CS269I: Incentives in Computer Science
Lecture #12: Asymmetric Information and Reputation Systems*

Tim Roughgarden†

November 2, 2016

1 Preamble

Previous lectures have skirted the topic of reputation systems—the BitTorrent reference client (Lecture #5) can be viewed as using a primitive reputation system, and reputation systems also came up in passing in our discussion of crowdsourcing systems (Lecture #10). In this lecture we tackle reputation systems head-on.

Reputation systems have been around as long as trading itself. They are also a crucial ingredient of online platforms, as is familiar from eBay, Amazon, Uber, Airbnb, and so on.

2 The Market for Lemons

Much of this lecture is about identifying the problem(s) that a reputation system is designed to solve. The first one is called the "adverse selection" problem, which shows how market failure can occur when one side of the market (buyers or sellers) has more information about the value of products than the other side. We first describe the academically most famous example, and then relate the lessons learned back to online marketplaces.

George Akerlof devised the following story, known as the "market for lemons," to illustrate the perils of information asymmetry [1].[1] Here "lemon" does not refer to the citrus fruit, but rather to a used car that no longer works.

Suppose there are multiple sellers, and each seller either has a good used car or a bad used car. Each seller knows whether or not her car is good or bad. Suppose there are also multiple identical buyers—and for concreteness, assume that there are more buyers than

†Department of Computer Science, Stanford University, 474 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: `tim@cs.stanford.edu`.

[1]See also the discussion in [2, Chapter 22].

sellers. (This is not an important assumption.) Suppose a good car has value 12 to a buyer and 10 to the seller, while a bad car has value 6 to a buyer and 4 to the seller. Note that buyers value all cars more than sellers, so the outcome we want is for all cars to be sold, at fair prices (between 10 and 12 for good cars, between 4 and 6 for bad cars).

The easy case is when the buyers have as much information as the sellers—when they can distinguish good cars from bad cars. Then we expect all of the cars to sell, the good cars for 12 and the bad cars for 6. (Since there are more buyers than sellers, the competition on the demand side should push prices up to the maximum buyers are willing to pay.)

The interesting and more realistic case is when buyers have less information about car quality than the sellers. To keep the story as simple as possible, let's make the extreme assumption that all of the cars look identical to the buyers. So if a buyer buys a car, it's a good car with probability equal to the fraction of the cars on the market that are good cars. Note that with indistinguishable cars, all cars should sell for roughly the same price.

Suppose that a $g \in [0, 1]$ fraction of the cars that could possibly be put on the market are good, and that everybody knows g. Let h denote the fraction of good cars among those that are actually put up for sale (a seller can elect to not make her car available for sale, if the going price is too low). So for example, if all bad cars are put on the market but some sellers of good cars withhold their cars from the market, then $h < g$.

So what is the value of h at equilibrium? As usual, we appeal to intuition about what constitutes an equilibrium. Basically, the trades that occur should benefit both parties, and all sellers not in the market do not want to enter.[2]

One possibility is that $h = 0$. This means that all of the cars for sale are bad cars. Since a buyer is assured that the car she is buying is bad, she is willing to pay up to 6 for a car. The seller of a bad car benefits from selling her car for 6, so no owner of a bad car wants to drop out of the market. The opposite is true for sellers of good cars: the price of 6 is below the seller's value, so they do not want to enter the market. Thus $h = 0$ (with a price of 6) is an equilibrium, in the sense that it is self-reinforcing. (This is true no matter what g is.)

We might also hope or expect that $h = g$ constitutes an equilibrium. This will depend on the value of g. To see this, note that if buyers know that a g fraction of the cars for sale are good, then the price a buyer is willing to pay is

$$12g + 6(1 - g) = 6 + 6g.$$

(For simplicity, we assume that buyers only care about expected values.) As expected, the willingness to pay interpolates between 6 and 12, linearly with the fraction of good cars for sale. Are sellers willing to accept a price of $6 + 6g$? A seller with a good car (who values it at 10) would accept this price if and only if it is at least 10, that is, if and only if $g \geq \frac{2}{3}$. Obviously, sellers of bad cars would be ecstatic to sell at such a price.

