Trust and reputation

in the peer-to-peer platform economy

Judith Kas

Trust and reputation in the peer-to-peer platform economy

Judith Kas

Manuscript committee:

prof. dr. ir. V.W. Buskens (Universiteit Utrecht)
prof. dr. B. Derks (Universiteit Utrecht)
prof. dr. S. Rosenkranz (Universiteit Utrecht)
prof. dr. T. Teubner (Technische Universität Berlin)
prof. dr. R.P.M. Wittek (Rijksuniversiteit Groningen)

Trust and reputation in the peer-to-peer platform economy – Judith Kas

Cover design: Roel Stoltenborg (RS Vormgeving)
Printing: Ridderprint
ISBN: 978-94-6416-308-7

Trust and reputation in the peer-to-peer platform economy

Vertrouwen en reputatie in de peer-to-peer platformeconomie

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de rector magnificus, prof.dr. H.R.B.M. Kummeling, ingevolge
het besluit van het college voor promoties in het openbaar te verdedigen

op vrijdag 12 februari 2021 des middags te 2.30 uur

door

Judith Kas

geboren op 22 september 1992
te Wageningen

Promotor:
Prof. dr. A. van de Rijt

Copromotor:
Dr. R. Corten

Dit proefschrift werd (mede) mogelijk gemaakt met financiële steun van de Nederlandse Organisatie voor Wetenschappelijk Onderzoek (projectnummer 452-16-002). Het Swaantje Mondt Fonds heeft financieel bijgedragen aan een bezoek aan Stanford University in 2019. Dit proefschrift is afgedrukt met financiële steun van de J.E. Jurriaanse Stichting.

Contents

CHAPTER 1

Synthesis

Trusting a perfect stranger

In the spring of 2019 I was in the US for a research visit as a part of my PhD. I used the peer-to-peer hospitality platform Couchsurfing to find others who would like to go on weekend trips to the national parks around San Francisco. One of the advertisements I found was titled 'Grand Canyon, Hiking 3 Days'. The description of the event explained that the author of the ad would go to the Grand Canyon by motorcycle (which takes more than ten hours), and that he would have space for one more person on his motorbike and in his two-person tent.

This ad nicely illustrates how important interpersonal trust is in the peer-to-peer platform economy. *Platform economy* is an umbrella term that encompasses several activities, such as selling, exchanging, borrowing and renting of goods and services. It is difficult to measure the size and impact of the platform economy, but estimates suggest that the platform economy is large and rapidly increasing (McKinsey, 2016; Roof, 2019). Compared to traditional ecommerce, buyers and renters in the new platform economy are trading with individual sellers rather than with businesses. This leads to increased risks, because strangers often interact without the prospect of future interactions (Kuwabara, 2015; Parigi, Santana, & Cook, 2017), and because these platforms are characterized by lower levels of institutionalization (Katz, 2015; ter Huurne, Ronteltap, Corten, & Buskens, 2017).

Although most platform users will have good intentions, there is always a chance that you meet someone who may do harm to you. If I would go on a trip to the Grand Canyon, I would be concerned about my personal safety, because I do not know how well my travel partner can drive a motorcycle and because sleeping in a small tent in the middle of nowhere would make me quite vulnerable. Of course, the same applies to him: since he did not know me, he did not know about my intentions, so I could be a threat to his personal safety as well. On many platforms the individuals who rent out expensive goods, such as cars and motorcycles, run the risk that the renter does not return the good, or does not return it in a good condition. The problem here is that platform users cannot distinguish between trustworthy others, with whom they can safely interact, and untrustworthy individuals, who may abuse their trust (Macy & Skvoretz, 1998).

Trust problems are not unique to exchanges in the platform economy. They are common to everyday life, and especially to economic interactions. There is a rich literature on solving these problems (Akerlof, 1970; Diekmann, Jann, Przepiorka, & Wehrli, 2014; Kreps, 1996). In this dissertation I use the platform economy as a case to study this classical problem in a real-life case. Doing so both deepens our understanding of the theoretical trust problem and provides insight in solutions to the trust problem in the specific case of the platform economy.

Platforms try to alleviate the trust problem in a number of ways, of which reputation systems are one of the most common (Cook, 2005). Reputation systems collect, aggregate and distribute feedback about individual's past decisions on the platform (Resnick, Kuwabara, Zeckhauser, & Friedman, 2000). According to classical theories on reputation systems, people are more likely to interact with another individual if that person has been trustworthy in the past (Buskens, 2003; Buskens & Raub, 2002; Cook, Hardin, & Levi, 2005; Weigelt & Camerer, 1988). Reviews on Couchsurfing for example contain phrases like 'This host is a reliable person'. These reviews are provided by other Couchsurfing users who joined the host on a trip in the past. The positive effect of past trustworthiness on trust is well established in the existing literature (Boero, Bravo, Castellani, & Squazzoni, 2009; Bolton, Katok, & Ockenfels, 2004a; Charness, Du, & Yang, 2011; Duffy, Xie, & Lee, 2013; Fehrler & Przepiorka, 2013; Vincenz Frey & Van de Rijt, 2016).

This approach to studying the effectiveness of reputation systems is rigorous, but narrow. In this dissertation I challenge this view by studying three new ideas of how reputation systems affect behavior. These ideas are inspired by three typical, but not unique characteristics of the platform economy.

First, when deciding whether I would join the Couchsurfing trip, I did not only have access to reviews. The organizer of the trip also provided some pictures of himself, and his personal profile contained his first name. This emphasis on personal information is commonplace in the platform economy and is meant to enhance trust (Dubois, Willinger, & Blayac, 2012; Guttentag, 2015). However, earlier research has shown that the use of personal photos and real names in these profiles enables discrimination (Edelman & Luca, 2014; Edelman, Luca, & Svirsky, 2017). There is for example some evidence that

Couchsurfers prefer to stay with other females or with families rather than with male hosts (Tan, 2010). In the case of the Couchsurfing trip I could be more concerned about my personal safety because the organizer of the trip was male. Other examples of discrimination in the platform economy are ethnic discrimination (e.g. Edelman & Luca, 2014) and discrimination based on people's sexual orientation (Ahuja & Lyons, 2019). People generally prefer to interact with others who are similar to them (Balliet, Wu, & De Dreu, 2014; Currarine & Mengel, 2016; Tajfel, Billig, Bundy, & Flament, 1971). This is not only a problem for individuals who suffer from negative discrimination, but also for platforms and for the potential of the platform economy to serve larger audiences.

The first question that is addressed in this dissertation, is therefore whether reputation systems can help solve discrimination. Some authors have argued that reputation systems reduce the reliance on demographic information and are therefore a solution to discrimination in the platform economy (Abrahao, Parigi, Gupta, & Cook, 2017; Cui, Li, & Zhang, 2017; Ert, Fleischer, & Magen, 2016; Mohammed, 2017; Tjaden, Schwemmer, & Khadjavi, 2018). Applied to the context of the Couchsurfing trip, I could be less concerned about the gender of the organizer of the trip if he had some positive reviews that proved that he could be trusted. By doing so, the reputation system may reduce the extent to which I relied on demographic characteristics, and thus reduce differences in opportunities to participate between people with different demographic characteristics.

However, this is not the whole story. The process of reputation building is not exogenous: reviews are only written after completed interactions. Since not everybody is equally likely to participate in a first interaction (Cui et al., 2017; Edelman et al., 2017), there may be differences in the extent to which individuals with different characteristics can accumulate reviews, which in turn affects their future chances to participate on the platform. This may imply that it is more difficult for male Couchsurfers to participate in interactions, and thus to collect reviews. Through this mechanism, reputation systems may reinforce rather than reduce inequality. In chapters 2 and 3 we studied how reputation systems affect inequality. We propose that the reputation formation process is not bias free, and that it is therefore not obvious that reputation systems help reducing inequality.

Second, the presence of personal profiles does not only lead to discrimination based on names and pictures. Most peer-to-peer platforms allow their users to communicate before starting a transaction, both through their profiles and product descriptions and through the chat function. In the example of the Couchsurfing trip the organizer wrote that he was a good driver. Earlier research shows that communication by trustees significantly increases buyers' willingness to pay for a product (Anderson, Friedman, Milam, & Singh, 2007; Lewis, 2011; Rawlins & Johnson, 2007). However, these promises should not affect decision making under the assumption of rationality and selfishness, since there is no way to know if the seller speaks the truth. Any Couchsurfer can claim to be a good driver and there is no way to check if that is true. Reputation system may help here: The reviews of the author of the Couchsurfing ad may not only say something about the absolute quality of the interaction ('he's a good driver'), but also about the extent to which the author has spoken the truth in the past ('he said he would adhere to the speed limits, but he drove way too fast'). I could have used that information to assess if the author was speaking the truth when he told me he was a good driver. In chapter 4 we study if reputation systems may help to make this type of 'cheap talk' credible. We contribute to the literature by providing an elaborate discussion of the theoretical mechanisms leading to this prediction.

Third, reputation systems have mostly been studied in cases where the roles of providers and consumers are assumed to be strictly separated, while this assumption does not hold true in the context of the platform economy. The peer-to-peer nature of online platforms allows for individuals to act in more than one role. Relaxing the assumption that individuals act in only one role allows for the emergence of a mixed-role reputation system, in which an individual's reputation is based on reviews collected in both roles. The reputation of the Couchsurfer who wanted to go to the Grand Canyon could contain both reviews that he received as a host and reviews he received as a guest. This new reputation system changes the way trust is formed, because it provides more and different information, and because it allows for different ways of strategic reputation building. The last question that I address in this dissertation (chapter 5), is whether this new type of reputation system helps to increase trust. We deviate from prior research by proposing that reputation systems do not only increase trust by providing information about a trustee's

behavior in the current role, but also by informing trustors about a trustee's past behavior in different roles.

In sum, the goal of this dissertation is to study how reputation systems affect trust in the context of peer-to-peer platforms, where individuals fulfill multiple roles and where personal profiles are omnipresent. This research provides insight in the way platforms solve trust problems, and thereby increases our understanding of trust problems in general, and reputation systems specifically. I challenge the classical theory of reputation systems that reputation systems affect trust by providing information about a trustee's past trustworthiness by discussing three new mechanisms through which reputation systems affect trust. In the remainder of this chapter I will give a brief overview of the literature on trust, reputation systems and the platform economy. Then I briefly discuss the results of the four studies, before moving on to a general discussion.

Previous work on trust and reputation systems

Trust is commonly defined as the willingness to accept vulnerability due to others' actions based on expectation about their intentions and skills (Mayer, Davis, & Schoorman, 1995; Rousseau, Sitkin, Burt, & Camerer, 1998). Users of platforms run a risk, because the rules and regulations that apply to individuals are less clear than for companies (Katz, 2015; ter Huurne et al., 2017), and because individuals are generally involved in fewer interactions than businesses. In the example of the Couchsurfing trip, I had no guarantees that the trip would take place and that the quality would be as promised. Moreover, there is a lot of heterogeneity in the goods and services offered on these platforms, so even though consumers are protected against large losses by insurances and secure payment systems, the quality of their purchases may still be disappointing. The same holds for providers: they may be covered against theft and observable harm done to their products, but wear and tear and other damage that cannot (directly) be observed in the short term are not covered by insurances. Many of these online exchanges are followed by physical meetings between providers and consumers, which may cause a personal safety thread (Ranchordas, 2015). This is very apparent in the Couchsurfing trip to the Grand Canyon. The problem is not that all platform users are selfish, but rather that

the costs of occasionally encountering an untrustworthy user may be many times higher than the benefits of a successful interaction.

Describing the trust problem as a Trust Game

To better understand and theorize on trust formation on these platforms, interactions on platforms can be described as a Trust Game (Figure 1.1), which is a simplified, abstract version of the trust problem in these markets (Coleman, 1994; Corten, 2019; Weigelt & Camerer, 1988). The left panel of Figure 1.1 shows the Trust Game. The right panel provides a numerical example. In the basic form of the Trust Game, a trustor and a trustee participate in an anonymous one-shot interaction in which they sequentially take a decision.

I will use the example of the Couchsurfing trip to explain the Trust Game. In the example I could choose to join the trip if I trusted the organizer, or I could choose not to join. If I decided to stay at home, there would be no trip for me, and the organizer would not have a travel partner. The game would end in that case and we would both receive a payoff of P_i. If I would decide to join, the organizer of the trip could honor my trust by driving safely and by respecting my personal space. In that case we would both receive a payoff R_i, which is higher than if no transaction would have taken place: $R_i > P_i$. The organizer could also abuse my trust. The costs of this abuse for me may quickly exceed the advantages of joining the trip[1]. If the organizer abuses my trust, my net payoff S_1 will be lower than if I had refrained from joining the trip: $P_1 > S_1$. The organizer of the trip however may benefit from deviant behavior, so his payoff T_2 of abusing trust may exceed his payoff from honoring trust: $T_2 > R_2$. This means that the trustor can maximize their payoff by placing trust when the trustee honors trust. However, because the payoff of the trustee is higher when they abuse trust, rational, selfish trustees are expected to always abuse trust. The best response of the trustor to that expectation is to not place trust, which is a sub-optimal outcome, since both players could have received a higher payoff if the trustor placed trust and the trustee honored trust.

[1] Platforms have taken several measures that protect trustors, such as insurances (ter Huurne et al., 2017). However, misbehavior by trustees can lead to damage that is not easy to identify, and even if the owner if financially compensated, it may not outweigh the hassle.

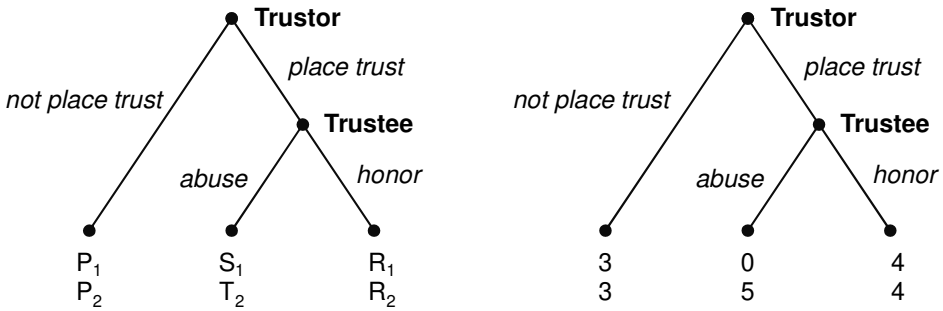

Figure 1.1: left panel: one round of the Trust Game without reputation system, where $R_1 > P_1 > S_1$, $T_2 > R_2 > P_2$ and $P_1 + P_2 > S_1 + T_2$. Right panel: numerical example of the Trust Game.

The Trust Game is a sequential game in which only the trustor runs a risk, while only the trustee can choose to behave opportunistically (Coleman, 1994; Dasgupta, 1988; Kreps, 1996). At first sight this may not seem very applicable to real platforms, because most interactions involve a risk for both parties. A Prisoner's Dilemma is an example of a game that models two-sided dilemmas, in this game the dominant strategy of players is always to defect (Kollock, 1998). In platform economy interactions, that is not necessarily true. Both providers and consumers can benefit from an interaction, provided that the other is trustworthy. While this is never a dominant strategy in a Prisoner's Dilemma, Trust Games model these strategic interdependencies between actions of different actors. Moreover, rather than viewing interactions on these platforms as a single Prisoner's Dilemma, these interactions can be viewed as a sequence of one-sided Trust Games, in which sometimes the provider runs a risk and sometimes the consumer. This simplification of real-life platforms helps to analyze the decision setting in detail and to distill the effects I am interested in.

Reputation systems are commonly used to alleviate the trust problem on online platforms. Existing theories on reputation systems usually assume that reputation systems increase trust by providing information about the past trustworthiness of a trustee. Trustors can use this information to distinguish between untrustworthy and trustworthy trustees (we call this 'learning'), and they provide an incentive for trustees to be trustworthy (trustors have control over the behavior of the trustee, Buskens, 2003; Cook, Hardin, & Levi, 2005; Weigelt & Camerer, 1988)

My theoretical contributions

Classical theories about reputation systems focus on the effect of an individual's history of trustworthiness on trust. The emergence of the platform economy has highlighted several characteristics of reputation systems that have been previously overlooked and that may affect the mechanisms through which reputation systems influence trust. Although the use of personal profiles, communication and the fluidity of the division of roles stand out in the context of the platform economy, they are not unique to this context. Information about providers and consumers, communication and reciprocal relations also affect interactions in more traditional, offline contexts. Studying reputation systems in the context of the platform economy thus raises three new questions and teaches us more about classical trust problems: what information do these reviews and ratings contain and how does this information affect behavior?

The role of reputation systems in digital discrimination

The first problem I address in this dissertation is to what extent reputation systems can help solving digital discrimination. Many platforms allow their users to view personal profiles of potential interaction partners before deciding to engage in an interaction. The aim of this information is to create a sense of personal contact and provide information about potential transaction partners that may help trustors make a choice between seemingly trustworthy and untrustworthy trustees, and between high-quality and low-quality products (Dubois et al., 2012; Guttentag, 2015).

The use of names and photos on platforms has been found to facilitate (ethnic) discrimination, which is a problem on many platforms. In the context of the Couchsurfing trip described in the introduction, there is some evidence that female guests are more likely to be trusted on Couchsurfing (Tan, 2010). Besides gender discrimination, ethnic discrimination is an issue on many platforms. Non-black hosts on Airbnb can for example charge higher prices than black hosts in the US (Edelman & Luca, 2014), and native German drivers on a German carpooling platform attract more attention than drivers with an Arab/Turkish/Persian name (Tjaden et al., 2018). Discrimination may be disadvantageous for both platform users and for the platforms, because they are missing out on potentially fruitful transactions.

In earlier studies on digital discrimination it is argued that reputation systems may provide a solution to this problem, because they reduce the reliance on demographic indicators (Abrahao et al., 2017; Cui et al., 2017; Ert et al., 2016; Mohammed, 2017; Tjaden et al., 2018). It had been argued that demographic characteristics are generally weak and diffuse indicators of unobservable characteristics, such as trustworthiness, while ratings and reviews are context-specific and therefore a more reliable indicator of these characteristics. Previous research indeed shows that trustors rely less on demographic information when a trustee has at least one positive review (Abrahao et al., 2017; Cui et al., 2017; Tjaden et al., 2018)

In chapters 2 and 3, we argue that earlier research has overstated the extent to which reputation systems mitigate discrimination. On most platforms, reviews can only be written after a completed transaction. We know that not all trustees are equally likely to participate in a transaction, and therefore that they are not equally likely to obtain a first review. We also know that especially that first review is very important for future chances to participate, since reputation tends to cascade: as soon as an individual has one positive review, it is more likely that (s)he will be selected for future exchanges, which in turn results in more reviews (Duffy et al., 2013; Frey & Van de Rijt, 2016). Data on the distribution of reviews over Airbnb hosts is indeed highly skewed, with a few hosts that have a large number of reviews, and many hosts that have no reviews (Teubner & Glaser, 2018). In chapters 2 and 3, we tested if reputation systems are indeed able to reduce inequality between different ethnic groups.

Reputation systems and cheap talk

The second feature of platforms that may affect the way reputation systems work is the use of written text in personal profiles and other forms of communication before an interaction. Users can use these texts to make promises about the product they have on offer, or about themselves or their own behavior during the interaction. The organizer of the Couchsurfing trip for example wrote that he is a skilled motor rider. Previous research shows that this type of information significantly affects trust when consumers determine prices, such as in auctions (Anderson et al., 2007; Lewis, 2011; Rawlins & Johnson, 2007). However, theories that assume that individuals are rational and selfish would predict that this information does not affect decision making,

since there is no guarantee that the author will keep their promises (Akerlof, 1970).

