DIGITAL MARKET DESIGN AND INEQUALITY

Digital Market Design and Inequality

Digital marketplaces and platforms are important and their design affects the rest of the economy. These markets are also constantly evolving due to the design decisions of firms, regulation from government, and new technological capabilities. The power of platform companies and the dependence of users on these platforms has raised concerns of fairness and accountability. For example, market design changes by Uber, Etsy, and TaskRabbit[1] have, at times, resulted in vocal complaints by users. Furthermore, there is a vigorous debate regarding the effects of digital markets on economy wide inequality.

In this paper, I discuss whether and when reducing inequality should be a criteria for the design of these markets. I start by describing the special role of data in the design of digital platforms and the potential effects of privacy regulation. Next, I discuss firms' market design decisions more broadly and argue that that firms should care about inequality for reasons of managing complexity, increasing transparency, and reducing regulatory pressure. I end by discussing recent work on the effects of platforms on the rest of the economy and the implications of this research for policy.

Data and Digital Market Design

Perhaps the largest difference between digital market design and traditional market design is the importance of data analysis and data intensive algorithms. For example, the design approach in applications such as school choice, medical matching, and kidney exchange cited in the Nobel Prize of Al Roth and Lloyd Shapley[2] is theoretical and focused on incentivizing agents to truthfully report preferences and participate in exchange.

Contrast this to the design of an online search engine. The search engine retrieves the information that is most relevant for each query combination. The quality of search engine results is a function of an algorithm, which is optimized by combining massive amounts of data, novel machine learning methods, and large scale experimentation. A naive view of the algorithm design is that it is purely an optimization problem. However, this algorithm serves as a medium for the strategic actions of independent website owners, who optimize their website to maximize

rankings, and advertisers who bid on queries. Therefore, both the data used to determine the ranking and the objective function against which the ranking is evaluated affect economic outcomes for a variety of agents in the economy.

More generally, data is used not only in algorithmic design but throughout other parts of the product (e.g. reviews and likes), as an input into future product design changes, and to help with high level management decisions. The types of data which are useful vary by application and industry. For example, Data on product failures for hardware and user clicks for website and apps is critical for product design. Data also has option value. It is not always clear which signal or user action will be a useful component of an algorithm or analysis that aids in decision making. Lastly, data storage is also useful for auditing, both for internal and regulatory purposes.

However, the collection of more data has its downsides. Data can be leaked, stolen, or seized by the government. This can have effects on firms' business success and valuations.[3] Furthermore, firms with reputations for privacy may be able to attract more consumers. A reputation for privacy may be especially important to users at risk of government enforcement actions, whether for legitimate or illegitimate reasons. Lastly, data usage that conforms with customer expectations may be good for its own sake.

Consequently, the decision of what data to collect is a key strategic decision by the firm. The importance of this decision can be seen in the diverging data strategies of technology companies. On the side of less data collection and more privacy, Apple touts its focus on privacy preserving technology and Duck Duck Go is a search engine which does not store or use personal data. On the other side, Google's Allo messaging product records and stores all conversations between users. However, most aspects of data use policy, such as the effort of companies to secure data, the internal process for using data, and the compliance of companies with government requests for data access, remain opaque.

Regulation, Privacy, and Inequality

Many governments have chosen to regulate the use of data by companies. The typical motivation for this regulation is to ensure that data is used in accordance with customer expectations. Governments argue that without regulation, firms over-invest in data collection and use relative to the preferences of consumers. In

"Inequality, Privacy, and Digital Market Design", Goldfarb and Tucker discuss the effects of these privacy regulations. Their paper treats the government as a market designer whose job is to balance the benefits of data usage by firms with the privacy costs of that data usage.

Goldfarb and Tucker use examples from a variety of industries to illustrate the fact that privacy regulation has economic effects and that there is heterogeneity in effects across income groups. It is not surprising, a priori, that any policy, regardless of whether it concerns privacy, affects inequality. Therefore, whether an inequality effect should affect policy decisions depends on the exact effects of the policy. The authors make a convincing case, using examples such as insurance and electronic medical records, that privacy regulation sometimes hurts lower income groups and that sometimes it helps. Consequently, they suggest that privacy policy should be done on an industry by industry basis.

The most striking aspect of the paper is how little evidence there is of the equilibrium implications of privacy policy. Take for example the case of targeting in advertising. The authors discuss research which shows that European privacy regulation in the early 2000s reduced the effectiveness of online advertising for general interest sites which have a diverse set of visitors. On the other hand, product-specific sites such as travel sites did not experience a decrease in ad effectiveness. The authors use this result to argue that the equilibrium mix of sites will shift away from general interest sites and towards sites frequented by rich individuals. While this may be true on the margin, the paper's discussion gives little sense of the quantitative importance of this adjustment. If it is minor, as seems likely given the low costs of operating a typical website, then this should not be a key determinant of the policy decision. However, it is also easy to imagine that general interest sites which support resource intensive investigative journalism or other publicly valuable enterprises could be particularly hurt. More evidence is needed.