So in the regime with lots of good cars ($g \geq \frac{2}{3}$), the case of $h = g$ is also an equilibrium (with price $6 + 6g$)—all trades make both sides happy, and no seller wants to drop out of the market. The interesting point here is that there are two very different equilibria, each self-reinforcing. If all buyers believe that there are only bad cars and hence only make low

[2]This type of equilibrium is sometimes called a *rational expectations equilibrium*.

offers, then indeed no seller of a good car will participate. But if all buyers believe that all good cars are on the market and hence are willing to make higher offers, then sure enough, it will be the case that all of the good cars will be on the market. In either case, buyers' beliefs are a self-fulfilling prophecy.[3]

In the regime with less good cars ($g < \frac{2}{3}$), the price $6 + 6g < 10$ is too low to be accepted by owners of good cars, so they will exit the market. Thus $h = 0$ is the only equilibrium in this case. This is a "market failure," in the sense that the outcome of the market is not Pareto optimal. (Since buyers value cars more than sellers, the only way to be Pareto optimal is to sell all the cars—otherwise, there is a mutually profitable exchange that has been left unexecuted.)

This particular type of market failure is called *adverse selection*. The culprit is asymmetric information—the goods on the market have different qualities, but only one side of the market knows what the quality is. The result is that only low-quality goods remain on the market (for a low price). Thus the market adversely selects for the worst goods.

It can get worse. Suppose in addition to good and bad used cars, there are also *lemons*, which are valued at 0 by both sellers and buyers. Suppose there is initially a (publicly known) $\frac{1}{3}/\frac{1}{3}/\frac{1}{3}$ split between good/bad/lemons. If all cars are put on the market, then the going price will be 6 (why?). With such a low price, all the sellers of good cars will exit. This depresses the selling price further: with a 50/50 split between bad cars and lemons, buyers will only be willing to pay 3 for a car. At this price, even the sellers of bad cars will exit, leaving a market consisting solely of lemons.

3 Asymmetric Information and Adverse Selection

3.1 More Examples

The key trigger of adverse selection is asymmetric information, when the two sides of a market have different information about the value of the goods being sold. The result is that only the lowest-value products remain in the market. There are a number of famous examples of adverse selection. For example, consider the market for health insurance.[4] In this market, the goods being sold are health care policies. The sellers are insurance companies, and the buyers are individuals. The value of the good to a buyer (and also the cost to the seller) is the expected amount of health costs that the insurance company will need to reimburse. Who has more information, buyers or sellers? Generally, the *buyers*—an individual is better-informed about her own health than an insurance company is. (Note this is flipped from the market for lemons example.) If the insurance company cannot distinguish between different individuals, then it has to offer the same price to everybody. And at a given price, who will buy it? The buyers for whom it has the most value—*the least healthy*

[3]When you hear the term "consumer confidence" on the news, this is usually the type of phenomenon they're getting at.

[4]When Obamacare was originally proposed, it was common to see the phrase "adverse selection" in relatively mainstream media.

people. That is, if 50% of the population buys health insurance, then it will be the least healthy half of the population. Knowing this, the insurance company must charge more for health insurance. But this drives still more people out of the market (the most healthy ones who remained), which drives up the price further, and so on. In the extreme, the end result of adverse selection is very high prices for health insurance, with only very unhealthy people purchasing it. One of the points of universal health care is to mitigate this adverse selection problem—if everyone purchases health insurance (especially healthy people), then health insurance can be less costly.

Another famous example is the labor market. Here, buyers are firms, sellers are workers, and the goods are also the workers. One can imagine that some workers are more productive than others, and accordingly more valuable to a firm. Like in the market for lemons, here the sellers are typically better informed about the value of the goods (i.e., their own productivity) than the firms. In the extreme case, if a firm cannot distinguish different workers at all, then it has to pay all of them a common wage. This can lead to the exit of the most productive workers (who perhaps can make more money than the common wage by staying self-employed), which in turn depresses the wages firms are willing to pay, which leads to the exit of more workers (the most productive ones who remained), and so on. The end result is a market with only low-productivity workers, all earning low wages.