A few papers have hinted at the idea that reputation systems may 'make cheap talk credible' (Ba & Pavlou, 2002; Diekmann et al., 2014; Houser & Wooders, 2006; Lewis, 2011; Neto, Bloemhof, & Corbett, 2016; Snijders & Zijdeman, 2004). We are the first to provide an elaborate discussion of the theoretical mechanisms underlying this claim. We hypothesized that when buyers determine prices and when there is not only heterogeneity between sellers in the quality of products on offer, but also variation within sellers over time, reputation systems may provide an incentive for sellers to communicate honestly about the quality of the product they have on offer. Under these conditions, sellers may be able to maximize their long-term payoffs by providing a high-quality product whenever possible, and by communicating honestly when they only have products of inferior quality. By doing so, they may keep up a reputation for being trustworthy and honest, which may be beneficial in the long run. In chapter 4, we tested if the presence of a reputation system may indeed make this cheap talk costly within the context of online auctions.

Reputation systems in mixed-role markets

A third characteristics of peer-to-peer platforms is that the distinction between providers and consumers is less clear than in traditional e-commerce. Many users act in the role of consumer in one exchange, while they act in the role of provider in other exchanges (Cova, Dalli, & Zwick, 2011; Ritzer & Jurgenson, 2010). Couchsurfing users are for example known to alternate between staying at other people's places and inviting people to their house (Lauterbach, Truong, Shah, & Adamic, 2009). In these cases where the lines between actors in different roles are blurred, a new type of reputation system emerges that contains information on the user's behavior in both the trustor and the trustee role.

Mainstream theory on reputation systems argues that information about behavior in the trustor role is at best of secondary importance. In chapter 5, we argue that when there is no financial compensation for trustors, it is likely that individuals who are more trustful are also more trustworthy. Both placing trust and honoring trust incur a cost to the self, while they increase the payoff of

another individual. Trustees who have placed more trust in the past can thus be expected to be more trustworthy, and trustors can use information about past trustfulness behavior of trustees to learn about their trustworthiness. This also implies that the presence of a mixed-role reputation system may provide an incentive for trustors to place trust. We studied these claims in chapter 5 using a laboratory experiment.

Summary of the results

Chapters 2 and 3: How reputation systems may reinforce existing inequalities

In chapters 2 and 3 we focused on the question whether reputation systems may reduce existing inequalities between individuals of different ethnicity. We used simulations to define the conditions under which reputation systems may reinforce existing inequalities, thereby sustaining the ethnic gap, or even widening it. The simulations suggest that when initial trust in individuals with an ethnic minority background is low, reputation systems may not necessarily decrease inequality between users with different demographic characteristics.

Chapter 2: The role of reputation systems in digital discrimination: a field study

We tested the hypothesis that reputation systems affect inequality by analyzing the complete historical records of user activity on a Dutch peer-to-peer motorcycle sharing platform. We found that with time the reputation system failed to reduce the ethnic gap. Regardless of the number of requests a renter had submitted, requests from renters belonging to the ethnic majority had a larger probability of being accepted than requests from ethnic minority renters. This means that the reputation systems did not help ethnic minority renters to overcome the initial disadvantage that they have on the platform. Moreover, we did not find evidence that the presence of reputation information reduced the importance of demographic information.

Although the platform data allowed for a first test of the theory with a high external validity, there were several limitations related to the method used. First, the theory is about the comparison between a platform with and without a

reputation system. Instead of comparing two platforms, one with a reputation system and one without a reputation system, we explored how the effectiveness of the reputation system changed over time. Another limitation of the dataset was the small fraction of individuals belonging to the ethnic minority, and especially the number of ethnic minority members with positive reviews. Although this is a direct implication of the relation between ethnicity and reputation, it limited the opportunity to study the interplay between these characteristics. The platform data also provided information about a single case and only one sequence of reputation building, there was no 'resetting' of the reputation system. A last limitation of the dataset was that the decision to submit another request was conditional on outcomes of earlier requests, which distorted interpretation of the results.

Chapter 3: The role of reputation systems in digital discrimination: an experiment

I addressed these limitations in chapter 3 using an online experiment that is inspired by the experiment of Frey and Van De Rijt (2016) on reputation cascades in competitive markets. Using an experiment allowed me to directly compare a situation with and a situation without a reputation system, and to draw causal conclusions about the effect of reputation on demographic inequality based on multiple sequences of reputation accumulation. In the experiment American and Indian trustors took turns in selecting a trustee from two American and two Indian trustees. When there was a reputation system, they were informed about previous decisions of each trustee to honor or abuse trust. American trustors were generally likely to select an American trustee.

The results from the experiment corroborated the finding from chapter 2 that reputation systems fail to overcome differences in the opportunity to participate in platform exchanges between individuals with different demographic backgrounds. Inequality between American and Indian trustees was not smaller when there was a reputation system than when there was no reputation system. Moreover, I found that decisions made in earlier rounds of the game strongly affected decisions made in later rounds. If in early rounds of the experiment an Indian trustee was selected, chances were high that trustors in later rounds also selected an Indian trustee. Again, I did not find evidence

for the 'compensation effect': the presence of reputation information did not decrease the reliance on demographic information.

The results from chapters 2 and 3 suggest that reputation systems do not reduce existing inequalities. Biased decisions that are made earlier in time affect who gets to accumulate reviews, and thus who will be trusted in the future.

Chapter 4: Reputation systems, trust and the identification of liars

In chapter 4 we tested the hypothesis that communication in heterogeneous markets between trustors and trustees may be credible because of the presence of a reputation system. We used data on eBay auctions of second-hand cars collected by Lewis (2019). The analyses provide only weak support for the hypothesis that bidders in auctions respond stronger to claims made by more reputable sellers. To exclude the possibility that the relation between promises made and the effect of reputation is spurious and to test the effect of the mere presence of a reputation system, we conducted a second study.

We used an online experiment in which trustors (i.e. buyers) and trustees (i.e. sellers) played one of two versions of a Trust Game. The first version is a regular Trust Game (see Figure 1.1), while in the other version of the game the trustee does not have the opportunity to honor trust when the trustor places trust. We varied whether trustees could send a binary message to the trustor before the trustor decided to place trust or not. We also varied whether buyers were informed about previous decisions made by sellers.

The results from the experiment show that trustee behavior followed the predicted patterns. Trustors could maximize their payoff by only placing trust in trustees who promised to honor trust and who had been trustworthy and honest in the past. However, trustors in the experiment did not follow this optimal strategy. They relied on messages sent by the trustee even when there was no way to check if this trustee spoke the truth. Based on the two studies we conclude that the presence of a reputation system cannot explain why platform trustors rely on non-binding communication by trustees.

Chapter 5: Reputations in mixed-role markets

In chapter 5, we analyzed whether trustors behave differently in the presence of a mixed-role reputation system. Such a reputation system does not only include information on the trustee's behavior in the trustee role, but also information on the trustee's behavior in the trustor role. We expected that when trustors are not financially compensated for their trust ('lending'), more trust would be placed when there is a mixed-role reputation system than when there is a single-role reputation system. When trustors can increase their own payoff by placing trust ('letting'), placing trust cannot necessarily be viewed as an altruistic act, so placing trust is a weaker signal of trustworthiness in these cases. We hypothesized that the effect of mixed-role reputation would be weaker when trustors are compensated for being trustful than when they are not compensated.

We tested this hypothesis using a laboratory experiment in which we introduced a new game in which trustors were not compensated for being trustful: The Lending Game. We compared the effect of a single-role reputation system and a mixed-role reputation system in a regular Trust Game and in the Lending Game. We found that, indeed, mixed-role reputation systems increased trust levels, and only for scenarios of lending, not letting. We also found evidence that providers are more likely to place trust in users who themselves have placed more trust in the past. We expected the effect of the mixed-role reputation system in paid exchange to be smaller, and in practice we found no effect at all. These findings suggest that in modern day platforms where the roles of consumers and providers are mixed, individuals can develop their reputation in different ways: by showing they can be trusted and by showing they are willing to trust others.

General discussion and conclusion

When deciding if I wanted to join the Couchsurfing trip, I could read the reviews written by the past travel partners of my potential travel companion. According to classical theory on reputation systems, these reviews would increase trust by informing me about the extent to which the author of the advertisement had been trustworthy in the past (Buskens, 2003; Buskens &

Raub, 2002; Cook et al., 2005; Weigelt & Camerer, 1988). The emergence of the platform economy has highlighted three characteristics of the context in which trust problems are studied that affect the mechanisms through which reputation systems work. Although the emergence of the platform economy drew our attention to these issues, these issues apply to many more situations in which trust plays a role. The goal of this dissertation was to study how these characteristics affect the mechanisms through with reputation systems increase trust.

A first characteristic that affects the effectiveness of reputation systems is personal information that is provided about potential interaction partners. Ethnic-, gender- and age discrimination may be based on names and appearance of potential interaction partners. In chapters 2 and 3 we showed that reputation systems do not necessary reduce inequality between (ethnic) groups, contrary to what concluded in earlier research (Abrahao et al., 2017; Cui et al., 2017; Ert et al., 2016; Mohammed, 2017; Tjaden et al., 2018). This implies that the introduction of a reputation system may not lead to a reduction in discrimination. Applied to the example of the Couchsurfing trip, this would mean that the preference for female users would not be smaller for users with many reviews, and that males do not get the opportunity to collect enough reviews to overcome the difference.

In chapter 4 we tested if reputation systems may provide an incentive for trustees to make honest promises. We did not find evidence for that hypothesis. Trustors relied on quality claims made by trustees, even when these trustees did not have a reputation for honesty. In the context of the Couchsurfing trip this would mean that I would trust the organizer to be a good driver, even if the reputation systems did not allow me to find out if he generally tended to speak the truth.

Besides personal profiles, the effect of reputation systems may also be affected by fading boundaries between consumers and providers. This mixing of roles allows for the emergence of a new type of reputation system in which information about both roles is displayed. In chapter 5 we found that this mixed-role reputation system positively affects trust formation on platforms where providers do not benefit from providing their goods, but that this mechanism ceases to work on paid platforms. We propose that in addition to

sheer goodness, mixed-role reputation systems play an important role in facilitating free exchange.

Theoretical implications

Reputation systems may reinforce existing differences between (groups of) people

Earlier studies on the effect of reputation systems on inequality concluded that reputation systems help reducing discrimination (Abrahao et al., 2017; Cui et al., 2017; Ert et al., 2016; Mohammed, 2017; Tjaden et al., 2018). This conclusion was based on the finding that the attention for demographic characteristics, such as ethnicity, decreased once a trustee had at least one positive review. This conclusion implicitly assumed that reputation scores are exogenously generated. In chapter 2 and 3 we showed that reputation systems are not free from bias. Reputation systems are not simply a tool that help trustors decide. They are not objective measures of quality: it is not necessarily the case that those individuals that are most trustworthy will accumulate the most positive reviews (Frey & Van De Rijt, 2016).

Rather, reputation building is a sequential process: decisions made early on affect decisions made later. Individuals who have a larger likelihood of being trusted, have a larger chance of acquiring a first review. That first review will help them to be trusted more, which will in turn help them accumulate reviews. This means that reputation systems may reinforce initial differences between individuals. In the best scenario, these initial differences are random, and the resulting inequality is also randomly distributed among individuals and groups (Frey & Van De Rijt, 2016). In a more realistic scenario, initial differences are based on structural differences between individuals, such as their demographic characteristics. Future research on reputation systems should therefore account for this dynamic nature of reputation systems.

Trustors may improve their decision by using different types of reputation information

According to classical theory, reputation systems convey information about the trustee's past trustworthiness (Buskens, 2003; Buskens & Raub, 2002; Cook et al., 2005; Weigelt & Camerer, 1988). Trustors can use that information to assess the probability that a trustee will be trustworthy in the future. In chapters 4 and

5 we showed that, contrary to what previous research has focused on, reputation systems do not only include information about absolute quality of past interactions. What information is included in the reputation system depends on the context in which the reputation systems is implemented. In both chapters we showed that trustors could use different types of reputation information to improve their estimate of the probability that a trustee will be trustworthy. We extended the knowledge about the functioning of reputation systems by systematically varying the type of information in the system and by studying the effect of the different information sources on trust.

In chapter 4 we studied the effect of reputation systems in markets where trustors and trustees could communicate before a transaction. The experimental results show that trustors could improve their decisions by using a combination of reputation information and promises made by trustees. Interestingly, trustors in both the field study and the experiment did not optimally use the available information: they always placed more trust in the trustees who promised to honor trust, even if the trustee did not have a good reputation. In chapter 5 we studied reputation systems in markets where individuals participate in interactions in multiple roles. We found that trustors used information about a trustee's past behavior in a different role when deciding to place trust or not.

These findings raise the question why trustors sometimes show the predicted behavior and sometimes not, since trustees generally behave as expected, and since both theories tested in chapter 4 and 5 relied on assumptions of rationality and selfishness. One explanation for why trustors did not always make an optimal decision, is that their bounded rationality limits their ability to fully comprehend the consequences of their choices (Buskens, Raub, & van der Veer, 2010; Kahneman, 2002; Simon, 1955). Knowing why trustees who have placed more trust in the past will be more trustworthy in the future, is cognitively less demanding than understanding how information about a trustee's past behavior and current promises can be combined to reach the best decision.

Another explanation is that trustors are not selfish, or that they believe that trustees are not selfish. If trustors believe that the trustees in chapter 4 generally spoke the truth, there would be no reason to distrust them. Rational trustors

should learn over time that not all trustees speak the truth and adjust their behavior accordingly. I indeed observed a downward trend in trust in most of the experiments, but this can also be explained by a shorter shadow of the future. (Classical) theories on reputation systems mostly rely on assumptions of rationality and selfishness. Relaxing these assumptions may pave the way for theories that rely on different mechanisms, such as reciprocity and the use of heuristics. Future research may investigate if the limited rationality or selfishness of trustors explains the difference in the results between the studies. An example of such a study would be an experiment in which trustors are informed about which decision is expected to yield them the highest payoff. The goal of that study would be to rule out limited rationality as an explanation for why trustors do not optimally use the available information.

Information in reputation system affects strategic reputation building

The presence and reliance on different types of information implies that individuals can build their reputation in different ways. The findings from chapters 4 and 5 show that trustees cannot only improve their reputation by being trustworthy, but also by communicating honestly and by placing trust in others. The increased opportunities for strategic reputation building may have consequences beyond building trust.

First, the presence of incentives for cooperative behavior may weaken the correlation between people's intrinsic motivation and their behavior. Without a reputation system, only individuals who are intrinsically motivated to behave cooperatively will do so. The presence of a reputation system may provide an incentive to behave cooperatively for those who would not otherwise do that. For example, in chapter 5 we found that trustors are more likely to place trust when future trustors will be informed about that. This implies that placing trust is not a signal of an intrinsic motivation to behave cooperatively anymore. This may in turn limit the usefulness of reputation information, especially if the incentives to cooperate are not constant. Given that reputation systems contain more information than only about a trustee's past trustworthiness, the presence of a reputation system may induce strategic reputation building not only by being more trustworthy, but also by being more honest and more trustful.

I will explain this using the example of the platform Peerby, where individuals can borrow goods from people in their neighborhood. If their reputation

system only contained information about the number of goods users have borrowed in the past, users would only lend goods if they wanted to help others. However, the reputation system on Peerby also includes information about the number of goods a user has lent out to others. If a user expects to need to borrow e.g. a power drill in the future, they might be concerned about keeping up a good reputation. With the current reputation system, they can do so by lending goods to others. The reputation system may extrinsically motivate people to lend out goods that they would otherwise not lend out. This means that lending out goods is not an indication of sheer goodness anymore. A history of lending may indicate that a user wants to help others, but also that a user believes that lending out goods will give them the credit to borrow goods in the future. The incentive to lend out goods created by the reputation system thus undermines the informativeness of the reputation system when trustors aim to distinguish between intrinsically motivated altruists and extrinsically motivated opportunists.

A related issue is that the presence of a reputation system may crowd out intrinsic motivation to cooperate. I hypothesize that the presence of a reputation system may change the motivation of individuals that would cooperate in the absence of a reputation system, such that they are less motivated to cooperate when the reputation system disappears, or even when the reputation system is still present. There are numerous examples of studies that show that the presence of external incentives may crowd out intrinsic motivation (Fehr & Gächter, 2005; Titmuss, 1970). In this dissertation I did not investigate if this also applies to the presence of reputation systems. Future research could investigate if reputation systems indeed crowd out motivations of platforms users to be trustworthy, honest and trustful, and to what extent there are spillover effects to the broader society.

The influence of reviews varies across contexts

In all four chapters we found that trustees with a better reputation are trusted more. However, the extent to which each review increases trust varies greatly across the studies. On the one hand, a single review or even a few reviews may not be enough to convince trustors that the trustee can be trusted. This argument seems to be mostly applicable when there is a large trust problem because the goods that are exchanged are expensive, e.g. when renting out

motorcycles (chapter 2) or buying cars (chapter 4). On the other hand, reviews may substitute each other. Every extra review adds less information than the previous review. Previous studies suggest that a first review is indeed extremely important for a trustee's future chances on a platform (Cui et al., 2017; Frey & Van De Rijt, 2016).

The empirical data in the different chapters provide mixed support for the two mechanisms. In chapter 2 renter reputation did not affect the probability that owners trusted them, unless the reputation contained more than three positive reviews. This seems to support the idea that one review is not enough to convince a trustor that a trustee will be trustworthy. However, the results from the eBay study (chapter 4) show the opposite: the first reviews had a larger effect than later reviews on the probability that the auctioned car was sold. This supports the idea that reviews are substitutable. In the experimental studies the relationship between the trustee's reputation and the decision of the trustor to place trust was more or less linear with a strong end-game effect. It may be possible that the duration of the experiments was too short to observe any of the trends seen in the field studies.

There may be several explanations for these differences. First, the ease of getting a (first) review depends on the context. If it easier to get a review, e.g. because the temptation to abuse trust is not large, the significance of a single review is limited. Second, the content of the reviews differs. On some platforms only numerical ratings are provided, while some platforms allow users to write extensive reviews. The level of detail of the reviews may affect the extent to which a single review is valuable. If the reviews are not very detailed, trustors may need more of them in order to make an informed choice. Third, if the reputation scores on a platform are highly left-skewed, as is generally the case in the platform economy (Teubner, Hawlitschek, & Dann, 2017; Zervas, Proserpio, & Byers, 2015), a single review may not be very informative. There is little variation in the numeric ratings, so a single review may not be enough to convince trustors that a trustee is trustworthy. Fourth, on most platforms trustors have access to multiple sources of information about the trustee. Trustees generally provide profile pictures and write something about themselves and in some cases they have also linked social media accounts to their profile on the platform. If trustors have more information about a trustee, the effect of a single review is expected to be limited. Fifth, the risk for the

trustor of placing trust may differ between platforms. The larger the risk, the more careful the trustor will be when making the decision, and the more information they would like to consider before making a decision. Future research could focus on the question how these contextual factors affect the extent to which every review influences the trustor's decision.

Methodological implications

This dissertation relies on a combination of abstract experimental studies and data generated on real platforms. To test the hypotheses about effects of reputation systems, data at the individual level are necessary. To explain how reputation systems affect (groups of) individual users, one needs to understand the composition of the user base of a platform as well as existing stereotypes, the size of the trust problem and other sources of information available to trustors. Ideally, different institutional designs would be studied on real platforms. However, since real platforms either have or do not have a reputation system, it is not possible to compare different institutional designs within a platform, unless a platform is willing to vary their institutional design across different groups of users. Alternatively, one could acquire data from a large number of platforms that differ in multiple dimensions to study the effect each of these dimensions have on trust.

Due to its digital nature, real-life platforms are potentially an important source of relevant user-generated data that could be used to answer these questions. In chapters 2 and 4 we used data from eBay and from a motorcycle sharing platform. These data allowed me to study the effect of reputation systems on individual trust decisions while taking the context of the platform into account.

There are several drawbacks to using field data. First, many platforms are not willing, able or allowed to share data in a format that can be used for scientific research. For example, large platforms such as Airbnb are very selectively sharing data. The data that they share may not be representative of the average activity on the platform (Cox & Slee, 2016). Another issue is that the data may not be suitable to answer the research question, because platforms do not always save and store data in ways that are useful for research. A third reason why platforms do not share their data is that the General Data Protection Regulation (GDPR) limits data sharing between platforms and third parties, including researchers.