More generally, loose privacy regulation supports business models where consumers pay nothing and advertising generates the platform's revenue. This arrangement seems particularly helpful to lower income individuals who are more price sensitive. However, the key question is not whether data is useful, but whether particular privacy policy has important effects on the margin. One particular concern is with regards to entry costs. Businesses such as Google or Facebook may still be able to offer a comparable quality service even if their data

use were curtailed by regulation. However, smaller firms may not have the resources to comply with privacy regulations and offer a high quality service. In this case, privacy regulation may serve as an inequality increasing barrier to entry in these industries.

A similar dearth of evidence exists for the other examples mentioned in Goldfarb and Tucker's paper including price discrimination, search personalization, and insurance. A priori, regulation regarding privacy has ambiguously signed and sized effects in these industries. Any given study, especially one which studies short-run partial equilibrium effects, is unlikely to be dispositive on the topic. One potential solution to this problem is to look for complimentary regulations. For example, the authors point to evidence that lower income individuals are targeted more frequently by manipulative advertising. If manipulative advertising is so bad, then it should be regulated directly rather than in a circuitous route using privacy policy.

Lastly, the authors omit important considerations relating to data regulation and privacy: data safety, usage transparency, and information disclosure. One of the reasons that consumers do not want their data shared is that that data may be leaked or stolen. If this is a primary concern, then regulation regarding the manner in which data is secured within the company may reduce the fear of data usage by firms while retaining the benefits of data usage. Relatedly, it may be efficient for governments or other organizations to audit the data storage practices to ensure minimal risk. If these policies succeeded in increasing the trust that consumers place in firms, then strict rules regarding the usage of data could be less important.

Another policy response may be to combine choice based approaches with better information regarding the potential costs and benefits of data sharing. Consider the case of privacy regulation and electronic medical records. The work of Miller and Tucker shows that EMRs improved health outcomes but that privacy reduced the spread of EMRs. Specifically, EMRs were less likely to be used in states which required the active consent of users regarding data sharing. One potential way to alleviate this effect is to clearly state the benefits of EMRs when the consumer makes a choice. If consumers are responsive to this information, then the costs of policies requiring consent may be mitigated.

Digital Market Design and Inequality: The Firm's Perspective

Choices regarding market design are critical for a diverse set of digital platforms including search engines, retailers, social networks, and peer-to-peer marketplaces. Consequently, these platforms have been introducing novel market mechanisms since their inception. In this section, I use Hahn and Metcalfe's discussion of ridesharing as a starting point in discussing the role of design decisions on inequality. I argue that there are good reasons why firms may want to sacrifice short-term gains in order to keep their platforms more equitable.

Hahn and Metcalfe document the novel design features of ride-sharing platforms. These include a GPS based matching mechanism, a reputation system, "surge pricing", and a payments platform. The first order effect of these features is positive for the ride-sharing platforms and is a likely reason why the generate so much consumer surplus.[4][5][6] GPS based matching makes it easier and faster to find a ride, reputation reduces the problem of asymmetric information and moral hazard, and `surge pricing' clears the market when the amount of drivers is relatively small. However, there is a lot of nuance in the design of each of these systems, which may affect within platform inequality in a meaningful way.

Within platform inequality is worth considering when choosing market designs. As an example consider the effects of a change in the reputation system design of a ride-sharing platform. Suppose that a ride-sharing platform experiments with raising the threshold rating for drivers participating in a platform. Also suppose that an A/B test conducted for two weeks shows that riders in the treatment left better reviews and matched just as fast on average. This policy may seem unambiguously good when looking at simple treatment effects. However, the information relayed by the A/B test is not sufficient to determine whether a new ratings policy should be enacted.

The true effect of a policy may be complicated by the following concerns. First, it may take a while for the market to adjust to a new equilibrium. Lower rated drivers may have placed downward pressure on prices on the platform, and this effect could take time to manifest itself. Second, simple marketplace experiments often have problems of interference which make experimental treatment effects invalid [7]. Third, the experiment may have redistributed the supply on the platform in a consequential manner. For example, lower rated drivers may have been serving different neighborhoods than higher rated drivers. This could affect the competitive position of the platform in an area. Lastly, drivers hurt by this policy could protest the new policy as unfair. These concerns may potentially be valid. For example,

ratings may be driven by illegitimate reasons such as race or accent. In fact, even a rating mistake by a rider may potentially lower the average rating of a driver below a threshold. Driver complaints about unfair practices may be damaging to the platform's reputation and may solicit regulatory pressure.

Even with the above concerns, the platform may find it advantageous to increase the ratings threshold. However, a focus on the inequality effects of a policy guides the market designer to important concerns which are easy to miss by looking at simple dashboards.

Understanding the effects of design on inequality may also aid in finding complementary market mechanisms. For example, while surge pricing can improve market efficiency, it may be disproportionately harmful to lower income individuals. Those with liquidity constraints may choose to not rely on ride sharing due to the risk of high prices. Recognizing this, market designers may find it advantageous to offer surge price insurance for these individuals. Indeed, Uber has begun offering monthly passes which guarantee low prices for riders regardless of market conditions.[8]

The above discussion has barely skimmed the complex effects of market designs. Other important market designs which have inequality effects include the choice of search ranking algorithm, the use of the instant book feature on Airbnb, the specific rules of auction bidding, and many others. What is clear is that the effects of market design on equilibrium outcomes are complex, and require a nuanced understanding of not only the simple treatment effects but also of the distributional consequences.