For an example more in the scope of this course, consider the online advertising market, such as the market for sponsored links on a search engine results page. The sellers are the advertisers, and the goods being sold are clicks to various landing pages. The buyers are search engine users. Clicking on a sponsored link corresponds to buying a good. The value of the good (to a buyer) corresponds to the quality of the landing page. Note that the search engine itself is neither the buyer nor the seller, it's just the platform that enables the transaction, and it collects a transaction fee in the form of a payment for each click. There is again asymmetric information, with advertisers much better informed about the quality of their own landing page than the buyers. If the buyer is unable to distinguish between high-quality and low-quality ads (and the search engine doesn't do it for her), then her propensity to click will be driven by the average ad quality. If high-quality ads are not treated as such, it can lead to the exit of the highest-quality ads from the market. This lowers the average ad quality, and hence the propensity of users to click on ads, and hence the search engine's revenue.[5] The ultimate result would be a market of pure clickbait, with users almost never clicking on ads.[6]

[5]It's tempting to speculate that something similar has been going on at Yahoo, where search ad revenue has dropped for two consecutive years even as it has increased for all other major search engines.

[6]Don't ever forget that in a system with self-interested users, changing the system also changes how the users behave in the system. You might wish that user behavior would remain fixed, but it won't. For example, it may be tempting to promote clickbait to increase clicks (and hence revenue) in the short term, but disappointed users will adjust and click less often on ads in the future.

3.2 Mitigating Adverse Selection

If adverse selection is driven by asymmetric information, then an obvious solution is to reduce information asymmetry by exposing more information about the value of the goods being sold. This boils down to giving high-quality sellers or buyers the ability to differentiate themselves from lower-quality buyers and sellers. Let's think about how this might work in the examples described above.

In the market for lemons, sellers of good cars need to be able to distinguish themselves. One way of doing this is to have the car quality verified by a trusted third party (e.g., an appropriately certified mechanic). Another way is to offer a warranty (covering all repairs in a given time window). Since the expected cost of offering a warranty is much higher to sellers of bad cars than sellers of good cars, a warranty is a credible signal of a car's quality.

In the labor market, highly productive workers need to differentiate themselves to firms. Education can be viewed as one way of doing this. Suppose that there is a positive correlation between the productivity of a worker and the ease of performing well at school. (An imperfect assumption, to be sure.) Then, obtaining degrees is less costly for productive workers than unproductive workers, and the former can signal their likely productivity by obtaining more academic degrees. Thus differentiated, firms can then pay workers with more education a higher wage. The point is that, while of course education has many direct benefits (learning lots of new ideas and skills, meeting lots of new people, etc.), independent of these, it can also be interpreted as a signaling device that mitigates adverse selection in the labor market.

Finally, in online advertising, a search engine would be wise to form accurate estimates of the quality of various ads, and to export this information to its users (e.g., by giving the highest-quality ads the most prominent placement). Google, for example, has done precisely this for many years. The benefits include: high-quality advertisers stay in the market (i.e., adverse selection is mitigated); users continue to click on ads at a healthy rate, and typically see highly relevant ads when they do so; and the search engine protects its future search ad revenue.

Now that we understand what adverse selection is, we can recognize a reputation system as a mechanism for mitigating it. A good example is Yelp. Suppose no restaurant reviews existed, and more generally that diners had no way of differentiating between restaurants without actually trying them. (Whereas a restaurant knows how good it is.) A diner would only be willing to pay for the expected value of a meal at an average restaurant, which might not be high enough for higher-quality restaurants to stay in business. (Perhaps higher-quality restaurants use more expensive ingredients or pay higher wages.) When the high-quality restaurants exit, diners will be willing to pay even less for a meal. The endgame is a scenario where the only restaurants still in business are the cheapest and worst ones. Reviews on Yelp help avoid this culinary apocalypse by exporting the information about restaurant quality to the buyers, reducing the information asymmetry.

4 Moral Hazard

In the market for lemons, the seller of a bad car wasn't a bad person—she just didn't have a very good car. In that model, sellers of bad cars couldn't just snap their fingers and change their car from bad to good. Adverse selection, in its purest form, is not caused by anyone's deliberately harmful actions.

Thinking about other online marketplaces, like eBay and Amazon, we can see that mitigating adverse selection is not the most important goal of a reputation system. In eBay, for example, the issue is not that some sellers are inherently less capable of carrying out an honest transaction than others; it's that some sellers choose to rip off buyers for their own personal gain. This problem bears more resemblance to the Prisoner's Dilemma (Lecture #5) than the market for lemons, and it has its own name: *moral hazard*.