The limited extent to which platforms share data is not the only drawback of relying on user-generated data. Another issue is that the extent to which these data can be used to answer the research questions is limited. Because real-life platforms either have or do not have a reputation system, it is not possible to study the effect of the presence of (different types of) reputation systems within a platform, unless a platform is willing to implement different types of reputation systems for different groups of users.

The data from the platforms are therefore complemented with online and laboratory experiments in which we studied a large number of abstract representations of real-life platforms. These studies allowed me to make clear causal inferences and to vary certain aspects of the institutional design, while keeping everything else constant. The positive effect of reputation on trust was not only found on real platforms, but also in game theoretic experiments using Trust Games. The similarity between the results of the experimental studies and the studies using empirical data suggests that experimental using Trust Games are indeed suitable to study the effect of reputation systems on trust on peer-to-peer platforms.

Practical implications

Reputation systems affect inequality

From the findings of chapters 2 and 3 follows that platforms that wish to reduce discrimination on their platforms should not only focus on making their reputation systems more effective. They should also consider making demographic information less visible in the initial booking process, like Airbnb recently started doing, and for which there is some evidence that it is effective (Mohammed, 2017). They should also help trustees to obtain an initial review, especially if a trustee belongs to a group that is discriminated against on the platform. A potential way to do this is by implementing a mixed-role reputation system (as discussed in chapter 5). This reputation systems allows trustees to build a reputation by helping others, rather than by behaving well when they are being helped by others. By doing so, trustees may acquire the first review necessary for future interactions, while they do not need to be trusted by others first.

In general, platforms should consider that reputation systems are an effective way of increasing interpersonal trust, but that they do not benefit every user to an equal extent. As explained previously in the discussion section, reputation systems may reinforce initial differences in the opportunity to participate on these platforms. If these initial differences are random, groups of people that share certain characteristics should not be structurally advantaged over groups of people with different characteristics. However, when initial differences are not random, but based on underlying personal characteristics, reputation systems may reinforce those differences. Platforms should consider these potential negative side-effects when implementing a reputation system. They should realize that reputation systems are not necessarily unbiased.

Platforms should consider alternative designs for reputation systems

When designing their reputation systems, platforms should consider that different types of reputation systems affect decision making in different ways, and that the reputation system does not function in isolation: they interact with other features on the platform. This means that there is no one size fits all solution to what reputation system works best on a platform. A platform should consider possible interactions with other platform features and potential side-effects before implementing or revising a reputation system.

The findings from chapters 4 and 5 suggest that a trustee's history of trustworthiness is not the only predictor of a trustee's future trustworthiness. The reputation system may also provide information about their past trustfulness and past honesty, and trustors may use this information to improve their estimate of the probability that a trustee will be trustworthy, although trustors do not always optimally use this information. Real-life platforms should consider these alternative forms of reputation building as a potentially fruitful method to increase trust.

Limitations

This dissertation has a number of limitations. First, while platforms take many measures to solve the trust problem for their users, I have solely focused on methods that reduce the information asymmetry between trustors and trustees, and not on measures that reduce the costs of interacting with an untrustworthy trustee, such as insurances and secure payment systems. These measures may

affect the extent and the mechanisms through which the reputation system affects trust. In the field studies (chapters 2 and 4) I have studied the effect of reputation systems in the presence of these cost-reducing measures, while in the experimental studies (chapters 3, 4 and 5) I have looked at reputation systems when these measures were not present. While I did not study the interaction between reputation systems and cost-reducing measures, the results from the empirical studies and the experiments were largely similar. This suggests that the results are robust to these differences between the different methods, and ultimately, between different platforms.

Second, while the experimental studies allowed me to test if trust was misplaced, it is not possible to do this in studies using field data, simply because no objective information about the behavior of the trustee is available. Reports of actual damage or abuse could give an insight here, but that would only give an indication of severe damage, and on most platforms such events are rare. In real-life, the costs of abuse are larger than in the monetized experiments. While trustors in the experiments could lose at most a few euros, trustors in in real life may lose a lot of money, or their personal safety may be threatened. Overall, trustees are probably more likely to abuse trust in a laboratory setting than in the field. This implies that the nature of the risk differs between the two situations: the probability that a trustee abuses trust is lower in the field, but the costs may be much higher. It is unclear how this difference in risks affects the effectiveness of the reputation system, but the similarity between the results from the field studies and experiments suggest that reputation systems affect trust in more or less the same way.

Third, the user base of real platforms is not stable: people enter and leave the market. The results from the field studies thus apply to those situations where some individuals enter the market later than others. However, in the experimental studies the composition of the groups was stable. New entries may especially affect the extent to which reputation systems lead to inequality between users. The process of entering and leaving a market at endogenous to a user's success in that market. We know that users with a good reputation usually stay in the market, while users who do not have a reputation, or who have a bad reputation tend to leave the market (Norbutas, Ruiter, & Corten, 2020a; Teubner & Glaser, 2018). It can thus be expected that the relative advantage of established sellers is even larger in markets with a flexible user

base than in the experiments with fixed groups. Future research could specifically focus on the effect of reputation effects in markets with new entrants.

Fourth, most of the theory about reputation systems assumes that trustors use reputation information to learn about the incentives of trustees. However, an alternative explanation that we discussed in chapter 5 is that reputation information allows trustors to reciprocate decisions made by trustees. Both mechanisms led to the same hypothesis and it is not possible to disentangle these two mechanisms with the available data. Future research should focus on these two mechanisms to better understand how reputation information affects trust. To do so, laboratory or online experiments could be used that are designed in a way that different motivations lead to different behavioral outcomes.

Future research

A question that could be addressed in future research, is why we did not find evidence for the 'compensation effect' in chapter 2 and 3, while earlier studies supported this hypothesis (Abrahao et al., 2017; Cui et al., 2017; Ert et al., 2016; Tjaden et al., 2018). These other researchers found that individuals with a minority background benefit more from having a good reputation than individuals belonging to a majority group. Part of the explanation for the difference may be that earlier studies did not look at ethnicity, but at other demographic characteristics (Abrahao et al., 2017), or that they did not look at actual transactions, but at the number of clicks on an offer (Tjaden et al., 2018). Another explanation may be that they used fake guest profiles that contained no information besides names and reputation scores (Cui et al., 2017), or reputation scores that were experimentally varied to increase their variance (Ert et al., 2016). This artificially carries users over the initial hurdle of getting the first positive review.

Future research could investigate if the compensation effect only exists for specific demographic characteristics, possibly even only for demographic characteristics that only have a small effect on trust. Future research could also look into the effect of the presence of more and less personal information: does the compensation effect get smaller with every extra source of information that is added? In chapter 2 we found that the first reviews a trustee acquired had

little effect on trust. Since most renters only had a few reviews, this may explain why reputation cannot compensate any negative discrimination. Future research could investigate if the average number of reviews and the effect each review has on trust affects the extent to which the reputation system reduces discrimination.

A second direction of future research about the role of reputation systems in digital discrimination is related to the use of (ranking) algorithms. The results from chapters 2 and 3 show that initial differences made on a platform affect future outcomes. Individuals that are selected for an exchange early on, benefit from reputation accumulation (Frey & Van De Rijt, 2016). Ranking algorithms that favor individuals with a better reputation, like Airbnb seems to do (Fradkin, Grewal, & Holtz, 2018), may further reinforce these differences, leading to even strong cascading effects. Future research should look into the dynamics of discrimination, reputation systems and algorithms to better understand how these newly emerging tools affect societal outcomes.

Lastly, future research should explore the extent to which reputation data portability changes interpersonal trust. Platform users are rather dependent on a single platform, because they would lose their reputation when moving to another platform. This creates a lock-in: users stay with a given platform, because it is too costly to give up their reputation and to make a fresh start at another platform (Dellarocas, 2010; Demary, 2015). Recently a few initiatives arose that allow platform users to include their reputations from different platforms in one overall profile (e.g. Traity, 2020). This development leads to several questions, for example on what reputation information users include in their profile and how useful this information is when assessing someone's trustworthiness (e.g. Norbutas, Ruiter, & Corten (2020b) and Teubner, Adam, & Hawlitschek (2019)). Another question is how this affects individuals without any reviews and individuals who have a negative reputation on one platform, but a positive reputation on another platform. On the one hand portable reputation systems may offer an opportunity for newbies on a platform to kickstart their participation on the platform. On the other hand, they may make it even more difficult for individuals who have no experience with participation in the platform economy before to join.

Final notes

A few months before the Couchsurfing trip to the Grand Canyon I joined a comparable trip to Yosemite National Park. The organizer of that trip had more than thirty positive reviews, communicated very openly and was very flexible in his plans. After this positive experience I was hesitant to join the trip to the Grand Canyon. Although the organizer of the trip to the Grand Canyon had some positive reviews, these reviews were not very detailed, and he did not provide a lot of information about his plans. In the end I decided not to join, partly because the timing did not suit my schedule, but also because I was not convinced that the organizer could be trusted.

Finally, a note on the structure of the remainder of this dissertation. The four chapters in this dissertation have been written as self-contained articles. This implies that there may be overlap and small inconsistencies between the chapters.

CHAPTER 2

The role of reputation systems in digital discrimination: a field study

This chapter has been co-authored by
Rense Corten and Arnout van de Rijt.

This chapter has been submitted to an international scientific journal.

Abstract

Reputation systems are commonplace in online markets, such as peer-to-peer sharing platforms. These systems are often proposed as the most promising solution to (ethnic) discrimination in these markets. This claim is based on earlier studies showing that the ethnic gap is smaller for users with reviews than for users without reviews. Using simulations, we show that this conclusion may be premature, as minorities have a harder time accumulating reviews. Hence, the greater benefit of reviews to minorities may be offset by their troubles acquiring any, thereby diminishing the potential for reputation systems to reduce discrimination. We use a unique dataset that contains information on all interactions that took place on a peer-to-peer motorcycle rental platform to test this hypothesis. We find that the reputation system does not reduce the initial inequalities between otherwise comparable renters of different ethnicity. Platforms that aim to reduce discrimination should not only make their reputation system more effective, they should also reduce initial differences between users.

Introduction

Across a wide range of markets, individuals have been found to avoid transacting with minorities. Such discrimination is a widespread phenomenon in online marketplaces for goods (Edelman & Luca, 2014; Nunley, Owens, & Howard, 2011). Examples of these marketplaces are second-hand selling platforms, such as eBay, and platforms where individuals can rent and borrow goods (e.g. cars and tools) from other individuals. Many of such marketplaces are supported by Internet platforms that enable direct communication between peers. Such platforms rely extensive personal profiles that include names and photos of users to foster interpersonal trust (Bente, Baptist, & Leuschner, 2012; Ert et al., 2016), but this abundance of personal information also leads to discrimination (Cui et al., 2017; Edelman & Luca, 2014; Edelman et al., 2017; Ert et al., 2016; Ge, Knittel, MacKenzie, & Zoepf, 2016; Jaeger & Sleegers, 2020; Jaeger & van Beest, 2019; Laouénan et al., 2017; Mohammed, 2017; Pope & Sydnor, 2011; Tjaden et al., 2018; Wu & Jin, 2018; Wu, Ma, & Xie, 2017) Discrimination in the platform economy results in fewer opportunities for users with certain demographic characteristics to buy, rent, sell and borrow goods. This outcome is disadvantageous both for the people who are discriminated against and for the platforms and possibly even for those individuals that discriminate, as it leads to suboptimal market outcomes where otherwise fruitful interactions are not realized.

It has been argued that reputation systems temper or even solve this problem of digital discrimination (Abrahao et al., 2017; Cui et al., 2017; Ert et al., 2016; Mohammed, 2017; Tjaden et al., 2018)[2]. The argument goes that discrimination on platforms mostly originates from a lack of information about other users. This information is especially important for interactions in the (peer-to-peer) platform economy, as these platforms are characterized by lower levels of institutionalization and higher levels of interpersonal trust (Katz, 2015; ter Huurne et al., 2017). By providing user-specific information on performance and trustworthiness through reviews from past transactions, reputation systems would supersede group stereotypes as a basis of partner choice and reduce or even eliminate unfounded inequalities. Whereas the use of photos and names

[2] Reputation systems collect, aggregate and distribute feedback about trustees' past decisions to trustors (Resnick et al., 2000).

always entails the risk of discrimination based on demographic characteristics, reviews are often considered a better and fairer way of reducing information asymmetry between users and providers.

Several studies have provided empirical support for the claim that reputation systems help overcome discrimination by showing that a difference between (ethnic) groups in the acceptance chances of a transaction request is reduced once users have received at least one positive review. For example, in an online experiment with Airbnb users, reputation was found to partly offset the tendency of people to trust others who are similar to them more (Abrahao et al., 2017). In another study it was even found that while guests on Airbnb with African American-sounding names were 19.2 percentage points less likely to be accepted than those with white-sounding names, this difference completely disappeared when both guests had a positive review (Cui et al., 2017).

However, here we argue that earlier research may have overstated the extent to which reputation systems mitigate discrimination. It is critical to consider that not all users are equally likely to achieve that initial reputation necessary for acquiring trust on the platform. Most platforms only allow users to write a review *after* a transaction via the platform. As previous research shows, the probability that a first request from a user without any reviews is accepted, and thus the probability of receiving a first review, strongly depends on a trustee's (ethnic) background (Abrahao et al., 2017; Cui et al., 2017). While a reputation system may decrease the importance of information retrieved from names and pictures once users have acquired reviews, it may altogether fail to reduce inequality in transaction volume as it gives majority members the ability to more quickly build a good reputation.

In the current research, we first demonstrate these inequality-sustaining effects of reputation systems with a simple model. To study the empirical relation between reputation systems and discrimination, we then analyze a unique dataset containing the complete historical records of user activity on a Dutch peer-to-peer motorcycle renting platform. This platform is similar in function and design to many other platforms, such as hospitality platform Airbnb and various carsharing platforms. However, whereas previous studies have analyzed platforms in a static way, e.g. comparing differences in the number of clicks an offer received (Tjaden et al., 2018) or prices Airbnb hosts could charge

(Edelman & Luca, 2014), the full historic record of all interactions on the online platform we study allows us to investigate the *dynamics* of inequality. In order to test whether reputation systems with time decrease or sustain discrimination on this platform, we analyze data on interactions at the platform at different timepoints as users accumulate reviews.

Theory

Modern day exchange is increasingly mediated by online platforms. 'Platform economy' is an umbrella term encompassing several activities, such as selling, exchanging, borrowing and renting of goods and services. Leong and Belzer (2017) identified two distinctive features of platforms. First, platform economy businesses make money not by providing goods and services per se, but rather by connecting people who need particular goods and services with people who want to provide them. Second, to facilitate this connection efficiently, platform economy businesses rely on online platforms.

Trust and reputation formation are crucial on sharing platforms, as interactions pose a high risk for the owners of the goods. Owners may thus be strongly inclined to derive trustworthiness of others from past experiences of other owners as well as borrower demographics. When an owner decides to rent out their goods to another user, who is generally a stranger, he or she runs the risk of not getting back the good in a good state. Although most sharing platforms usually offer some support in solving problems between users, the legal structures owners can rely on are usually limited (ter Huurne et al., 2017) Owners will therefore be motivated to carefully consider which renters can and which cannot be trusted before deciding to accept or reject a rental request. Our theoretical focus is the decision of the owner to accept or reject a request of a potential renter.

Discrimination in the platform economy

Unlike traditional exchange, in online markets people often exchange goods with perfect strangers from around the world (Frenken & Schor, 2017). There is limited opportunity to get to know a person before a transaction or to acquire information about the other person's unobserved qualities. The online nature of the platform economy implies several information asymmetries (ter Huurne

et al., 2017). First, consumers and providers are unsure about each other's intentions, leading to perceived personal safety risks, especially when the two actors meet offline after the online interaction. Second, consumers cannot check upfront whether the quality of the good they are buying, renting or borrowing is good. Following Akerlof's (1970) classical lemons problem, the risk of buying a low-quality good will result in market failure.

To mitigate this risk, platforms allow their users to create extensive user profiles, including names and profile pictures. Compared to traditional e-commerce companies, anonymity is lower in these platforms (Abrahao et al., 2017). Names, photos and descriptions are used as a means of identity verification and are intended to foster an increased sense of personal contact (Dubois et al., 2012; Guttentag, 2015). They also convey information about someone's ethnicity, gender and age and may thereby lead to discrimination on the basis of these demographic characteristics. Discrimination is defined as the unequal treatment of individuals or groups in the basis of their (demographic) characteristics (Pager & Shepherd, 2008).

The theoretical literature on discrimination distinguishes between two types: Statistical discrimination (Arrow, 1973) refers to discrimination on the basis of the perceived association between a person's demographic characteristics (such as ethnicity) and other characteristics that cannot be directly observed that impact the value or merit of that person. Taste-based discrimination on the other hand refers to a preference for certain characteristics over other characteristics, without an underlying expectation of qualities related to these characteristics (Becker, 1957). In practice it is difficult to distinguish between the two types of discrimination, because they are both manifested through a disadvantage for individuals with specific demographic characteristics. Empirically distinguishing between the two is not the aim of this chapter.

Discrimination in online markets is believed to be mostly statistical: it is caused by a lack of information about the a potential transaction partner's characteristics that impact the value or merit of that person, such as trustworthiness (Abrahao et al., 2017; Cui et al., 2017; Tjaden et al., 2018). Individuals fill in those gaps on the basis of the perceived association between a person's demographic characeristics (such as ethnicity) and those relevant characteristics. Applied to trust problem in the platform economy, statistical

discrimination is the tendency to derive expectations about the trustworthiness of exchange partners from demographics on their personal profiles. When owners assume that renters with specific dempgraphic characteristics are less trustworthy (e.g. because of stereotypes), they may place less trust in these individuals.

In the current study we focus on ethnicity as a socio-demographic dimension of discrimination, because we know from the labor market literature and earlier studies on the platform economy that this is a widespread and persistent form of discrimination (Rich, 2014; Zschirnt & Ruedin, 2016). We use data from a Dutch platform and follow the definition of Statistics Netherlands that classifies individuals' ethnicity by their biological parent's birth country. In 2019, 76.4% of the Dutch residents had parents who were born in the Netherlands (Statistics Netherlands, 2019). The second and third largest group consists of (children of) migrants from Turkey and Morocco. Labor market research shows that these migrants are disadvantaged in the Dutch labor market (Gracia, Vázquez-Quesada, & Van de Werfhorst, 2016). Discrimination then entails owners placing more trust in renters belonging to the ethnic majority (i.e. native Dutch renters) and renters belonging to the ethnic minority (i.e. non-native Dutch renters).

Reputation systems as a solution to statistical discrimination

Earlier research proposes that reputation systems may be the most promising solution to statistical discrimination. They are believed to provide more specific and therefore more accurate information about a renter's behavior in a transaction through the platform than more diffuse sociodemographic information (Resnick et al., 2000; Robbins, 2017). This would then allow minority renters to acquire the same interest as equally trustworthy majority members. After an interaction, users are asked to leave a rating and a review about their interaction partner. These ratings and reviews are displayed on their profile page for potential future interaction partners. Reputation systems allow users to assess the expected quality of an interaction with a potential partner ('learning' according to Buskens and Raub (2002)). Users with better reviews are more likely to have prosocial interests and are therefore more likely to live up to expectations and honor agreements, while users with negative reviews

likely care less about the interest of their partner, and are therefore more likely to abuse trust again in the future. In this way, ratings and reviews serve as a costly signal of trustworthiness and other unobservable characteristics of individuals, while they at the same time provide an incentive for prosocial behavior (Bolton, Katok, & Ockenfels, 2004b; Fehrler & Przepiorka, 2013; Resnick & Zeckhauser, 2002; Tadelis, 2016). For these reasons, requests from renters with a better reputation (i.e. more positive ratings) should be more likely to be accepted by the owners, and requests from renters with more negative reviews less likely.

