female Musicians." *American Economic Review* 90 (4): 715–41.
- Hall, Jonathan, and Alan Krueger. 2015. "An Analysis of the Labor Market for Uber's Driver-Partners in the United States." IRS Working Paper no. 587, Industrial Relations Section, Princeton University.
- Hitt, Lorin and Xinxin Li (2008). "Self-selection and information role of online product reviews." *Information Systems Research*, 19: 456–474.
- Hoffman, Elizabeth, Kevin McCabe, and Vernon Smith. 1996. "Social Distance and Other-Regarding Behavior in Dictator Games." *American Economic Review* 86 (3): 653–60.
- Katz, Lawrence, and Alan Krueger. 2016. "The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015." Working Paper.
- Lang, Kevin, and Jee-Yeon K. Lehman. 2011. "Racial Discrimination in the Labor Market: Theory and Empirics." *Journal of Economic Literature* 50 (4): 959–1006.
- Luca, Michael. Forthcoming. "User-Generated Content and Social Media." *Handbook of Media Economics*.
- Luca, Michael, and Georgios Zervas. 2015. "Fake It till You Make It: Reputation, Competition, and Yelp Review Fraud." Harvard Business School Working Paper no. 14-006, Harvard University. Forthcoming in *Management Science*.
- Masterov, Dimitriy, Uwe Meyer, and Steven Tadelis. 2015. "Canary in the e-Commerce Coal Mine: Detecting and Predicting Poor Experiences Using Buyer-to-Seller Messages." *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. doi: 10.1145/2764468.2764499.
- Mayzlin, Dina, Yaniv Dover, and Judith Chevalier. 2014. "Promotional Reviews: An Empirical Investigation of Online Review Manipulation." *American Economic Review* 104 (8): 2421–55.
- Nunley, John M., Mark F. Owens, and R. Stephen Howard. 2011. "The Effects of Information and Competition on Racial Discrimination: Evidence from a Field Experiment." *Journal of Economic Behavior & Organization* 80 (3): 670–79.
- Ott, M., Y. Choi, C. Cardie, and J. Hancock. 2011. "Finding Deceptive Opinion Spam by Any Stretch of the Imagination." In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, 309–19.
- Pope, Devin, and Justin Sydnor. 2011. "What's in a Picture? Evidence of Discrimination from Prosper.com." *Journal of Human Resources* 46 (1): 53–92.
- Scott Morton, Fiona, Florian Zettelmeyer, and Jorge Silva-Risso. 2003. "Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?" *Quantitative Marketing and Economics* 1 (1): 65–92.
- Shu, Lisa L., Nina Mazar, Francesca Gino, Dan Ariely, and Max H. Bazerman. 2014. "Signing at the Beginning Makes Ethics Salient and Decreases Dishonest Self-Reports in Comparison to Signing at the End." *Proceedings of the National Academy of Sciences* 109 (38): 15197–15200.