Moral hazard basically refers to an incentive to take an action that is harmful to others. A bit more formally, it refers to an action whose cost is not fully borne by the decision-maker. Moral hazard can also be thought of as a problem of asymmetric information, where the decision-maker knows about what action she will take better than anyone else.

Let's see how moral hazard can come up in some of our running examples. In the health insurance market, we previously assumed that each individual had some immutable level of health. But your actions affect your health. If you have to bear the full cost of your medical expenses, there is a strong incentive to exert effort to stay healthy. If this cost is largely borne instead by an insurance company, then there is less incentive to stay healthy. Similarly, in the labor market, we talked about workers as if they were born into the world with a given productivity level. But workers have some control over the amount of effort they exert and their consequent productivity. If exerting effort is costly and a firm cannot directly observe it (and hence cannot compensate workers based on the chosen effort level), then a worker has an incentive to work less than she would otherwise. (In this case, the cost of not working hard is largely borne by the firm, not the worker.) In an online platform like eBay, the moral hazard problem faced by a seller is whether to honor a transaction after receiving the buyer's payment. Absent some future karmic retribution, the cost of ripping off a buyer is borne by the buyer, not the decision-maker (the seller).

A natural approach to mitigating moral hazard problems is to expose more information about the actions decision-makers take. This enables differentiation based on past actions, and provides a lever for karmic retribution. This is the primary purpose of a reputation system in an online platform like eBay or Amazon.

Many scenarios exhibit a blend of adverse selection and moral hazard. For example, consider the reputation system in Airbnb. It mitigates adverse selection through reviews of sellers—absent this information, sellers would be much more informed about the quality of their accommodations than buyers. It also mitigates some moral hazard problems—for example, if some renters throw a party and trash the apartment they rented, then they will be rewarded with a negative review and may be refused service by other sellers in the future.

5 Reputation in the Repeated Prisoner's Dilemma

The repeated Prisoner's Dilemma (Lecture #5) provides a very clean way to view the role of reputation in mitigating moral hazard problems.

5.1 Review of Repeated Prisoner's Dilemma

First, a brief recap from Lecture #5. The (single-shot) Prisoner's Dilemma is described by the following payoff matrix:

	Cooperate (C)	Defect (D)
Cooperate (C)	$2, 2$	$-1, 3$
Defect (D)	$3, -1$	$0, 0$

In each entry of the matrix, the first number is the payoff to the "row player" and the second the payoff to the "column player." Recall that defecting is a dominant strategy, meaning that it always maximizes a player's payoff, no matter what the other player does. This means there is a unique Nash equilibrium (i.e., an outcome from which neither player has an incentive to deviate unilaterally), with both players defecting. This outcome is not Pareto optimal; both players would be better off if both cooperated. Thus the Prisoner's Dilemma represents a fundamental conflict between what's good for an individual and what's good for the collective.

When the Prisoner's Dilemma is played only once between two players, one expects to see players defecting. (Recall from Lecture #5 that this was one view of free riding in Gnutella.) When the Prisoner's Dilemma is played multiple times, however, cooperation can emerge. The specific model we looked at had each player maximizing her total discounted payoff:

$$\sum_{i=1}^{\infty} [\text{utility at stage } i] \cdot \gamma^{i-1}, \tag{1}$$

where $\gamma \in (0, 1)$ is the *discount rate.* One interpretation of discounted payoffs is that, after each stage, there is a $1 - \gamma$ chance that no more stages will be played.

In a repeated game, a strategy consists of a mapping from the history-so-far (from previous stages) to a chosen action in the current stage. We looked at both the grim trigger (GT) strategy, where the player defects if and only if her opponent defected at any time in the past, and the tit-for-tat (TFT) strategy, where the player defects if and only if her opponent defected in the previous stage. (In both cases, one cooperates in the first stage.) We saw that when $\gamma \geq \frac{1}{2}$, the best response to an opponent playing GT or TFT is to always cooperate. The reason is that a defection nets you an extra payoff of 1 in the current stage, but causes a loss of at least 2 in the next stage (which is not worth it when $\gamma \geq \frac{1}{2}$).