Ratings on online platforms are typically highly left-skewed, with many very high ratings and only few ratings lower than 4.5 out of 5 stars (Teubner & Glaser, 2018; Teubner et al., 2017; Zervas et al., 2015). We therefore define a positive rating as having five stars and a negative rating as having fewer than five stars (Przepiorka, Norbutas, & Corten, 2017). This skewness may be due to underreporting of negative experiences (Fradkin et al., 2018), or may be a direct effect of the reputation system preventing negative experiences. As users' present behavior on the platform may affect their possibilities for future behavior, they may be motivated to behave well to build and maintain their reputation (Buskens & Raub, 2002). The lack of variation in average rating reduces the informativeness of the average rating of a user, but may increase the importance of the number of ratings a user has.

Previous research has confirmed that trustees with a better reputation are trusted more often (Boero et al., 2009; Charness et al., 2011; Duffy et al., 2013; Fehrler & Przepiorka, 2016; Gong & Yang, 2010). Especially the first review seems to matter (Duffy et al., 2013; Frey & Van De Rijt, 2016). One study found that users with one positive review received between 8.4% and 29.5% more trust than users without any reviews (Cui et al., 2017).

When a renter has no reviews, the owner relies on the limited information that is available, such as names and photos. Yet demographic information is only a proxy for trustworthiness, while reviews contain direct information about a renter's trustworthiness, so when reviews are available this should trump demographics in trust assessment. This supremacy of reviews over demographic information is confirmed in empirical studies: The difference between individuals of different ethnicity in the probability to participate on

platforms is much smaller when reputational information about the user is available (Abrahao et al., 2017; Cui et al., 2017; Ert et al., 2016; Mohammed, 2017; Tjaden et al., 2018).

> **H1:** The ethnic gap decreases with the number of positive reviews the renters have received.

Because of this 'compensation effect', reputation systems have been suggested to constitute a solution to digital discrimination by researchers (Abrahao et al., 2017; Cui et al., 2017; Ert et al., 2016; Mohammed, 2017; Tjaden et al., 2018) and practitioners (Murphy, 2016) alike. Minority renters would be discriminated less and less as they file more requests and thus accumulate positive ratings.

> **H2:** The ethnic gap decreases with the number of requests renters have filed.

Reputation systems as an amplifier of discrimination

However, while reputation systems may indeed decrease discrimination for those users who have positive reviews, the expectation that reputation systems decrease overall discrimination on the platform is based on the implicit assumption that all users are equaly likely to obtain a good reputation. Yet, as reviews can generally only be written after an interaction on the platform, not all users are equally likely to obtain reviews. The results from previous research suggest that users who already have a good reputation are more often selected for new interactions, which in turn improves their reputation and thus the probability that their next request will be accepted. In a laboratory experiment Frey and Van De Rijt (2016) show that these 'reputation cascades' lead to arbitrary inequality between equally trustworthy individuals. The chance to build a reputation is thus dependent on the probability that an initial request is accepted. If this probability differs between people with different ethnicity, members of some ethnic groups may have fewer opportunities for reputation building than others. We thus argue that minorities' requests for a transaction are less likely accepted, so that per filed request (either accepted, declined, or not responded), ethnic minority renters receive fewer reviews than ethnic majority renters. As such, the greater benefit of positive reviews to minority members enabled by reputation systems is counteracted by minorities' greater struggles to obtain such reviews.

We use simulations to better understand this interplay between reputation and discrimination. In the simulations, both minority and majority group renters propose transactions. These proposals are accepted and rejected based on a combination of minority group status and the number of positive reviews received. When a request is accepted, a renter receives a positive review. We vary the extent to which there is initial trust in members of the majority group (12-88%) and minorities (2-73%), the extent to which reputation has a positive effect on the acceptance probability, and the extent to which ethnic minority renters may compensate their initial disadvantage with reputation. We simulate a sequence of ten periods in each of which every renter submits a rental request. After ten periods we evaluate to what extent the success rate of majority and minority renters differs[3]. A detailed description of the model and the analyses can be found in the Appendix.

The simulations shows that under a broad range of conditions inequality between renters of different ethnicities on a platform with a reputation system does not decrease and sometimes even increases, in contrast to Hypothesis 2. Whether inequality decreases with the number of requests renters have filed, is perpetuated, or increases turns out to predominantly depend on the initial level of trust in ethnic minorities. Figure 2.1 shows the average rate with which requests from ethnic majority and ethnic minority renters are accepted over the rounds in the simulations for three different levels of initial trust in the minority. The converging lines in panel C of Figure 2.1 show that when the initial level of trust in ethnic minority renters is high, their initial disadvantage can be compensated through reputation, consistent with what past scholars have argued (H2). Panels A and B, however, show that when requests from ethnic minority renters are less frequently accepted, these renters do not get the opportunity to build a reputation and can therefore not benefit from the

[3] On many platforms, like eBay, Airbnb and Couchsurfing, most users but especially consumers (i.e. buyers, renters), transact only a handful of times (Lauterbach et al., 2009; Resnick & Zeckhauser, 2002; Teubner, 2017). Renters on the motorcycle sharing platform that we study in the next sections on average submitted 2.5 requests since the start of the platform, so with ten periods in our simulation model we study a time frame that includes the vast majority of user histories on our platform and most other platforms. Most renters will never submit ten requests, and if the reputation system fails to reduce the inter-ethnic inequality in acceptance chances in a few rounds, most minorities will never experience equal chances.

reputation system and the acceptance rates of the different renters diverge. Hence intial differences between renters in the probability to be accepted for an interaction may accumulate over time, thereby diminishing the potential of reputation systems to decrease discrimination. Only when trust is less of an issue to begin with, the reputation system may reduce inequality.

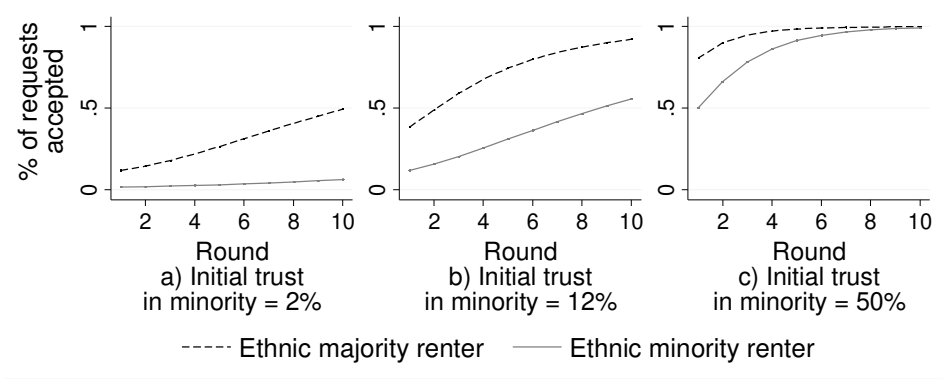

Figure 2.1: Simulation results: average acceptance rate of native and non-native Dutch renters over the rounds.

Based on these simulation results, we predict that given sufficiently low initial success chances of minorities, the ethnic gap will not diminish on a platform with a reputation system, in contrast to H2. This is the main contribution of the chapter: contrary to what is commonly believed, we argue that reputation systems do not necessarily lead to a decrease in inequality in online markets. Prior research has focused on static differences in trust between individuals of different ethnicity with and without a reputation score. They compared how much trust was placed in individuals without reviews with how much trust was placed in individuals with at least one positive review. They did not discuss why some individuals have positive reviews and others do not. However, as we have argued, an individual's reputation is not exogenously determined: it is conditional on an individual being selected in interactions, which in turn may depend on that individual's ethnicity and prior reputation depend. In the current chapter we therefore view reputation through a dynamic lens. Instead of comparing individuals with and without reviews, we study the process of reputation formation itself. The purpose of the analyses is to test if the previous findings that reputation system reduce discrimination still hold true when

taking into account that the process of reputation formation itself may be biased.

Data and methods

We study a Dutch peer-to-peer motorcycle sharing platform that was founded in 2016. The platform operates within the so-called sharing economy and has a similar design and functionality as peer-to-peer carsharing platforms and hospitality platforms such as Airbnb. It is common on sharing platforms for users to advertise their goods and also this platform allows motorcycle owners to advertise their motorcycle. Renters can browse through the listed motorcycles and the personal profiles of the owners and send rental requests for specific time slots and for a predefined price. Before accepting or declining the request, the owner can view the personal profile of the renter, including first name, photo, personal description and reviews. These personal profiles are commonplace on online platforms. When the request is accepted, the renter pays the rental price to the platform and the owner and renter meet offline to hand over the motorcycle. Similar to the payment system on other platforms, the platform keeps the renter's payment until the transaction has been completed. After the rental period the platform transfers the money paid by the renter to the owner. Both renter and owner are asked to write a review that is publicly displayed in their user profile. Many platforms automatically arrange full insurance options for rental goods, and also the present platform automatically insures motorcycles during the rental period and checks drivers licences and fraud histories of all users.

The platform facilitates the sharing of a fragile good of high value, motorcycles, with strangers. Risk is thus particularly high. Moreover, motorcycle owners attach certain non-material values to their motorcycles that cannot be reimbursed by insurances, as illustrated by the following quotes: "How do you let some random person ride your motorcycle? I would never do that! It's quite hard to watch your baby go the first time. No matter how attached you are to your motorcycle, it's difficult watching someone else ride it off. Even when you sell it!" (Hooshmand, 2018). "The first mental hurdle to get over is the fact that you'd be "cheating" on your bike if you decide to ride someone else's" (Klinger,

2018). In addition to the material and emotional risk of entrusting a stranger with a motorcycle, the transaction requires that the parties meet offline before and afterwards which may pose personal safety concerns, especially in case of conflict.

We analyze the complete historical records of user activity on the platform. The dataset provided to us by the platform contains information on all (11,418) interactions that took place since the start of the platform, May 2016, to July 2017[4]. We excluded unfinished requests that were not filed, requests that were cancelled by the platform[5], and requests from renters who had rented from the same owner in the past (as we are interested in trust between strangers). The remaining data include 7,181 requests for 973 motorcycles sent by 2,896 renters to 851 owners.

Dependent variable

Our dependent variable is the decision of the owner. They can either actively accept (2,626 requests, 36.6%) or decline (2,443 requests, 34.0%) a request, or not send a response at all, after which the request expires (2,112 requests, 29.4%). An earlier study of discrimination on the basis of sexual orientation shows that discrimination on Airbnb is mostly driven by non-responses rather than outright rejection (Ahuja & Lyons, 2019). Because both result in an unfulfilled request, we combined the latter two categories into the category 'declined'. A robustness check shows that the results do not change when excluding expired requests from the analyses.

Independent variables

To operationalize the reputation of the renter, we created two continuous variables, indicating the number of positive ratings and the number of negative ratings. Since 92.1% of all reviews are 5-star ratings, we define a positive rating as having five stars and a negative rating as having fewer than five stars (Przepiorka et al., 2017). The variable "number of positive reviews" counts the

[4] Utrecht University and the motorcycle sharing platform signed a confidentiality agreement to ensure GDPR compliance.

[5] The platform automatically cancels requests for two reasons: 1) the renter does not have the required driver's license; 2) renters can send multiple (similar) requests at the same time. When one of the requests is confirmed, the other requests are automatically cancelled.

number of 5-star reviews a renter received. The variable 'number of negative reviews' indicates the number of reviews with less than five stars a renter received. Renters without any reviews serve as the reference category.

We make use of data from the the Dutch Civil Registration (DCR) to operationalize the users' ethnicity (Edelman et al., 2017; Hofstra & de Schipper, 2018; Laouénan et al., 2017). All names are verified by the insurance company with which the platform collaborates, so the names visible to the owners are the renters' real names. The DCR data are register data of those who have Dutch nationality and were alive and living in the Netherlands. We have aggregated DCR data that comprises 3,800 unique first names (85.0% of the unique names in the platform dataset, covering 93.0% of the users of the platform[6]). Per name we know the frequencies of combinations of the parents' birth countries. Based on the definition of Statistics Netherlands, country-combinations are classified into one of two ethnic origin groups: Native Dutch or non-native Dutch (Statistics Netherlands, 2018). In cases where both parents were born in the Netherlands, the combination is assigned to the native Dutch group. If only one or neither parent was born in the Netherlands, the combination is classified as non-native Dutch. Per name we then calculate the probability that a user with that name is native Dutch. We include this continuous variable in the analyses. For some of the tables and figures we used dichotomous ethnicity variables rather than continuous ones. In such cases we assigned users their most likely ethnicity. The final independent variable is the cumulative request count of the renter, which is the total number of requests made by the renter at the time of the request, including the current request.

Control variables

We included all other information visible to motorcycle owners: the price and duration of the rental, the number of years a renter had been a member on the

[6] 123 of the unique names in the platform data contained special symbols (√, ∏ etc.) and could therefore not be matched with the DCR-data. The appearance of these symbols is probably caused during the transition from the platform to Excel. There does not seem to be a relation between the ethnicity of the renter and these symbols (two examples: 'G√∂khan' and 'Mari√´lle'). Among the remaining 37 names that could not be matched non-native Dutch names seem to be overrepresented ('Dzahid' and 'Carone'), but there are also names that are more likely to belong to native-Dutch renters ('Juergen' and 'Paulpeter').

platform and the age and gender of the renter[7], whether the renter had a profile picture and a linked Facebook account. We included the same information about owners and we also added the following control variables related to the motorcycles: weight, engine displacement (cc), power (hp), age of the motorcycle, and the number of positive and negative ratings of the motorcycle. We also included the total number of requests submitted by a renter, to account for differences in the tendency to submit a new request[8]. The final control variable is time. As requests are not evenly distributed over time, with peaks in summer months and virtually no activity in winter, we operationalize time by the total number of completed requests at the time of the new request (divided by 1000).

Analytical strategy

The data have a cross-classified structure: requests are nested within renters and owners, but renters and owners are not nested within each other: the same renter interacts with multiple owners and vice versa. To test the hypotheses, we ran a cross-classified multilevel linear probability model with as dependent variable a dummy variable indicating whether the request was accepted or not. We used a Bayesian estimator with two MCMC chains and noninformative priors in Mplus (Muthén & Muthén, 2017). We used the default convergence criterion of Mplus (Proportional Scale Reduction factor lower than 1.1). We used 50,000 iterations of which the first half is considered a burn-in phase.

[7] We apply the same methodology to estimate the gender of the users as for estimating their ethnicity.

[8] We tested if past experiences at the platform affected renter's tendency to submit another request using a cross-classified multilevel logistic regression regression with the dependent variable indicating whether a renter submitted at least one more request after the current one. The renter's ethnicity, the fraction of previous succesful requests, and the interaction between the two are included as independent variables. We found that renters who have experienced more rejections are indeed more likely to file a new request (b = -0.131, 95% Conf. Int. = -0.246, -0.015). Moreover, we found that the interaction between the historical success rate and the renter's ethnicity is significant and negative (b = -0.043, 95% Conf. Int. = -0.091, -0.003). That means that even though renters who submit another request tend to have a lower success rate than renters who do not file another request, this is less so for non-native Dutch renters than for native Dutch renters.

We included random intercepts for renters and owners in all analyses. In the second and third model we added random slopes for reputation. We used full information maximum likelihood (FIML) to deal with missing values. This method includes partial information in observations with missing values, which allowed us to use all available information (Collins, Schafer, & Kam, 2001).

In the first model we included the renter's ethnicity and reputation to check if there is indeed inequality between renters of different ethnicity, and to see if reputation increases trust. We control for the total number of requests submitted by the renter in all models. To the second model we added the interation between the renter's ethnicity and reputation to test if renters with a better reputation are discriminated less, testing Hypothesis 1. The third model included the renter's reputation and ethnicity, which is interacted with the cumulative request count of the renter, testing Hypothesis 2.

We ran the same models with and without control variables. We ran the models in several steps. First we only included variables at the level of the transaction (level 1). We then removed insignificant variables before adding renter- and owner-level variables (level 2 and 3). We also removed insignificant variables before adding cross-level interactions, except when we were testing cross-level interactions that involved insignificant variables.

Results

Descriptive statistics

Table 2.1 contains the descriptive statistics on all variables.

Figure 2.2 shows the relationship between the acceptance rate and the renter's ethnicity. Requests from renters with more Dutch-sounding names are accepted more often. Native Dutch requesters are about twice as likely to receive a positive response as non-native Dutch requesters.

Table 2.1: Descriptive statistics

Variable	Level	N_{miss}	Mean	SD	Min	Max
Dependent variables						
Accepted	I	0	0.37	0.48	0	1
Declined	I	0	0.34	0.47	0	1
No response	I	0	0.29	0.46	0	1
Independent variables (renter characteristics)						
Dutch ethnicity	R	223	0.79	0.30	0	1
# of positive reviews	I	0	0.24	0.70	0	12
# of negative reviews	I	0	0.03	0.19	0	2
request #	I	0	2.84	2.75	0	27
Control variables						
# of rental days	I	0	1.81	1.49	1	33
Motorcycle characteristics						
Day price	I	2	66.1	25.5	20.9	184.6
# of reviews: positive	I	0	3.60	6.60	0	53
# of reviews: negative	I	0	0.78	1.54	0	9
Age	I	4	12.83	7.47	0	61
Weight	I	189	217.97	40.53	93	585
Engine displacement (CC)	I	12	865.2	269.1	113	2,294
Horsepower (HP)	I	0	87.74	29.99	10	173
Renter characteristics						
Age	I	2	35.56	10.81	20	80
Member #years	I	0	0.20	0.36	0	2.12
Female	R	191	0.12	0.31	0	1
Profile picture (%)	R	0	0.61	0.49	0	1
Facebook verification	R	0	0.26	0.44	0	1
Cumulative total # of requests (/1000)	I	0	3.591	2.073	0.001	7.181

(Continued on next page)

Table 2.1: Continued

Variable	Level	N_{miss}	Mean	SD	Min	Max
Owner characteristics						
# of reviews: positive	I	0	2.31	3.67	0	33
# of reviews: negative	I	0	0.16	0.50	0	6
Age	I	2	35.74	10.21	18	73
Member #years	I	0	0.64	0.53	0	2.29
Dutch ethnicity	O	142	0.85	0.22	0	1
Female	O	117	0.08	0.25	0	1
Profile picture (%)	O	0	0.81	0.39	0	1
Facebook verification	O	0	0.37	0.48	0	1

Note: Variables are measured at the level of the interaction/request (I), renter (R) or owner (O).
$R_{equests}$ = 7,181
$N_{renters}$ = 2,896
N_{owners} = 851

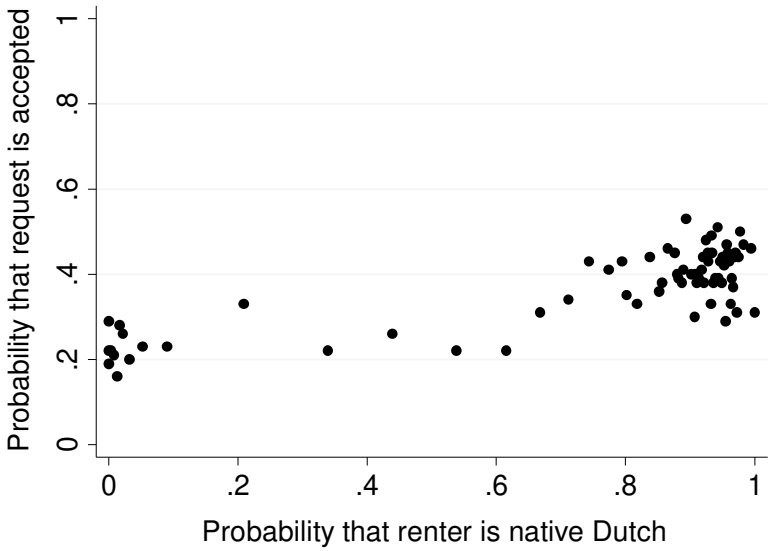

Figure 2.2: Average fraction of requests accepted, by renter's ethnicity. Each point represents 100 requests.

Figure 2.3 shows the relationship between a renter's reputation and the probability that a request is accepted. Having more positive reviews seems to have a positive effect on the acceptance rate, both for native Dutch and non-native Dutch renters (panel A). Having negative reviews seems to be detrimental for the probability to receive trust, especially for non-native Dutch renters (panel B). A note here is that there are no non-native Dutch renters with more than five positive reviews, suggesting that indeed they have a harder time establishing a reputation.