Market Design and Inequality: A Societal Perspective

Digital platforms have pervasive effects on the rest of the economy and, as has been stress throughout this paper, also have distributional consequences. There are three types of agents affected by new digital platforms: consumers, suppliers, and asset owners. Given these complex effects, I will focus the discussion on the ride-sharing vertical.

Successful digital platforms such as Uber must, by the virtue of their success, provide at least some consumer surplus above those of the old suppliers (taxis).[9] As emphasized in Hanh and Metcalfe, there has been little study as to the

distribution of this surplus. It seems that both discriminated against minorities and those living in poorer areas have benefited from having the ride-sharing service. Yet racism still persist on these platforms.[10] Furthermore, ride-sharing is still relatively expensive compared to public transport in a lot of cases. So the very poor, especially those without a cell phone, may have been overlooked by these platforms.

Perhaps more interesting is to speculate about the future effects of the platforms. It seems likely that the market penetration of smart phones will be ubiquitous in the near future, meaning that even the poorest would have access to services such as Uber and Lyft. However, ride-sharing may also affect equilibrium demand for buses and trains. This could reduce middle-class and elite support for public transport, hurting the poorest people. More research on the substitutability of demand for public transport and ride-sharing is needed to understand these concerns.

Ride-sharing is also likely to interact with congestion in a meaningful way. Hahn and Metcalfe summarize the existing evidence on congestion and conclude that it is inconclusive. One intriguing possibility is that due to the digital nature of ride-sharing, it may be easier to implement congestion pricing to combat congestion. The effectiveness of congestion pricing in combatting congestion is well known, and the possibility of implementing congestion pricing is an important benefit of the digitization of transport. However, yet again, congestion pricing is likely to have inequality effects and may especially hurt poorer individuals who commute to work.

Next, I consider the effects of digital platforms on suppliers. It is apparent that ride-sharing platforms hurt the owners of taxi medallions, although there is no rigorous academic evidence on this point. The effects of ride-sharing on taxi drivers are more ambiguous because taxi drivers can also supply their services on ride-sharing platform. New platforms both increase market demand and lower the barriers to entry in these markets, creating more competition. Which of these mechanisms dominates likely varies by market. Cramer (2016) shows that, at least for the US, traditional labor in the taxi industry hasn't been hurt by Uber up to this point. Another aspect of the job which may be affected is it's quality. As discussed in the previous section, depending on the platform rules, drivers may have more or less choice regarding whom to drive, how much to drive, and where to drive. This amenity is important to consider in analyzing these markets and it's value may change with the composition of drivers in equilibrium.

Lastly, and perhaps most importantly for inequality, digital platforms have tended to generate highly concentrated returns for founders, early investors, and early employees. This raises concerns about proper returns for these early platform participants. There are two extremes views regarding this topic. The first view states that platforms are innovators and are constantly under the threat of competition. In that case, typical economic arguments suggest that the returns to successful platforms should be high. On the other hand, if these platforms constitute natural monopolies with overwhelming network effects, then these rewards are less deserved. Understanding which of these scenarios is closer to the truth will be an important topic of research and policy discussion for many years to come.

Conclusion

Digital platforms are ubiquitous. The decisions of engineers in Silicon Valley and other technology hubs have broad societal effects. In this paper, I've tried to synthesize a nascent literature focusing on the distributional effects of digital market design. As is evident throughout the paper, we are solely lacking in evidence. However, even without good evidence, market design activities by firms and regulation by governments will continue. It is the unique responsibility of academics to produce objective research to guide these decisions.

---

1. "Hunting Task Wabbits" https://medium.com/ondemand/hunting-task-wabbits-c60679bad0f6#.2egas1w2j
2. "Stable matching: Theory, evidence, and practical design" http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2012/popular-economicsciences2012.pdf
3. Is there a cost to privacy breaches? An event study http://aisel.aisnet.org/cgi/viewcontent.cgi?article=1215&context=icis2006
4. Cohen, P., R. Hahn, J. Hall, S. Levitt and R. Metcalfe (2016), "Using Big Data to Estimate Consumer Surplus: The Case of Uber," Working paper.
5. Chen, M. K. and M. Sheldon (2015), "Dynamic Pricing in a Labor Market: Surge Pricing and Flexible Work on the Uber Platform." Working Paper.

6. Cachon P., K. M. Daniels, and R. Lobel (2016). “The role of surge pricing on a service platform with self-scheduling capacity.”
7. Andrey Fradkin. “Search Frictions and the Design of Online Marketplaces.” Working Paper.
8. https://newsroom.uber.com/us-dc/introducing-the-uberpool-monthly-ride-pass/
9. Farronato and Fradkin (2016) show that Airbnb has also generated substantial consumer surplus in the accommodations market.
10. Ge, Knittel, MacKenzie, and Zoepf. Racial and Gender Discrimination in Transportation Network Companies. 2016.