5.2 Repeated Prisoner's Dilemma with n Players

In an online platform like eBay, participants transact in pairs (one buyer and one seller), but there are far more than two players in the system. Here's a natural extension of the repeated Prisoner's Dilemma to such settings with a large number n of players:

- Repeat:
  - Pick two random players, who play one stage of the Prisoner's Dilemma.[7]
  - The actions chosen by the two players are observed by all n players.

A player can again implement a strategy like GT or TFT by using her opponent's past actions (against all other players) as a history to respond to. So in the new version of GT, a player defects if and only if her opponent defected at some point in the past (possibly against a different player). In TFT, a player defects if and only of her opponent defected the last time she played.

It is again true that when $\gamma \geq \frac{1}{2}$, and all other players are playing TFT or GT, then the best response is to always cooperate.[8] The reason is that, from a player's perspective, payoffs are exactly as if she were playing a single player using the GT or TFT strategy.

It's important that the actions chosen by players are observed by all other players. If a player only observed the actions taken in the games that she herself plays, then she will not have an opportunity to punish her opponent until the next time they meet. Since her opponent is expected to play around n stages in the mean time, the future cost of defection is discounted by γ^n. Thus the "$\gamma \geq \frac{1}{2}$" assumption needs to be changed to "$\gamma^n \geq \frac{1}{2}$," which means that the discount rate would need to be around $1 - \frac{1}{n}$, which is unrealistically high when n is large.

We can interpret the incentive to cooperate in the n-person Prisoner's Dilemma through the lens of reputation. The assumption that all actions are observed by all is effectively assuming that a good reputation system has been implemented. (Recall one of the primary roles of a reputation system is to expose information about actions taken.) For example, with the grim trigger strategy, we can think of a player's reputation as being either "good" (if she never defected in the past) or "bad" (if she did). With the tit-for-tat strategy, a player's reputation toggles between good and bad according to the most recent action played. In both cases, every player starts with a good reputation.

5.3 Whitewashing

In Lecture #10 we talked about Sybil attacks, which is an attack where a single user creates multiple identities. This is easier in some systems (like eBay, which uses pseudonyms) than others (like Airbnb, where some degree of authentication is required). When it is easy, like in eBay, it can lead to undesirable behavior. For example, with cheap multiple identities, it is no longer a best response to cooperate in the n-person Prisoner's Dilemma when other players are playing GT or TFT. A better strategy (with 50% more payoff) is to repeatedly

[7]One can also randomly pair up all of the players and have them all play one stage in parallel.

[8]When evaluating a player's discounted payoff as in (1), by "stage i" we mean the ith stage in which this player played, not the ith stage overall.

defect, exit the system, and rejoin. This attack exploits the fact that players are initially granted a good reputation.

The attack of exiting and then rejoining a system is called *whitewashing*. There are various strategies to mitigate whitewashing. One is to increase the difficulty in obtaining multiple identities (e.g., requiring the legal name, phone number, social security number, etc.). Another is to explicitly charge an entry fee when people join. In the n-person Prisoner's Dilemma setup, with $\gamma > \frac{1}{2}$, an entry fee of 2 is sufficient to deter whitewashing.

Another idea is to punish newcomers within the reputation system itself. For example, we can allow a player to have one of three types of reputation: "good," "bad," and "new." (We're assuming that newly joined users can be identified as such, as is usually the case.) The intended behavior is for a player to defect against someone with a new or bad reputation, and to cooperate against a player with a good reputation. Any player that deviates from this intended behavior acquires a bad reputation forevermore. This is known as the *pay-your-dues* strategy [3]. Since newcomers are always defected against, if all other players are playing this strategy, then the best response is to always cooperate and never whitewash (assuming $\gamma > \frac{1}{2}$).

5.4 Other Undesirable Behavior

The ability to cheaply obtain multiple identities causes problems beyond whitewashing. A different attack is to create lots of identities, and have each of them transact with each other many times to build up good reputations. Strategies for mitigating this attack include: making it harder to obtain multiple identities (as above), imposing non-trivial transaction costs (so that building up a good reputation carries a significant cost), and working hard to identify fake accounts (based on user behavior and/or the pattern of interactions between different users).

A second problem is if users withhold feedback or submit dishonest feedback. The more benign reason that this can happen is just the standard underprovisioning of a public good (Lecture #6)—since the cost of leaving feedback (the time involved, if nothing else) is borne only by the rater, while the benefit is enjoyed by all of the users, one would expect users to contribute less than the optimal level. There are also more nefarious reasons for withheld or dishonest feedback, as the following case study makes clear.