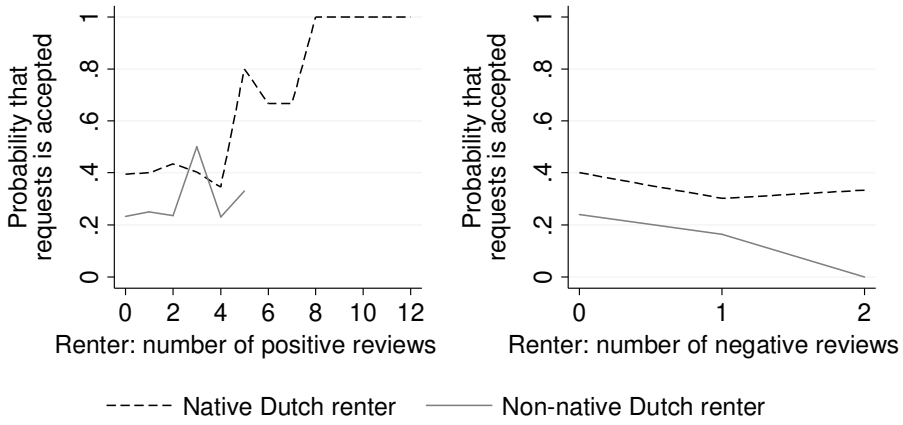

Figure 2.3: Average fraction of requests accepted, by renter's ethnicity and number of reviews.

Figure 2.4 shows how the acceptance rate changes with every request filed for renters of different ethnicity. Only renters who have submitted five or more requests are included to control for the changing composition of the sample, which may for example be caused because renters who were previously unsuccessful are more likely to submit another request[9]. Per renter we randomly drew one request in order to account for individual differences in the probability that their request is accepted due to characteristics other than ethnicity. The figure is based on the average acceptance rate for 500 draws. The figure shows the acceptance chances for every request filed for native Dutch (N = 255) and non-native Dutch renters (N = 87). The ethnic gap is again

[9] The figure looks similar when including only renters who submitted at least three or seven requests.

visible: native Dutch renters seem to have a higher probability of getting a positive response than non-native Dutch renters. The figure shows an increase in the probability that a request was accepted with every additional request filed. Crucially, the acceptance rates of native and non-native Dutch renters do not converge. Apparently, the reputation system is not able to overcome to persistently lower acceptance rates of minority requesters.

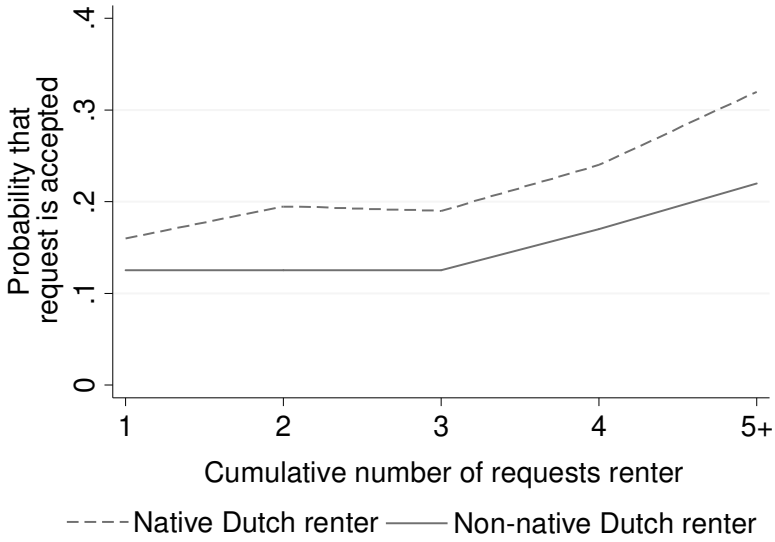

Figure 2.4: Fraction of requests accepted, by the cumulative number of requests made by the renter and the renter's ethnicity (i.e. '1' on the x-axis refers to the first request made by a renter, '2' to the second request made by the same renter). Only renters who submitted five or more requests are included. Datapoints representing the five or more requests are collapsed into one category.

Regression results

Table 2.2 contains the results from the regressions without control variables and Table 2.3 the results with control variables. We find that requests from renters with more Dutch-sounding names have a higher probability of being accepted by the owners (Table 2.2, Model 1). This gap is persistent even when control variables are added (Table 2.3, Model 4).

We also find that having more positive reviews increases the probability of getting a positive response (Table 2.2, Model 1), and this effect is persistent when control variables are added (Table 2.3, Model 4). Having more negative reviews does not affect the acceptance probability, neither in the regressions without control variables (Table 2.2, Model 1), nor in the models with control variables (b = -0.031, 95% Conf. Int: -0.089, 0.026, p=0.140).

We expected that the ethnic gap would be smaller for renters with more and positive reviews (Hypothesis 1), but we do not find evidence for this in either the regression models with or without controls (Table 2.2 and 2.3, Model 2 and 5). The findings suggests that having more positive reviews is equally beneficial for native Dutch and non-native Dutch renters. This implies that the ethnic gap does not become smaller once renters have reviews.

To test Hypothesis 2, we examine changes in the ethnic gap with every request filed by a renter. The results of this analysis are shown in Table 2.2, Model 3 (without controls) and Table 2.3, Model 6 (with controls). We find that the probability that a request is accepted increases with the number of requests a renter has already filed. However, we do not find that there is a difference between renters with Dutch-sounding names and non-Dutch sounding names in the extent to which the probability that a request is accepted changes with the number of requests done by that renter. This means that the disadvantage non-native Dutch renters have vis-a-vis native Dutch renters does not decrease as they file more requests. This leads us to reject Hypothesis 2.

Table 2.2: Results of the multilevel cross-classified linear probability model of outcome of the request: accepted (0/1) without control variables

Independent variables	Model 1	Model 2	Model 3
Intercept	0.321***	0.331***	0.385***
	(0.288, 0.352)	(0.285, 0.364)	(0.337, 0.435)
Main effects			
Renter Dutch ethnicity	0.144***	0.137***	0.161***
	(0.117, 0.175)	(0.107, 0.185)	(0.113, 0.205)
Renter # positive ratings	0.058***	0.070***	-
(ref. cat. = no reviews)	(0.042, 0.073)	(0.041, 0.113)	
Renter # negative ratings	0.002	-0.032	-
	(-0.053, 0.057)	(-0.149, 0.052)	
Renter request #	-	-	0.041***
			(-0.045, -0.037)
Interactions			
Renter Dutch * Renter #	-	-0.007	-
postive ratings		(-0.062, 0.025)	
Renter Dutch * Renter #	-	0.049	-
negative ratings		(-0.044, 0.207)	
Renter Dutch * Renter	-	-	-0.003
request #			(-0.014, 0.006)
Control variables			
Renter total number of	-0.024***	-0.025***	-0.041***
requests	(-0.026, -0.021)	(-0.028, -0.022)	(-0.045, -0.037)
Residual variance renter-	0.001***	0.002***	0.007***
level	(0.000, 0.003)	(0.000, 0.005)	(0.002, 0.015)
Residual variance owner-	0.039***	0.039***	0.042***
level	(0.033, 0.046)	(0.033, 0.046)	(0.035, 0.049)
Observations	7181	7181	7181

Note: 95% Confidence interval in parentheses. *** p<0.001, ** p<0.01, * p< 0.05 (one-tailed test).

Table 2.3: Results of the multilevel cross-classified linear probability model of the outcome of the request: accepted (0/1) with control variables

Independent variables	Model 4	Model 5	Model 6
Constant	-0.104*	0.078	0.068
	(0.000, 0.206)	(-0.019, 0.174)	(-0.033, 0.173)
Main effects			
Renter Dutch ethnicity	0.138***	0.135***	0.144***
	(0.108, 0.169)	(0.095, 0.170)	(0.100, 0.191)
Renter # positive ratings	0.028**	0.034	0.028**
(ref. cat. = no reviews)	(0.011, 0.045)	(-0.015, 0.072)	(0.010, 0.046)
Renter # negative ratings	-	-0.111*	-
		(-0.201, 0.002)	
Renter request #	0.034***	0.036***	0.036***
	(0.028, 0.040)	(0.030, 0.043)	(0.027, 0.044)
Interactions			
Renter Dutch * Renter	-	-0.003	-
#reviews: positive		(-0.051, 0.047)	
Renter Dutch * Renter	-	0.119	-
#reviews: negative		(-0.023, 0.207)	
Renter Dutch * Renter	-	-	-0.001
request #			(-0.011, 0.007)
Control variables			
Rental characteristics			
# Rental days	0.002	0.002	0.002
	(-0.005, 0.009)	(-0.005, 0.009)	(-0.005, 0.009)
Cumulative total # of	-0.017***	-0.017***	-0.017***
interactions (/1000)	(-0.024, -0.010)	(-0.024, -0.010)	(-0.024, -0.009)
Renter characteristics			
Renter total number of	-0.039***	-0.040***	-0.039***
requests	(-0.043, -0.035)	(-0.044, -0.036)	(-0.043, -0.035)
Renter female	-0.037*	-	-
	(-0.072, -0.001)		
Renter has profile pic	0.056***	0.056***	0.057***
	(0.032, 0.082)	(0.033, 0.082)	(0.032, 0.081)
Renter has Facebook	-0.035**	-0.036**	-0.037**
verification	(-0.061, -0.010)	(-0.060, -0.012)	(-0.062, -0.011)
Renter age	0.001**	0.001**	0.001**
	(0.000, 0.002)	(0.000, 0.002)	(0.000, 0.003)

(Continued on next page)

Table 2.3: continued

Independent variables	Model 4	Model 5	Model 6
Owner characteristics			
Owner Dutch ethnicity	0.138***	-	-
	(0.108, 0.169)		
Owner age	0.005***	0.005***	0.005***
	(0.004, 0.007)	(0.004, 0.007)	(0.004, 0.007)
Owner member # years	-0.071***	-0.073***	-0.073***
	(-0.102, -0.040)	(-0.104, -0.042)	(-0.103, -0.042)
Owner female	0.025	-	-
	(-0.039, 0.091)		
Owner has profile pic	0.140***	0.126***	0.126***
	(0.086, 0.195)	(0.073, 0.178)	(0.075, 0.177)
Owner has Facebook verification	-0.030	-	-
	(-0.068, 0.009)		
Motorcycle characteristics			
Motorcycle age	-0.003**	-0.003**	-0.003**
	(-0.006, -0.001)	(-0.006, -0.001)	(-0.006, -0.001)
Motorcycle # negative ratings	0.013*	0.013*	0.013*
	(0.001, 0.025)	(0.001, 0.025)	(0.001, 0.025)
Residual variance renter-level	0.002***	0.004***	0.004***
	(0.000, 0.005)	(0.001, 0.009)	(0.000, 0.012)
Residual variance owner-level	0.035***	0.035***	0.035***
	(0.029, 0.042)	(0.029, 0.042)	(0.029, 0.042)
Observations	7181	7181	7181

Note: 95% Confidence Interval in parentheses. *** p<0.001, ** p<0.01, * p< 0.05 (one-tailed test).

Robustness checks

When using a logistic regression instead of a linear probability model, there are only minor differences with the results in Table 2.2 and 2.3. The coefficient of the number of negative reviews in Model 2 is insignificant when using a linear probability model, but is negative and signficant when using a logistic regression. The interaction between having negative reviews and the ethnicity of the renter in Model 5 is significant and positive when using a logistic regression. This means that non-native Dutch renters suffer more from negative reviews. When using a logistic regression the positive effect of having positive reviews disappears in model 6.

In the main analyses we operationalized reputation as the number of positive and negative reviews. When reputation is operationalized as having no positive reviews versus one or more reviews the results do not change. When taking the natural logarithm of the number of reviews the coefficient of the interaction between the ethnicity of the renter and the number of negative reviews a renter has is positive and significant in the model with control variables, which means that non-native Dutch renters experience more negative effects of having negative reviews. The other results remain the same.

Excluding non-responses (i.e. requests that are neither accepted nor rejected) from the analyses does not affect the results.

Control variables

Table 2.3 includes the coefficients of the control variables that were included in the main analyses. We find that requests from older renters and from renters with a profile picture have a higher probability of being accepted, while requests from female renters and renters who have verified their Facebook account are less likely to be accepted. Native Dutch, female and older owners, as well as owners with a profile picture accept more requests. Owners who have been a member for a longer time and owners who have verified their Facebook account are less likely to accept a request. Requests for longer rental periods and for motorcycles with a higher number of negative reviews are accepted more often. Requests for older motorcycles have a lower probability of being accepted. The probability that a request is accepted decreases over time. The probability that a request is accepted does not vary with the number of rental days.

Table 2.3 does not include the coefficients for insigificant control variables at the request level, which were removed before adding higher level control variables. Table A2.3 in the Appendix shows the model with all request-level control variables. We found that the day price, weight, engine displacement (cc) and power of the motorcycle, duration of the membership of the renter, reputation of the owner and the number of positive reviews of the motorcycles do not affect the probability that a request is accepted.

Exploratory analysis

To better understand the mechanisms leading up to those the findings, we explored a number of potential explanations. First, we explored whether owners on the platform indeed place more trust in renters who have the same ethnicity as they have (homophily). Homophily is treated as a behavioral outcome measure here, since we do not have information about the mechanisms leading to this outcome. In another set of cross-classified linear probability models with the acceptance of requests as dependent variable, we included the difference between the renter's ethnicity and the owner's ethnicity (i.e. the difference in the extent to which their names sound native Dutch) as the independent variable. We also included the renter's ethnicity to see whether the distance between the ethnicity of the renter and the owner explains all of the discrimination. The results of these regressions are given in Table 2.4, model 7. When the renter and owner are of the same ethnicity, the probability that the request is accepted is higher. However, even when controlled for the difference between the owner's and renter's ethnicity, requests from renters with non-Dutch sounding names are accepted less often. This means that homophily is not the only explanation for the ethnic gap. Besides a preference for renters who have the same ethnicity as they have, non-native Dutch owners also have a preference for native Dutch renters. Table 2.5 shows the acceptance rates of different combinations of renters and owners. It seems that the homophily effect is mostly present for native Dutch owners and owners of Turkish ethnicity.

Second, we explored if renters without Dutch sounding names indeed receive fewer reviews. Using a linear regression of the number of reviews per filed request on the renter's ethnicity shows that this is indeed the case. Per request, renters with completely Dutch-sounding names receive 0.18 reviews more than renters with names that do not sound Dutch at all (b = 0.181, z = 7.34, p < .001). This difference is entirely caused by the difference in acceptance rate between renters with the Dutch ethnicity and renters with other ethnicity, as renters with less Dutch-sounding names receive slightly more reviews per accepted request (b = -0.055, z = -2.05, p = .041). There is a significant difference in the rating owners give to renters of different ethnicity: compared to names that are completely native Dutch, renters with non-native Dutch

names on average get ratings that are 0.13 points lower (*b* = 0.134, *95% CI* = (0.058, 0.207), *p* <.001*)*.

Table 2.4: Results of the exploratory analyses.

Independent variables	**Model 7: DV = Accepted**	**Model 8: DV = Submit another request**	**Model 9: DV = Accepted**
Intercept / threshold	0.939*** (0.763, 1.119)	-0.748*** (-0.838, -0.649)	1.038*** (0.883, 1.189)
Main effects			
Distance ethnicity	0.196* (-0.375, -0.013)	-	-
Renter Dutch	0.580*** (0.411, 0.756)	-0.131* (-0.246, -0.015)	0.737*** (0.572, 0.898)
Success rate	-	-1.682*** (-1.887, -1.508)	-
Experience owner	-	-	-0.092*** (-0.139, -0.050)
Interactions			
Renter Dutch * success rate	-	-0.043* (-0.091, -0.003)	-
Renter Dutch * experience owner	-	-	0.029 (-0.010, 0.071)
Variance intercept renter-level	0.058*** (0.021, 0.112)	0.016*** (0.002, 0.046)	0.115*** (0.066, 0.191)
Variance intercept owner-level	0.554*** (0.452, 0.681)	0.028*** (0.007, 0.056)	0.015*** (0.012, 0.020)
PPP	0.499	-	-
Observations	7181	7181	7181

Note: 95% Confidence Interval in parentheses. *** p<0.001, ** p<0.01, * p< 0.05 (two-tailed test).

Lastly, we explored whether more experienced owners behave differently than less experienced owners. Since there are very few cases in which the renter turns out to be untrustworthy (as reflected in the extremely high reputation scores and the low frequency of insurance claims), we would expect that more experienced owners are better able to select the relevant information for assessing the trustworthiness of the renter. Hence, we expect that more experienced owners trust more and discriminate less. To test this hypothesis

we ran a multilevel linear probability model with the outcome of the request (accepted or not) as the dependent variable. We included the renter's ethnicity and the number of completed transactions of the owner, as well as the interaction between the two. The results of this regression are in Table 2.4, Model 8. We find that more experienced owners are less likely to accept a request, and that they are not less likely to discriminate. This suggests that more experienced owners are more selective, possibly because they receive more requests in general.

Table 2.5: Number of accepted requests and acceptance rate (in parentheses) per renter-owner ethnicity combination.

		Owner				
		Dutch	*Moroccan*	*Turkish*	*Other*	*Total*
Renter	*Dutch*	2123	61	50	134	2296
		(39.9%)	(33.7%)	(32.7%)	(46.2%)	(39.7%)
	Moroccan	174	16	16	13	203
		(24.5%)	(21.1%)	(30.8)	(26.5%)	(24.3%)
	Turkish	114	11	7	9	137
		(28.6%)	(33.3%)	(36.84%)	(33.3%)	(29.3%)
	Other	107	5	1	16	134
		(24.1%)	(16.7%)	(7.1%)	(38.1%)	(24.7%)
	Total	2395	88	73	176	2626
		(36.8%)	(29.2%)	(30.5%)	(41.8)	(36.6%)

Discussion

Reputation systems are often proposed as the most promising solution to (ethnic) discrimination in online markets (Abrahao et al., 2017; Cui et al., 2017; Ert et al., 2016; Mohammed, 2017; Tjaden et al., 2018). In the current chapter we argue that reputation systems fail to solve discrimination except in cases where minorities have little problem finding a transaction partner to begin with. Using simulations, we developed a new argument about the interplay between discrimination and reputation systems. We argue that reputation systems may sustain inequality between renters with different ethnic backgrounds, even when discriminated renters can compensate their initial disadvantage with reputation. We explain this by the fact that reputation systems themselves are

biased: only individuals who are given the opportunitiy to participate in interactions can accumulate reviews.

Previous research concluded that reputation systems decrease inequality. We tested if this hypothesis is still supported when taking into account that the reputation formation process itself may be biased by using data from a Dutch motorcycle sharing platform. While previous researchers concluded that reputation systems help overcome inequality, we find that the reputation system fails to reduce the ethnic gap. The difference between majority and minority members in the likelihood of having a request granted persists even for renters with positive ratings. Regardless of the reputation of the renter, requests from renters with an ethnic minority background are less likely to be accepted. We show that this decreases their probability of getting a (positive) review, which in turn further decreases their chances to participate in future interactions. This implies that, contrary to what has been assumed in previous research, reputation system may be biased and therefore that the extent to which they may reduce inequality is limited.

We argue that our research setting allows for a particularly clean test of the hypothesis that reputation systems may increase discrimination, given that it involves a clear trust problem. The platform we studied is an examplary case of a sharing platform in which consumers grant each other temporary access to underutilitzed physical assets (Frenken & Schor, 2017). The platform is similar to many other sharing platforms in its procedures to match providers to consumers, the institutions in place that stimulate trust formation (e.g. secure payment system, insurances) and the risks that owners and renters encounter. Moreover, rather than investigating the effect of ethnicity and reputation on the number of clicks an offer receives or on a proxy of the number of bookings, we had access to a complete dataset of all interactions that ever occurred on the platform. Our dataset contains data on real and complete user profiles, which allowed us to study a more natural setting than in lab experiments where user profiles are constructed by the researchers.

A limitation of our test of the theory is that there are only very few ethnic minority renters with many positive reviews on the platform we studied. This limits the possibility to test to what extent reputation may compensate for the initial disadvantage of these renters. At the same time, this limitation in itself

serves as a proof of our argument: we predicted that ethnic minority renters are less likely to obtain reviews, which is reflected in the low number of ethnic minority renters with many positive reviews.

Our contibution to the literature on discrimination and reputation is twofold. First, we reevaluated the claim that if objective information derived from third parties (positive ratings) is available, in-group preferences no longer matter for acceptance rates or are at least significantly reduced (Abrahao et al., 2017; Cui et al., 2017; Ert & Fleischer, 2019; Laouénan et al., 2017; Tjaden et al., 2018). Although reputation is beneficial both for native ethnic majority and ethnic minority renters, we do not find evidence for this "compensation hypothesis". Motorcycle owners may generally attach more non-material value to their motorcycles than users of other platforms to their shared possessions, as the quotes in the introduction suggest. This may strengthen the preference for renters with specific characteristics, as owners may simply prefer not to have certain renter with certain characteristics ride their bike (their 'baby'), regardless of how trustworthy that renter has proven to be. The discrimination on this platform may not be statistical (i.e. caused by a lack of information), but based on other motives. This reasoning could also explain why more experienced owners do not discriminate less. Another explanation for why we do not find this 'compensation effect' may be that reputation information is not the type of information that owners are looking for. While reputation information did not decrease the ethnic gap, other types of information may succeed in doing so.

Second, and more importantly, we argued that rating systems may fail to overcome inequalities caused by discriminatory tendencies, even in the presence of a compensation effect. While past field experiments convincingly show that ethnic background is less of a determinant of success on online platforms among users with profiles that contain (artificially created) positive ratings than among profiles lacking such ratings, we emphasize that groups that are discriminated against are less likely to obtain positive ratings in the first place, precisely because they are less likely to be accepted. Our simulation model shows that, as a result, the ethnic gap in acceptance chances may actually grow rather than recede within empirically reasonable time spans. Even if ethnic minority users can compensate their initial disadvantage with a good reputation, initial differences between (ethnic) groups may be sustained due to

a process of cumulative advantage. Evidence of similar cumulative advantage processes has been found in a wide variety of fieds, such as education and poverty (Diprete & Eirich, 2006). Our dataset provides a unique opportunity to assess the dynamic process in which initial inequalities may be reproduced, because we observe complete trajectories of acceptance rates of individual users, across all their interactions on the platform. While we find that reputation benefits renters of all ethnicities, we show that with time the reputation system maintains the disadvantage of ethnic minority renters.

This is important knowledge in the light of efforts to mitigate discrimination in the rapidly growing platform economy. The emphasis on an experience of "personal contact" through revealing personal details by peer-to-peer market platforms has the unintended consequence of opening the door to discrimination based on precisely this personal information, and our results indicate that the effectiveness of "institutional" solutions as implemented by platforms to mitigate such discrimination may be limited. In particular, practitioners should be aware of the possibility that reputation systems may exacerbate rather than reduce inequalities between users with different (ethnic) backgrounds. While reputation systems may improve the chances for some individuals to participate on the platform, they may, as an unintended consequence, reduce the chances for others. Platforms that aim to reduce discrimination should not only make their reputation system more effective, they should also reduce initial differences between users with different backgrounds, and help first-time users to acquire their first review, especially if a user belongs to a group that is discriminated against on the platform.

CHAPTER 3

The role of reputation systems in digital discrimination: an experiment

This study is approved by the Ethics Committee of the Faculty of Social and Behavioral Sciences of Utrecht University (FETC19-047).

Abstract

Recent studies suggest that reputation systems may decrease or even eliminate inequality between people with different demographic characteristics, because they decrease the importance of demographic information. However, because not everybody is equally likely to accumulate reviews, the extent to which reputation system reduce inequality may be limited. This study is the first to provide a direct test of differences in inequality between a situation with and without a reputation system. Using an online experiment, I show that inequality between American and Indian individuals is not lower when there is a reputation system than when there is no reputation system. Moreover, emergent inequality in initial rounds affect inequality in later rounds. This implies that platforms that wish to create equal opportunities for users with different backgrounds should not only try to make their reputation system more effective, but also to reduce initial differences in chances for individuals with different backgrounds to participate.

Introduction

Discrimination is a widely spread and persistent problem that is increasingly studied in the context of online peer-to-peer markets. Examples of digital discrimination are lower prices for black sellers on Airbnb (Edelman & Luca, 2014) and eBay (Nunley, Owens, & Howard, 2011). This outcome is negative both for individuals experiencing negative discrimination, as well as for platforms that miss out on potentially fruitful transactions.

Recent studies suggest that reputation systems, which collect, aggregate and distribute feedback about past behavior (Resnick et al., 2000) and that are widely used in online markets, may decrease or even eliminate discrimination (Abrahao et al., 2017; Ert et al., 2016; Tjaden et al., 2018). This claim is based on the finding that initial differences between individuals with different demographic backgrounds disappear once these individuals have a positive reputation. For example, Tjaden et al. (2018) find that the difference in the number of clicks per ride between drivers with a typical German name and drivers with a typical Arab/Turkish/Persian name disappears when the driver has a five-star rating. The mechanism underlying this finding is that reputation information is more relevant and specific to the choice of selecting a transaction partner, thereby decreasing the importance of more diffuse demographic information.

However, these findings at the individual level only tell part of the story of the relation between reputation systems and inequality. By drawing the conclusion that reputation systems reduce discrimination by making demographic information less important, the researchers overlook an important assumption underlying this conclusion, namely that equally trustworthy individuals are equally likely to obtain reviews. However, since reviews can generally only be written after a completed transaction, individuals who have demographic characteristics that are more preferred by others are more likely to be selected for transactions and thus more likely to obtain initial reviews.

Obtaining a first review is important for an individual's chances to participate in future interactions on the platform. As Frey and van de Rijt (2016) show, reputation tends to cascade: individuals who already have a positive review are more likely to be selected for new transactions, and will as a consequence

accumulate more reviews. Random initial differences accumulate over time, resulting in inequality between equally trustworthy individuals. This theory is especially applicable to competitive markets, where choosing one alternative means that another alternative cannot be selected.

If we consider that individuals with different demographic characteristics differ with respect to the probability of receiving a first review, and that a first review is critical for acquiring more reviews, the probability to be selected for a transaction for renters with different backgrounds may diverge over time. We know for example that requests from guests with distinctively African American names on Airbnb are less likely to be accepted than requests with distinctively white names (Edelman et al., 2017). As a consequence, guests with African American names may not only participate in fewer transactions, but they may also accumulate fewer reviews, which may result in a prolonged disadvantage. The results in chapter 2 using data from a peer-to-peer rental platform indeed support this claim.

In the current chapter I study if reputation systems may decrease inequality between individuals with different demographic backgrounds by using an online experiment that is inspired by the experiment of Frey and Van De Rijt (2016) on reputation cascades in competitive markets. Using an experiment allows me to directly compare a situation with and a situation without a reputation system and to draw causal conclusions about the effect of reputation on demographic inequality based on multiple sequences of reputation accumulation. The goal of this study is to evaluate if reputation systems manage to overcome discrimination, or if they may reinforce inequality between different groups.

Theory

Trust and discrimination in online markets

Discrimination is widely studied in the context of labor markets, and online peer-to-peer markets are a relatively new area for discrimination research. These markets are especially suited for the study of trust formation for several reasons. First, in these markets buyers interact with individuals rather than with

businesses, which results in increased risk and uncertainty. Online exchange is characterized by anonymity, and strangers often interact without the prospect of future interactions (Kuwabara, 2015; Parigi et al., 2017). Both buyers and sellers face a risk when trading with unknown others (Macy & Skvoretz, 1998; ter Huurne, Ronteltap, Guo, Corten, & Buskens, 2018). Buyers may have difficulties assessing the value of the goods or services traded, and assessing what the interests of the seller are (Akerlof, 1970). Likewise, sellers may face uncertainty regarding the intentions of the buyer with respect to the payment. Second, buyers and sellers interact through online platforms, rather than face-to-face. This limits the opportunity to get to know the other person before a transaction. Third, peer-to-peer interactions are less strictly regulated than traditional business-to-consumer trade, allowing for more interpretation of what acceptable behavior is (Katz, 2015; ter Huurne et al., 2017).

Platforms apply different methods to solve this trust problem. The majority of platforms tries to reduce anonymity by allowing their users to create a user profile that generally contains their name and a photo, as well as a short personal description. These profiles are found to foster a sense of personal contact (Dubois et al., 2012; Guttentag, 2015) and to provide information that individuals use to assess the trustworthiness of their potential transaction partners. However, an unintended consequence of these profiles is that they may lead to discrimination (Edelman & Luca, 2014; Edelman et al., 2017; Ert et al., 2016; Tjaden et al., 2018). Names and photos convey information about users' demographic characteristics, of which ethnicity and nationality, gender, and age are the most obvious.

Platform users may discriminate on the basis of this information for different reasons. First, users may infer certain unobservable qualities, such as trustworthiness, from these demographic characteristics. Individuals may believe that demographic characteristics correlate with e.g. trustworthiness, and form expectations about the other's trustworthiness based on his or her demographic characteristics. This type of discrimination is referred to as statistical discrimination (Arrow, 1973). This type of discrimination may be reduced by providing better information about the unobserved characteristics of the potential interaction partner. This more accurate and more specific information decreases the informativeness of more diffuse demographic

information and may therefore reduce the reliance on this information (Resnick et al., 2000; Robbins, 2017).

Users may also have different reasons to include demographic information in their assessment of others', without an underlying belief on the relation between demographic characteristics and other merits. This is referred to as 'taste-based discrimination' (Becker, 1957). Generally, people tend to have a preference for others that are similar to them (McPherson, Smith-Lovin, & Cook, 2001). The result of this homophily in groups is that the majority in numbers (with respect to ethnicity, nationality, gender and age) has a relative advantage compared to the minority in numbers.

Using reputation systems to reduce inequality

Reputation systems are proposed to be the most promising solution to statistical discrimination, by researchers (Abrahao et al., 2017; Ert et al., 2016; Tjaden et al., 2018) and practitioners (Murphy, 2016) alike. Most platforms allow their users to give a rating and write a review about their transaction partner after the transaction has been completed. These ratings and reviews are displayed on the targeted user's profile and are visible to potential transaction partners, who may use these reviews to assess the trustworthiness of the reviewed used (Buskens & Raub, 2002).

Reputation information is believed to improve the assessment of the trustee's trustworthiness, and to increase overall trust levels. The reliance on reputation information when assessing whether a trustee can be trusted or not assumes that past behavior of that trustee is related to the trustee's future behavior. Reputation is a costly signal: it is easy to acquire a good reputation for individuals who are sincerely trustworthy, while it is difficult to obtain good reviews for less trustworthy individuals. Individuals who have been trustworthy in the past are believed to be more trustworthy in general, and thus to be more trustworthy in the future. Therefore, individuals who have behaved well in the past can be expected to also behave well in the future, making reputation information valuable for those that are assessing other's qualities. Individuals who have a better reputation (i.e. more and more positive ratings) should be more likely to be selected to an interaction than individuals with a more negative reputation. The positive effect of a good reputation on trust is well established in the existing literature (Boero et al., 2009; Bolton et al., 2004a;

Charness et al., 2011; Duffy et al., 2013; Fehrler & Przepiorka, 2013). At the same time, reputation systems provide an incentive for users to behave well on the platform, since misbehavior in the present limits their opportunities for future transactions (Bolton et al., 2004a; Fehrler & Przepiorka, 2013; Resnick & Zeckhauser, 2002).

Reputation systems have not only proven to increase trust. They are also considered a better and more accurate source of information when assessing another user's trustworthiness on a platform than demographic information. While demographic information can be considered a diffuse type of information with limited informativeness for specific cases, ratings and reviews are very specific to the situation (Resnick et al., 2000; Robbins, 2017). When no reputation information is available, users will rely on the limited demographic information they have, but when they have both types of information at their disposal, they are expected to rely more on the relevant and specific information rather than the more diffuse information. Indeed, previous studies found evidence for this 'compensation effect': the presence of reputation information reduced the importance of demographic information (Abrahao et al., 2017; Ert et al., 2016; Tjaden et al., 2018). Based on these findings researchers concluded that reputation systems are a solution to discrimination in the platform economy.

How reputation systems may increase inequality

However, the difference between individuals with or without positive reviews only tells part of the story. The conclusion that reputation systems reduce inequality assumes that every individual is equally likely to receive their first review. In chapter 2 this assumption was contested, because reviews and ratings can generally only be written after a completed transaction, and not all users are equally likely to get a transaction. People generally prefer others who are similar to them (McPherson et al., 2001). In chapter 2 we showed that rental requests from renters with an minority background on a peer-to-peer motorcycle sharing platform were less likely to be accepted than requests from renters who belong to a majority, and, as a consequence, were less likely to receive a first review.

We also know that getting a first review is of utmost importance for the opportunity to participate in future transactions and for gathering more reviews

(Frey & Van De Rijt, 2016). Those who already have good ratings and reviews are more likely to be selected for future transactions, and may therefore accumulate even more reviews, while those individuals that do not have a reputation yet have a hard time getting both transactions and reviews. In their experiment Frey and Van De Rijt (2016) showed that reputation systems may lead to arbitrary inequality between equally trustworthy trustees. Initial random choices between trustees may give these selected trustees the benefit of building a reputation, while others, who may be equally trustworthy, but who may not get the chance to prove it, do not have this benefit.

In chapter 2 we used simulations to show that when these initial differences are not random but based on the demographic characteristics of the trustee, this initial disadvantage for individuals with a minority background may accumulate when there is a reputation system. This may lead to a reinforcement of the initial disadvantage of individuals with a minority background. Especially when the initial level of trust in individuals belonging to a minority is low, the reputation system may amplify initial differences between groups. In that case individuals belonging to a minority may not get the opportunity to retrieve a first review, which hampers the extent to which they can build a reputation and thus their future opportunities on the platform.

Figure 3.1 summarizes the mechanism discussed in the theory section. In order to draw conclusions about the effect of the presence of a reputation system on inequality, one should not only look at individual decisions, but consider the system as a whole. By only comparing subjects with and without reviews, as most previous papers that look at the effect of reputation on inequality have done, the endogenous nature of the relation between reputation and inequality is overlooked. In the current chapter I focus on national origin discrimination. I conducted an experiment in which American and Indian subjects interacted with each other. I investigate if individuals are indeed more likely to select others who are similar to them, and if this tendency is different when there is a reputation system than when there is no reputation system.

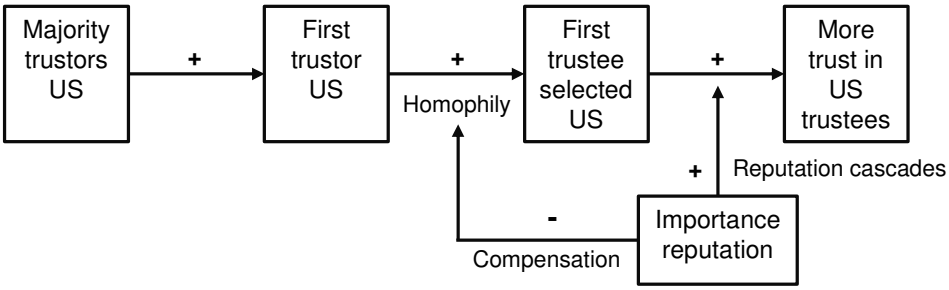

Figure 3.1: Conceptual model

On the one hand, reputation systems may thus decrease the importance of demographic information, thereby decreasing overall discrimination.

> **H1**: The relative difference in trust received between individuals of different nationalities is *smaller* when there is a reputation system than when there is no reputation system.

On the other hand, initial differences between users of different nationality may translate into differences in reputation, which in turn affect future possibilities for interaction on the platform.

> **H2:** The relative difference in trust received between individuals of different nationalities is *larger* when there is a reputation system than when there is no reputation system.

Methods

Design

The current study builds on previous work on the dynamics of reputation and discrimination by using a pre-registered online experiment. The experiment is inspired by the experiment of Frey and Van De Rijt (2016). In their experiment, subjects played a sequence of Trust Games of eight rounds each in groups of eight. The left panel of Figure 3.2 shows the structure of a one-shot Trust Game without a reputation system. In this figure, T_i, R_i, P_i and S_i represent the material payoffs of the actors. The right panel of Figure 3.2 contains the numbers used in the current experiment. Half of the group members played in the role of trustor, the other half in the role of trustee. Subjects played in the

same role throughout the entire game. The trustors took turns: trustor 1 played in round 1, trustor 2 in round 2 and so on. In round 5, it was again the turn for trustor 1. When it is their turn, the trustor could choose to place trust in one of four trustees. When a trustee was selected, they could choose to honor trust or abuse trust. The trustor's payoff was highest (50 points) when they placed trust and when the trustee honored trust, but lowest when the trustee abused trust (0 points). The trustee's payoff was highest when they abused trust (80 points). Trustors that were not active in a round and trustees that were not selected got a payoff of 30 points. When the trustor did not place trust, all players got a payoff of 30 points. At the end of the round all subjects were informed about the points they collected in that round. Under the assumption of rationality, trustors anticipate that selfish trustees will always abuse trust, and will therefore be motivated to select the most trustworthy trustee.

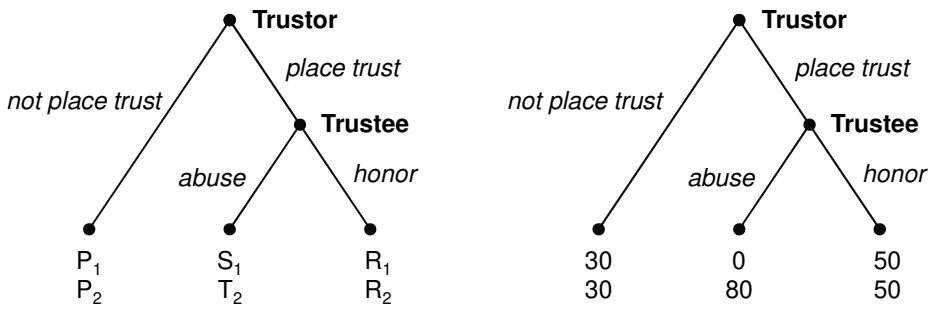

Figure 3.2: Left panel: one round of the Trust Game without reputation system, where $R_1 > P_1 > S_1, T_2 > R_2 > P_2$. Right panel: numerical example used in the current experiment.

Frey and Van De Rijt (2016) observed the emergence of reputation cascades 'that keep increasing the reputational advantage of one party while preventing others from building a reputation.' In their study, trustors had no prior information about the trustees, so the inequality that they observed was based on initial random choices of trustors in early rounds of the experiment. In the current study I extended their design by providing the nationality of the trustees to all group members to evaluate to what extent reputation systems may reinforce differences in outcomes between individuals with a different nationality.

Subjects participated either in the control condition or in the reputation condition. In the control condition, no information about past decisions of the other subjects was available. In the reputation condition full information about the decisions of the other subjects in the session was available. In all groups half of the trustees had the American nationality, and half had the Indian nationality. We expect trustors to have a preference for trustees who have the same nationality as they have. American trustors are thus expected to place more trust in American trustees, while Indian trustors are expected to place more trust in Indian trustees. In order to create a minority-majority situation, the distribution of the nationality of the trustors was counterbalanced across the groups: In half of the games the majority of the trustors had the American nationality, with three of four trustors from the United States, while the fourth trustor was Indian. In the other half of the games the majority of the trustors had the Indian nationality, with three trustors from Indian, while the fourth trustor was from the United States. In groups with more American trustors, we expect American trustees to be advantaged over Indian trustees and vice versa.