6 The Evolution of eBay's Reputation System

eBay started in 1995, which is nearly prehistoric in Internet years. For example, guess which Web browser people were using back then? Google, let alone Chrome, didn't exist in 1995. Neither did Firefox. Netscape Navigator was quickly taking over from the Mosaic browser, and Internet Explorer was just being invented. It was also way before Paypal, so many buyers paid for items using a personal check, sent via snail mail. The point is that the Internet was the Wild West back when eBay started, and eBay never would have survived without paying careful attention to its recommendation system.

In eBay, buyers rate sellers, and vice versa. This is not true in all online marketplaces. For example, in Amazon, buyers can rate sellers and specific products, but sellers cannot rate buyers. Given that in the early days of eBay, buyers and sellers transacted directly which each other (e.g., a buyer might mail a check that would later bounce), it made sense to allow buyers to be rated. In Amazon, all payments go through Amazon's system, so it can more easily deal directly with delinquent buyers.

A persistent issue in eBay is that buyer feedback is heavily biased toward positive feedback. For example, a conservative estimate places the fraction of unsatisfactory transactions on eBay at around 3.4% [4].[9] However, much less than 1% of the feedback left by buyers is actually negative.[10] Why might this be?

One tricky design issue is whether feedback by both parties is reported sequentially or simultaneously. In the sequential case, which is what eBay has used, feedback is immediately reported to the other party. This raises the possibility of retaliation, in which a seller that receives negative feedback subsequently leaves negative feedback on the buyer. And retaliation can be much worse than just leaving negative feedback: there have been reports of threatening phone calls in the U.S. to buyers who left negative feedback, and even a defamation lawsuit in the U.K (see [6]). Several years ago, eBay changed its reputation system so that sellers can only leave positive feedback and a comment about a buyer, and can't give negative feedback.[11]

eBay has also experimented with how it computes a reputation score from a seller's history. "Percent positive (PP)" is the standard reputational measure used in eBay, which is defined as

$$\frac{\text{\# of transactions with positive feedback}}{\text{\# of transactions with any feedback}}.$$

Having a PP score of 98% sounds pretty good, right? But actually only the worst 10% of the sellers have PP scores of 98% or below.

PP scores are misleading because they ignore transactions where no feedback was given. If an unusually large fraction of a seller's transactions did not result in any feedback, then it's natural to assume that some of these transactions were unsatisfactory. This motivates the "effective percent positive (EPP)" measure:

$$\frac{\text{\# of transactions with positive feedback}}{\text{total \# of transactions by seller}}.$$

Unsurprisingly, there is a much bigger spread in sellers' EPP scores than their PP scores. eBay did not make EPP scores visible, however—in part because it wasn't clear that buyers

[9]This estimate counts transactions where the buyer left negative feedback, or a dispute ticket was opened, or where "negative-looking" email was exchanged between the buyer and seller after the transaction.

[10]In Amazon, where sellers cannot rate buyers, the percentage of negative feedback is more in the 3-4% range.

[11]An alternative approach, adopted recently by Airbnb, is simultaneous feedback. This means that feedback is not revealed to either party until both parties have submitted feedback (or after the deadline for submitting feedback has expired). It is safer to provide honest feedback with the simultaneous implementation.

would know how to interpret them, and in part because they didn't want to encourage sellers to game the measure. eBay did, however, use EPP scores in their search algorithm, promoting sellers with high EPPs more prominently than those with low EPPs. This seemed to work well, with a significantly higher fraction of first-time users returning for a second transaction after the search algorithm was changed (see [5]).

References

[1] G. A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.

[2] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets. Reasoning about a Highly Connected World*. Cambridge University Press, 2010.

[3] E. J. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics & Management Strategy*, 10(2):173–199, 2001.

[4] D. V. Masterov, U. F. Mayer, and S. Tadelis. Canary in the e-commerce coal mine: Detecting and predicting poor experiences using buyer-to-seller messages. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 81–93, 2015.

[5] C. Nosko and S. Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. NBER working paper, 2015.

[6] S. Tadelis. Reputation and feedback systems in online platform markets. *Annual Review of Economics*, 8:321–340, 2016.