Procedures

In September and October 2019, 507 subjects completed the experiment, divided over thirteen sessions that each lasted for about 45 minutes. Subjects were recruited via Amazon Mechanical Turk (MTurk, 2019)[10]. Compared to a student sample, there is more demographic variation among workers on MTurk. I made use of the fact that the large majority of MTurk workers is from the US, followed by India (Difallah, Filatova, & Ipeirotis, 2018). Only MTurk workers aged 18 or above with one of these nationalities, who had completed 1000 HITs or more and with a HIT acceptance rate of 90% or higher could

[10] Each session was scheduled at a fixed date and time. Subjects could sign up for a specific session through a link that was posted on MTurk. Subjects who finished the instructions, but who could not be assigned to a group received a show-up fee of $2.00. In case a subject quit the experiment in the middle of a game, they did not receive a payment and the remaining group members received the points that they earned so far, plus the points they would have earned in the remaining rounds, based on their average number of points earned per round. In order to pay the subjects according to their earnings in the game, we had to collect their MTurk IDs. After completing the cleaning of the data, these IDs are stored separately from the experimental data.

participate in the experiment. Subjects could only participate once[11]. 41.8% of the subjects was female. 49.5 % had the Indian nationality and the remaining 50.5% had the American nationality. The average age of the subjects was 36.5 years (*SD* = 10.6). Most subjects held a high school degree (17.8%), undergraduate (28.8%) or a graduate degree (43.2%). Two pilot sessions were conducted. On average, subjects earned 6.96 dollar (256 points per game). All sessions were computerized using oTree (Chen, Schonger, & Wickens, 2016).

At the beginning of the study the participants were asked to fill out a short survey asking about their age, gender and nationality. They were informed about the duration, payment, and general procedures of the experiment and they were asked to give their consent for displaying their age, nationality and gender to other participants (see the Appendix for the full instructions). Their identity was never revealed, and no deception was used in the experiment. Subjects received on-screen instructions and could only proceed to the games if they correctly answered five comprehension questions to ensure that they carefully read the instructions. Six subjects who failed to reach a decision within the time limit (two minutes) during the game were excluded from the remainder of the experiment. The data from the groups in which this happened are excluded from the analyses. The subjects of the experiment played one or two games of eight rounds each. The nationality of the trustees was shown to all subjects in the group.

Variables

The main outcome variable is the relative inequality in the average level of trust placed in American and Indian trustees. Per group I calculated the fraction of trust placed in Indian trustees, which is the total number of times trust was placed in Indian trustees divided by the total number of times trust was placed in a group. If there is no discrimination, the expected fraction of trust placed in American trustees is the number of American trustees in a group divided by the total number of trustees in the group. Considering that in each group two of the trustees were American and the other two were Indian, on average 50% of trust would be placed in Indian trustees when trustors would randomly select one of the trustees, without taking their nationality into account. The more the

[11] Four subjects completed the study more than once. The data of the groups that they were part of in their second session are excluded from the analyses.

fraction of trust placed in Indian trustees deviates from 0.5, the higher the inequality, so inequality is calculated as the absolute difference between the fraction of trust placed in Indian trustees and 0.5. This difference is equal to half the value of the two-person Gini-coefficient[12]. The larger this difference, the more inequality there is in a group. This measure is robust to differences in discriminatory behaviors. It can be applied to cases in which trustors select trustees that have the same nationality as they have (homophily), as well as to cases in which trustors have other preferences. Ten groups in which trust was placed never or only once were excluded from the analyses.

Results

Descriptive statistics

Table 3.1 shows the average trust level for American and Indian trustees and the average inequality per group in the different conditions. Table 3.2 shows per condition the average trustworthiness level and success rate (the fraction of interactions in which trust is placed and honored) for American and Indian trustees.

[12] The inequality measure used here is $0.5 - S_2 / (S_1 + S_2)$, where S_1 is the frequency with which trust is placed in American trustees and S_2 is the frequency with which trust is placed in Indian trustees. When this measure is multiplied by two, the result is $1 - 2S_2 / (S_1 + S_2)$, which can also be written as $(S_1 + S_2)/(S_1 + S_2) - 2S_2 / (S_1 + S_2)$, which is equal to the two-person Gini coefficient: $(S_1 - S_2) / (S_1 + S_2)$.

Table 3.1: Average trust rate per group for Indian and American trustees, inequality rate per group, by condition

Condition		N	Trust rate US trustees	Trust rate Indian trustees	N	Inequality
Control	*Majority trustors US*	29	32.8 (17.2)	21.1 (17.7)	28	0.18 (0.17)
	Majority trustors India	28	23.2 (15.9)	23.7 (17.5)	24	0.24 (0.18)
	Total	57	28.1 (17.1)	22.4 (17.5)	52	0.21 (0.18)
Reputation	*Majority trustors US*	32	29.3 (23.9)	26.6 (25.9)	29	0.25 (0.18)
	Majority trustors India	28	31.7 (23.9)	24.1 (25.0)	26	0.27 (0.19)
	Total	60	30.4 (23.7)	25.4 (25.3)	55	0.26 (0.18)
Total		117	29.3 (21.7)	23.9 (21.8)	107	0.23 (0.18)

Note: Standard deviations in parentheses.

Table 3.2: Average trustworthiness and success rate for Indian and American trustees per group, by condition

Treatment		N	TW rate US trustees	N	TW rate Indian trustees	N	Success US trustees	Success Indian trustees
Control	*Majority trustors US*	28	60.7 (42.5)	25	26.4 (36.3)	29	23.3 (21.8)	6.9 (11.4)
	Majority trustors India	24	63.5 (42.9)	22	35.4 (37.2)	28	14.7 (14.9)	8.9 (10.7)
	Total	52	62.0 (42.3)	47	30.6 (36.6)	57	19.1 (19.1)	7.9 (10.9)
Reputation	*Majority trustors US*	27	51.6 (43.7)	25	41.4 (42.9)	32	19.9 (24.8)	15.2 (22.6)
	Majority trustors India	23	71.4 (37.3)	21	26.7 (40.7)	28	25.0 (24.3)	9.8 (21.1)
	Total	50	60.7 (41.7)	46	34.7 (42.1)	60	22.3 (24.5)	12.7 (21.9)
Total		102	61.4 (41.8)	93	32.6 (39.3)	117	20.7 (22.0)	10.4 (17.5)

Note: Standard deviations in parentheses.

All reported T-tests are conducted at the level of the group. An independent samples T-test shows that there is neither a difference in the average trust rate between groups in the control condition and groups in the reputation condition ($t(115) = 1.078$, $p = 0.283$), nor a difference in trustworthiness rate ($t(112) = 1.033$, $p = 0.304$). This suggests that the reputation system did not lead to an overall increase in trust and trustworthiness, but this does not imply that reputation did not affect decision making at the individual level, as will be shown in the exploratory analyses. American and Indian trustors are equally likely to place trust ($M_{US}= 54.3$, $M_{India}= 55.1$, $t(116) = 0.194$, $p = 0.847$). The average fraction of rounds in which trustees are trusted across all conditions does not significantly differ between American and Indian trustees ($t(116) = 1.76$, $p = 0.081$), but American trustees are on average more likely to honor trust ($M_{US}= 57.6$, $M_{India}= 31.5$, $t(80) = 4.109$, $p < 0.001$). Overall, this results in a much higher success rate for American trustees than Indian trustees ($t(116) = 3.675$, $p < 0.001$), but not in a difference in the average payoff of American trustees and Indian trustees in a group ($M_{US}= 34.4$, $M_{India}= 34.6$, $t(106) = -0.274$, $p = 0.785$), because abusing trust yields trustees a higher (immediate) payoff than honoring trust.

Figure 3.3 shows how inequality is distributed across groups in the different conditions. The bar in the middle indicates that the trustors in a group selected American and Indian trustees with the same frequency. Groups represented by bars more to the left placed more trust in Indian trustees than in American trustees. The bars more to the right represent groups that placed more trust in American trustees than in Indian trustees. The presence of a reputation seems to reduce the fraction of groups in which trust is equally divided between Indian and American trustees in groups where the majority of the trustors is American, but not in groups with an Indian trustor majority.

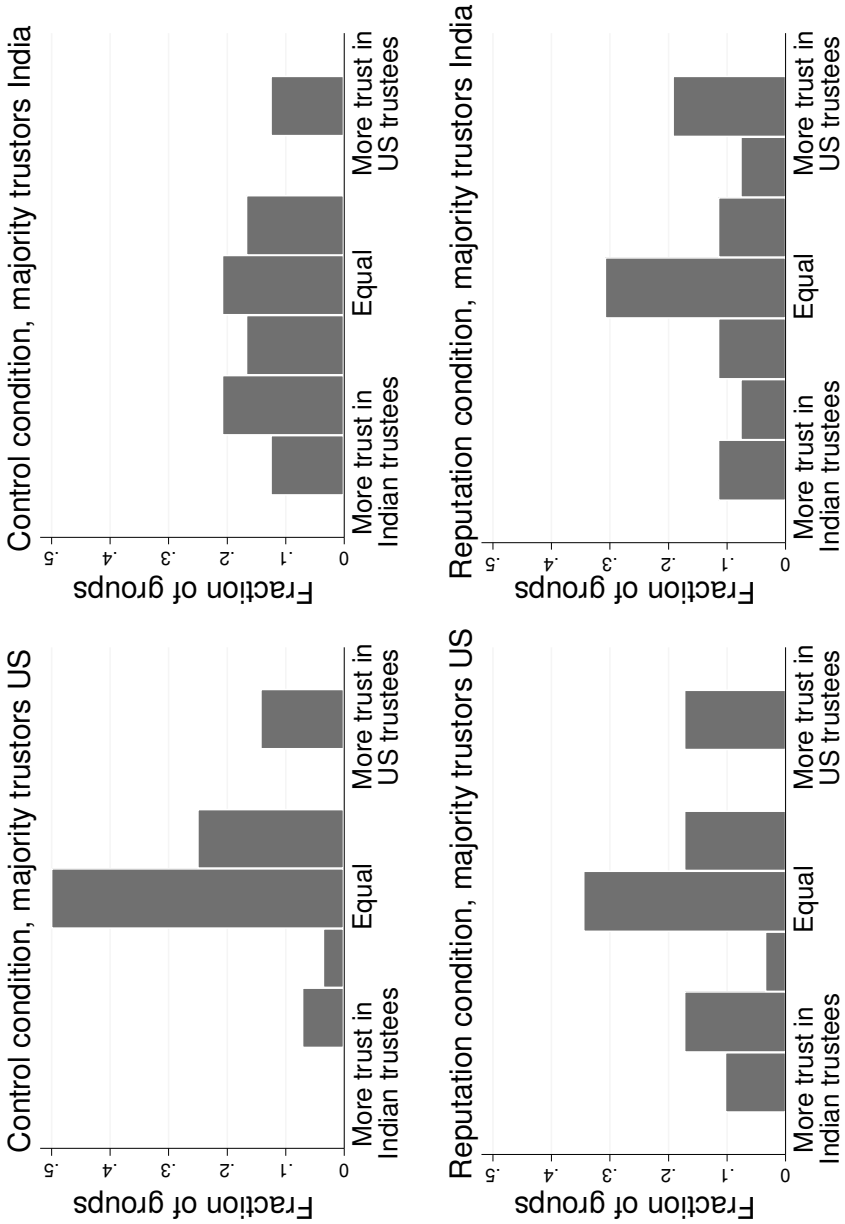

Figure 3.3. Distribution of inequality across groups in the different conditions.

Hypothesis test

An independent samples T-test at the level of the group reveals that the difference in inequality in the presence of a reputation system compared to the control condition is neither significant when the majority of trustors is American ($t(55) = 1.629, p = 0.109$), nor when the majority of trustors is Indian ($t(48) = 0.470, p = 0.640$). I cannot reject the null hypothesis that there is no difference between the control condition and the reputation condition[13].

Figure 3.3 suggests that trustees belonging to the minority benefit more from the reputation system. A T-test on the difference in the fraction of trust placed in Indian trustees between the control condition and the reputation condition reveals that this difference is neither significant in groups with an American majority ($M_{\text{Control}} = 0.36$, $M_{\text{Reputation}} = 0.44$, $t(55) = -1.163$, $p = 0.250$), nor in groups with an Indian majority ($M_{\text{Control}} = 0.51$, $M_{\text{Reputation}} = 0.41$, $t(48) = 1.194$, $p = 0.239$).

Robustness checks

There are various ways to quantify inequality. I tested if the results of the main hypothesis test are robust to using other commonly used inequality measures. The results change neither when using the Herfindahl index ($M_{\text{Control}} = 0.65$, $M_{\text{Reputation}} = 0.70$, $t(105) = -1.458$, $p = 0.148$), nor when a using the logged difference in the number of rounds that Indian and American trustees were expected to be selected, obtained from a negative binomial regression of the counts of trust placed in each trustee ($M_{\text{Control}} = 4.06$, $M_{\text{Reputation}} = 5.81$, Wilcoxon rank-sum test: $z = -1.299, p = 0.194$).

[13] Because every subject played one or two games, the groups are not completely independent. Moreover, the composition of some groups changes between the two games, while other groups remain the same. A regression with inequality as the dependent variable and with random intercepts for unique groups reveals that adding these intercepts does not significantly improve the model fit ($\chi^2(1) = 0.37, p = 0.270$). Adding a random intercept for each session also does not improve the model fit ($\chi^2(1) = 0.12, p = 0.366$). This suggests that no significant fraction of variance is explained at the level of unique groups and sessions, so in the main text the results from the regressions without these random intercepts are reported.

The number of rounds in which trust is placed differs across groups. In groups in which trust is placed in an uneven number of rounds, Indian and American trustees cannot be trusted an equal number of times. The number of rounds in which trust is placed thus affects the range of values the inequality variable can take. To account for this, a corrected inequality variable has been constructed. For groups in which trust was placed an even number of times there is no difference between the original and the corrected inequality measure. For groups in which trust was placed an uneven number of times, the corrected inequality variable took a value of 0 if the fraction of trust placed in Indian trustees could not be closer to 0.5. For example, in groups in which trust was placed five times, the corrected inequality variable was 0 when Indian trustees were trusted two or three times. When the fraction of trust placed in Indian trustees did not fall in this range, corrected inequality was calculated as the distance between the fraction of trust placed in Indian trustees and the closest boundary of that range, divided by the size of the range of values that were considered to be unequal (2 * 0.4 in the case of groups in which trust was placed in five rounds). The results do not change when using this new corrected inequality variable instead of the original inequality variable: there is no significant difference in inequality between two conditions ($M_{\text{Control}} = 0.18$, $M_{\text{Reputation}} = 0.23$, $t(105) = -1.207$, $p = 0.230$).

The inequality measure that is used in the main hypothesis test is based on all rounds in which trust was placed in a group. This is a rather conservative calculation of inequality, because it takes time for individuals to build a reputation, so the effect of the reputation system on inequality is expected to be stronger in later rounds. When constructing the inequality variable without the first round in which trust is placed in a group, the difference between the control condition and the reputation condition is not significant ($M_{\text{Control}} = 0.26$, $M_{\text{Reputation}} = 0.30$, $t(105) = -1.252$, $p = 0.213$). When also the second round in which trust is placed is excluded, there is significantly more inequality in groups in the reputation condition ($M_{\text{Control}} = 0.20$, $M_{\text{Reputation}} = 0.32$, $t(82) = -2.587$, $p = 0.011$), but when also the third round in which trust is placed there is again no difference between the conditions ($M_{\text{Control}} = 0.37$, $M_{\text{Reputation}} = 0.29$, $t(71) = 1.687$, $p = 0.096$). When excluding more rounds, the number of observations is not sufficient to draw any conclusions.

Reputation cascades are not expected to emerge when trust is abused. When only including rounds in which trust was placed and honored in the calculation of inequality within groups, inequality is higher than when also including rounds in which trust was not honored (mean difference = 14.4 percent point, $t(101) = 4.025$, $p < 0.001$), suggesting the cascading of reputation led to inequality between trustees of different nationality. Average inequality between trustors who honored trust is 0.35 in the control condition and 0.41 in the reputation condition. This difference is not significant at the 95% confidence level ($t(86) = 1.551$, $p = 0.125$).

Overall, these results suggest that the finding that reputation systems do not reduce inequality is robust to using different measures of inequality.

Exploratory analyses

To get a better understanding of the findings, I explored to what extent the mechanisms outlined in the theory section were at play in the experimental groups. First, I explore to what extent trustees belonging to the minority benefit more from reputation than trustees belonging to the majority because the presence of more relevant and specific reputation information may reduce the importance of more diffuse demographic information. This is the main argument of studies that propose that reputation systems may eliminate discrimination (Figure 3.1, bottom left part. e.g. Ert, Fleischer, & Magen, 2016). To test if the experimental data support this argument, I fitted McFadden's choice model for the control condition and the reputation condition to analyze how the reputation of the trustee in the reputation condition affects the decision of the trustor. This model allowed me to study how the choice of the trustor depends on characteristics of the trustees, as well as on variables that do not vary across trustees, such as the round number and the majority condition. Table 3.3 contains the results of these analyses. The models are constructed separately for American trustors (Models 1 to 4) and Indian trustors (Models 5 to 8). Table 3.3 only contains the outcomes for the variables of interest, the control variables are omitted. The full table can be found in the Appendix, Table A3.1.

Table 3.3: Results conditional logit models with the choice of the trustor as dependent variable. Logged odds. Estimates for the control variables are omitted from this table. A table with full results can be found in the Appendix, table A3.1.

	American trustor				Indian trustor			
	Control	**Reputation**			**Control**	**Reputation**		
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Main effects								
Trustee India	-0.77***	-0.82***	-0.87**	-0.84***	0.15	0.26	0.25	0.11
	(0.22)	(0.22)	(0.30)	(0.23)	(0.21)	(0.24)	(0.29)	(0.25)
Number of times trustee honored trust	-	0.81***	0.76**	0.81***	-	0.30*	0.29	0.28
		(0.16)	(0.25)	(0.17)		(0.14)	(0.21)	(0.15)
Number of times trustee abused trust	-	-2.70***	-2.72***	-2.83**	-	0.42	0.42	-0.30
		(0.71)	(0.72)	(0.89)		(0.22)	(0.22)	(0.54)
Interactions								
Trustee India * Number of times trustee honored trust	-	-	0.10	-	-	-	0.02	-
			(0.32)				(0.30)	
Trustee India * Number of times trustee abused trust	-	-	-	0.47	-	-	-	0.97
				(1.32)				(0.54)
Number of decisions	480	548	548	548	440	524	524	524
Number of trustors	120	137	137	137	110	131	131	131

Note: Standard error in parentheses. * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

The choice of a trustor in a round for one of four trustees was the outcome variable. Rounds in which no trustee was selected were excluded. For each trustee, their nationality was included as independent variable. For trustees in the reputation condition I also included the number of times they honored and abused trust in the past. I also included the game number and round number as control variables, and I added random intercepts for trustors, because they make multiple decisions across the game(s) and these decisions are not independent. Models 1, 2, 5 and 6 contain the outcomes of these analyses. The remaining models include the interactions between the reputation and the nationality of the trustee to test if the presence of reputation information reduces the reliance on the nationality of the trustee.

Supporting the findings from the Figure 3.3, models 1 and 5 show that American trustors place more trust in American trustees, while Indian trustors do not condition their decision on the nationality of the trustee in the control condition. Model 2 shows that American trustors place more trust in trustees who have honored more trust in the past, and less trust in trustees who have abused trust in the past. Model 6 shows that Indian trustors place more trust in trustees who have honored more trust in the past, but not less trust in trustees who have abused more trust in the past, compared to trustees who have neither honored nor abused trust in the past. The remaining models in Table 3.3 show that none of the interactions between the trustee's nationality and their reputation affects the decision of the trustor. This means that there is no compensation effect: trustees who are disadvantaged because of their nationality cannot compensate this disadvantage with a good reputation. These findings do not support the argument that the presence of more relevant and specific information reduced the reliance on demographic information.

Second, I explored how choices of individual trustors and trustees led to the outcomes at the group level. The bottom part of Figure 3.3 shows that in the reputation condition in some groups more trust is placed in American trustees, while in other groups more trust is placed in Indian trustees. I conducted a number of analyses to explain what factors determine which group of trustees is advantaged over the other.

The first mechanism outlined in the theory section is homophily (see the middle-top part of Figure 3.1): trustors are more likely to select a trustee that

has the same nationality as they have. The results from the conditional logit models in Table 3.3 already confirmed that this is the case for American trustors, but not Indian trustors. When there is a reputation system, the first round in which trust is placed in a round is expected to affect overall inequality because initial differences may be reinforced through accumulation of reputation. A proportions test confirms that also when trust is placed for the first time in the reputation condition, American trustors are more likely to select an American trustee (American trustees: 79%, $z = 3.024$, $p = 0.003$), while Indian trustors do not differentiate between American and Indian trustees (Indian trustees: 52%, $z = 0.186$, $p = 0.853$).

The second important mechanism that led to the prediction that a reputation system may not be an effective instrument to decrease inequality is reputation cascading (bottom-right part Figure 3.1). I analyzed the prevalence of reputation cascades in the reputation condition by estimating the probability that a cascade was continued. Consistent with the analysis by Frey and Van De Rijt (2016) I used a logistic regression with random intercepts for trustors and the variable indicating if the trustor selected the same trustee as was chosen on the last turn the trustor observed. Observations were excluded if no previous turn was observed, or if no trust was placed, or was abused in the previous round. If trustors would randomly select a trustee, the probability that a trustor selects the same trustee as was chosen on the last turn the trustor observed should be 25%, since there were four trustees in every group. In the control conditions this probability is indeed not significantly different from 0.25 (27.6%, $z = 0.430$, $p = 0.667$). The observed probability is 52.5% in groups with an American trustor majority, which is significantly higher than 0.25 ($z = 3.554$, $p < 0.001$). This means that in these groups the presence of a reputation systems led to reputation cascades. In groups with an Indian majority the observed probability is not significantly higher than 0.25 (31.3%, $z = 0.828$, $p = 0.408$), so we do not find evidence of reputation cascades in those groups. The probability that a cascade is continued in the groups in which the majority of the trustors is American is about 11.8% higher when the cascade length increases with one ($z = 2.08$; $p = 0.037$). The nationality of the reputable trustee does not affect the cascade continuation probability in those groups (mean difference = 15.8 percent point; $z = 1.38$; $p = 0.168$).

Next, I used a regression analysis at the group level to test if a higher likelihood of cascade continuation is related to more inequality. I included inequality in a group as the dependent variable and the probability that a cascade is continued as an independent variable. I controlled for treatment (control or reputation), trustor majority (US or India) and the game number (first or second). I find evidence that there is more inequality in groups with a higher likelihood of cascade continuation ($b = 0.136$, $t = 2.75$, $p = 0.008$).

The presence of reputation cascades leads to the prediction that the nationality of the first trustee who honored trust in a group has a large influence on the choices of future trustors (top-right part Figure 3.1). Figure 3.4 shows the relation between the nationality of the first trustee that honored trust and the nationality of the trustee chosen in subsequent rounds. A T-test at the group level shows that the difference in inequality between American and Indian trustees in the remaining rounds is significantly affected by the nationality of the first trustee that honored trust (mean difference = 0.336, $t(44) = 3.489$, $p = 0.001$). When the first trustee that honored trust is Indian, the probability that future trustors also select an Indian trustee is higher than when the first trustee that honored trust is American.

Together these findings suggest that in some cases trustees belonging to the majority benefit more from the reputation system, while in other cases trustees belonging to the minority benefit more. Which group will benefit more depends on the nationality of the first trustee who honored trust in a group. The selection of this first trustee depends on the nationality of the first trustor that places trust in a group, which in turn depends on the composition of the group: the probability that a trustor belonging to the majority is the first to place trust (top-left part Figure 3.1, 69.3% ($t(113) = 4.448$, $p < 0.001$). These results support the theory that differences in group composition, homophily, and cascading of reputation may lead to reinforcement of existing inequalities.

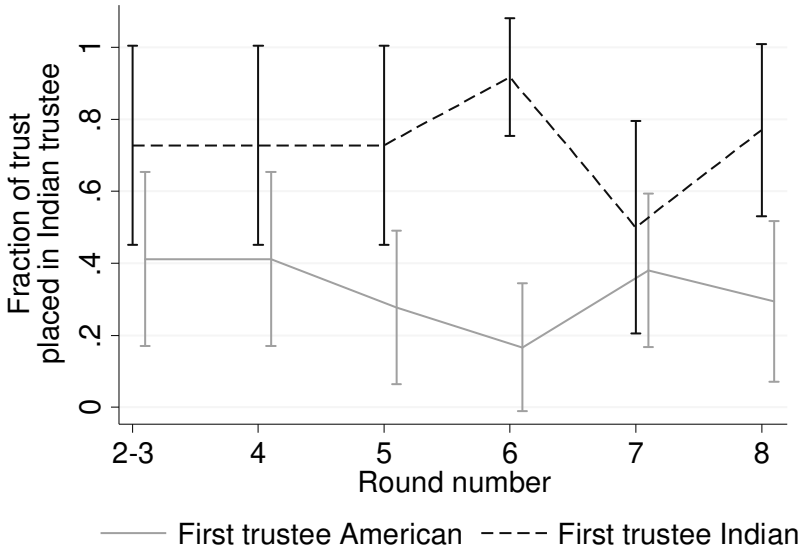

Figure 3.4. Relation between the nationality of the first trustee that honored trust in a group and fraction of trust placed in Indian trustees in the remaining rounds. Groups in which trust was placed never or only once are excluded.

Discussion

Reputation systems are often argued to be the most promising solution to discrimination in online platforms, because the presence of ratings and reviews may decrease the importance of demographic information (Abrahao et al., 2017; Ert et al., 2016; Tjaden et al., 2018). However, since not everybody may be equally likely to obtain these ratings and reviews, it is possible that reputation systems reinforce existing inequalities. In the presence of a reputation system, some individuals will accumulate reviews, while others do not. If individuals who belong to the national majority structurally get more opportunities for reputation building, their initial advantage may accumulate and result in an even larger advantage. In an online experiment, I found that the presence of a reputation system failed to decrease inequality between individuals of different nationality. This finding is robust to using different measures of inequality. Moreover, decisions made in initial periods determine to a large extent which individuals will be trusted more later on.

The argument that reputation systems reduce existing inequalities is based on the idea that reputation is a better indicator of another user's trustworthiness than more diffuse personal characteristics such as nationality. Contrary to findings of earlier experiments (Abrahao et al., 2017; Ert et al., 2016) and field studies (Tjaden et al., 2018), I did not find evidence that the presence of reputation information reduces the importance of demographic information. Neither American nor Indian trustors differentiate between reputation of Indian and American trustees in the experiment.

Although the difference in inequality between the control condition and the reputation condition is not significant, the exploratory analyses revealed that the mechanisms that led to the prediction that reputation system may increase inequality worked as predicted. The difference in the probability to be trusted between Indian and American trustees in the remaining rounds strongly depended on the nationality of the first trustee that honored trust in a group, because reputation tends to cascade. The higher the likelihood that a cascade is continued in a group, the larger the inequality in that group.

An explanation for why these mechanisms together did not lead to an overall difference in equality may be that the theory assumes that all trustees are equally likely to honor trust, while in the experiment American and Indian trustees did not honor trust at the same rate. One explanation for this latter finding may be that Indian trustees have a smaller incentive to honor trust because they are less likely to be selected, and therefore have a shorter shadow of the future. To avoid these differences in behavior between different users, their responses could be simulated. However, using simulations would require deception of the subjects and would lower the external validity, because in real markets there may also be differences between trustees. Another explanation may be that the observed trust levels are relatively high. The simulation study in chapter 2 showed that reputation systems are especially likely to increase inequality when the initial level of trust in users belonging to a minority is low. A last explanation may be that the number of observations per condition is limited, because inequality is measured at the group level.

This study is the first to compare the effect of the presence of a reputation system on inequality between individuals. The experimental setting allowed for a comparison of the situation in which there is a reputation system with a

control condition without a reputation system. This would not be possible when using data from real platforms, since they either have or do not have a reputation system, and there are too many differences between platforms to make a comparison across platforms. The subjects who participated in the experiment likely behaved differently than real platform users. They may have responded differently to in- and outgroup members, and they may have attached different value to reputation information. However, this does not mean that we cannot draw any conclusions about real-life platforms from the results. While the preferences and behaviors of individual subjects may be different, the theory is not about these preferences per se, but rather about the dynamics leading to reinforcement or a decrease of inequality. These dynamics are a product of individual behavior. Following the simulation results in chapter 2 we hypothesized that reputation systems may lead to an increase in inequality if, and only if, subjects strongly rely on reputation information and if initial differences between individuals with different demographic characteristics are large enough.

Overall, the experimental results suggest that reputation systems do not reduce inequality between individuals with different demographic backgrounds. Reputation systems help those that obtained at least one review but may negatively affect those individuals that have not participated in interactions before. Given that it is more difficult for individuals with certain characteristics than for others to obtain reviews, reputation systems may not be an effective means to reduce inequality between individuals with different demographic characteristics. The common use of (ranking) algorithms that favor individuals with a better reputation, such as used on Airbnb (Fradkin et al., 2018) may further reinforce existing differences. Platforms that wish to create equal opportunities for users with different backgrounds should not only try to make their reputation system more effective, but also to reduce initial differences in chances to participate for individuals with different backgrounds.

CHAPTER 4

Reputation systems, communication and the identification of liars

This chapter has been co-authored by Rense Corten and Arnout van de Rijt.

Both studies are approved by the Ethics Committee of the Faculty of Social and Behavioral Sciences of Utrecht University (Study 1: FETC-20-365; Study 2: FETC-20-194).

Abstract

Buyers in online auctions have been found to pay higher prices to sellers who promise to deliver a high-quality product, even though there are no guarantees that they will keep this promise. To explain this, past studies have hinted at the possibility that reputation systems make this form of 'cheap talk' credible, by giving buyers the opportunity to publicly rate sellers' past honesty, and sellers the ability to build a reputation for being honest. We provided an elaborate discussion of the theoretical mechanisms underlying this claim and we tested this hypothesis using both observational data from online auctions on eBay and an online experiment. In both studies we found that reputation information alone is not enough to convince buyers to pay a high price. Surprisingly, in neither study did we find that buyers trust quality claims from sellers more if the latter have a better reputation, while combining reputation information and communication could have improved buyers' decisions. Moreover, in the experimental study we did not find that the presence of a reputation system increases buyers' overall willingness to trust sellers' quality claims. The data thus provide no support for the theory that reputation systems help make cheap talk credible. The puzzle why buyers trust promises made by sellers remains.

Introduction

Online markets for used goods are much like Akerlof's (1970) "market for lemons". They are characterized by a large heterogeneity in the quality of products sold. Sellers have better information about the quality than buyers. This information asymmetry thus poses a risk of *adverse selection* that drives high quality products out of the market (Akerlof, 1970). Buyers may not be willing to pay a high price for a product if they are not certain about its quality. Sellers in turn may not wish to provide a high quality good for a low price. Sellers will then only provide low-quality products, and buyers will only pay low prices. And yet empirical research on online auctions shows that products promised to be of high quality are sold for high prices (Anderson et al., 2007; Lewis, 2011). How do platforms manage to prevent quality degredation of their markets?

At the beginning of an online auction, a seller can cheaply communicate details about product quality to interested buyers. Game-theoretic analysis would predict that these descriptions should not affect buyers' decisions if sellers and buyers meet only once, since product details provided need not be accurate or honest (Akerlof, 1970; Duffy & Feltovich, 2006). However, empirical research on online auctions shows that a product description written by a seller significantly increases the willingness of buyers to pay a high price (Anderson, Friedman, Milam, & Singh, 2007; Lewis, 2011; Rawlins & Johnson, 2007). Moreover, sellers voluntarily report deficits of the product they are selling (Anderson et al., 2007; Lewis, 2011), even though they may expect this to negatively impact price.

In the current chapter we review past cursory arguments that suggest that sellers can be trusted because they care about their reputation, but do not specify exactly why. We synthesize these past arguments into an account that draws on signaling theory to explain how reputations may lead rational and selfish buyers to trust seller promises. We test this explanation in an observational study of eBay data as well as in an online experiment.

Signaling theory states that information about the intentions of the sellers could help buyers distinguish between sellers who intend to provide a high-quality good and sellers who do not intend to do that (Gambetta, 2009; Przepiorka &

Berger, 2017; Przepiorka & Diekmann, 2013; Raub, 2004). Signals can be cheap or costly. Costly signals are type separating, i.e. sellers who intend to provide a high-quality good will produce the signal, while sellers who intend to send a low-quality good do not produce the signal (Przepiorka & Berger, 2017). Cheap signals on the other hand are not type-separating and can easily be faked by sellers who do not intend to provide a high-quality good. On online auction websites, sellers can inform interested buyers about the quality of their products at trivial cost, by providing a description of the product. As there is no way to find out if sellers tend to speak the truth, for example through sustained cooperation between two individuals, the information can be considered cheap talk, and buyers are not expected to use the information.

Previous studies have suggested that the presence of a reputation system may attach a cost to the otherwise cheap signals that product descriptions constitute (Ba & Pavlou, 2002; Bolton et al., 2004a; Houser & Wooders, 2006; Lewis, 2011; Neto et al., 2016; Snijders & Zijdeman, 2004). Platforms, including those for which product descriptions have been found to affect buyers' decisions (Anderson et al., 2007; Lewis, 2011), do indeed have reputation systems through which buyers can rate and review sellers. However, through what theoretical mechanism reputation systems and product descriptions would interact to allow quality products to sell at high prices has remained unclear. Moreover, we are not aware of any empirical study in which this hypothesis is tested.

In a nutshell, we theorize that reputation systems for online auctions communicate honesty more than quality. The quality of a product may be low, but a buyer will not complain about it unless it was unexpected. On eBay, reviews may contain phrases like: 'Just what I ordered!' or 'Item description was very accurate'. Buyers are likely to be more satisfied with their purchase if the product matches the product description than when the seller lied about the quality of the product. `More satisfied buyers will provide better reviews, and good reviews are important for sellers, because it increases their opportunity to sell their products at high prices in the future.

By providing the best quality possible, and by communicating about inferior quality products honestly, sellers can thus keep up a good reputation. For example, a seller selling a used car online may not have a car available that meets

the criteria of the buyer. If the seller in that case informs the buyer about the inferior quality, this may prevent the buyer from buying a car of inferior quality for a high price. In the short run honest communication may be costly to the seller, because sellers will likely receive a lower price if they communicate that their product is of low quality. However, sellers who currently only have low-quality products may have high-quality products on offer in the future. If they have a reputation for being honest, future buyers may be more inclined to trust that the product on offer actually has the quality as described by the seller. On the other hand, descriptions provided by sellers who deceived buyers in the past will not be believed. Honest sellers may thus receive higher prices in the future.

In the current chapter we use both observational data and an online experiment to study the interplay between reputation systems and promises made by sellers. The observational data we draw on were collected by Lewis (2011, 2019)(Lewis, 2019) and contain information on a large number of auctions completed on eBay. The online experiment allows a controlled test of the argument by systematically varying the presence of a reputation system and communication.

Theory

Sellers who sell second-hand products in auctions sometimes have high-quality products and sometimes only low-quality products. Buying second-hand products online involves several risks for buyers, due to intentional or unintentional defection by sellers (Axelrod & Dion, 1988). First, the quality of most products deteriorates over time, so the quality of used products will vary both between and within sellers. Second, after the buyer has paid for the product, the seller may not send the product, send it too late or send an inferior product. A classic example of this situation is Akerlof's (1970) market for second-hand cars in which buyers do not know if they are buying a good or a bad car. In many of these markets, sellers and buyers are strangers to each other and only interact once, and they have no incentive to invest in a good relationship. This lack of dyadic embeddedness is also typical for the modern-day exchanges that are mediated by online platforms.

The prices buyers are willing to pay depend on product quality: They are willing to pay a higher price for a high-quality good than for a low-quality good. They want to avoid paying a high price for a product of inferior quality. Both buyers and sellers benefit more when a high-quality product is sold for a high price, than when a low-quality product is sold for a low price. However, without the prospect of future interactions, rational and selfish sellers have interests that are opposed to those of buyers: they benefit more when they provide a low-quality product for a high price. In those markets, buyers are not willing to pay a high price for a product if they are not confident about its quality, and sellers will not provide a high quality good for a low price. This situation leads to adverse selection: it drives high quality products out of the market (Akerlof, 1970). To achieve better market outcomes, it would thus be beneficial if buyers could tell the difference between high- and low-quality products.

In online auctions, the final price is determined by the highest bidder. eBay is the largest platform for online auctions. Both individuals and businesses can sell their products via this platform. Sellers can list their products, choose a duration for their auction to run, set a starting price and provide a description of the products. Sellers can also add a reserve price, which is the minimum price for which the product will be sold. These reserve prices are not communicated to bidders (eBay, 2020b), so they do not function as a signal of quality. Bidders make bids that are publicly available. If the highest bid at the end of the auction is higher than the seller's reserve price, the buyer with the highest bid pays the price they offered and receives the product (eBay, 2020a).

Cheap and costly signals

Credible information may help a buyer estimate the probability that a seller will provide a high quality good. Different types of information can be communicated to buyers; here we distinguish cheap and costly signals. Costly signals are type separating, i.e. sellers who intend to provide a high-quality good will produce the signal, while sellers who intend to send a low-quality good do not produce the signal (Gambetta, 2009; Przepiorka & Berger, 2017). When the net benefits of producing a signal are positive for sellers who intend to provide a high-quality good, but negative for sellers who do not have that intention, a signal is type-separating. In that case only sellers who really intend to send a high-quality product will produce the signal, while other sellers will not produce

the signal. When a signal is not type-separating, it can easily be faked, and it is considered a cheap signal. The net benefits of producing a signal depend on the costs of producing the signal and on the expected (future) benefits of sending that signal.

Communication: information about intentions

One way for sellers to convince buyers of the high quality of their products, is to communicate this to buyers, for example through product descriptions. However, these descriptions are non-binding and can be produced at trivial cost by sellers who intend to provide a high-quality good and by sellers who intend to send a low-quality good alike. Hence, a description need not reflect a seller's true intentions and should be considered a cheap signal.

When a seller promises to provide a low-quality product, the best response for the bidder is to offer a low price, to avoid paying a high price for a low-quality good. When a seller promises to deliver a high-quality product, the best response of the bidders depends on whether they believe that the seller will really provide a high-quality product. If they believe the message, the best response is to offer a high price. Otherwise they can best offer a low price. In the short term, sellers benefit more when buyers offer higher prices. This means that there is no reason for any rational seller to describe the product as a low-quality product, since they can in that case be sure that the buyer will not pay a high price. Under the assumption of rationality and in the absence of a reputation system, sellers will always promise that they will provide a high quality good, so product descriptions are cheap signals.

However, empirical research shows that these product descriptions affect buyers' decisions (Anderson et al., 2007; Lewis, 2011; Rawlins & Johnson, 2007). For example, Lewis (2011) found that sellers on eBay who auction second-hard cars receive higher prices when they include in the product description that the car does not have scratches, dents or rust. Sellers who mention any of these deficits receive lower prices than sellers who do not mention these deficits. In this chapter we formulate the theory that the presence of a reputation system may explain this trust of buyers in product descriptions.

Reputation systems as an incentive to speak the truth

Reputation systems collect, aggregate and distribute feedback about sellers to buyers (Resnick et al., 2000). The dominant theory of reputation systems specifies a mechanism for how they enable trust in markets with product homogeneity. According to this theory, reputation information helps buyers learn about the incentives of sellers (learning; Buskens, 2003; Cook, Hardin, & Levi, 2005; Weigelt & Camerer, 1988; Yamagishi & Yamagishi, 1994). Sellers who do not have an incentive to deliver a low-quality product are less likely to do so than sellers who have that incentive. Buyers are expected to reflect their experiences in the reviews they provide: More positive experiences are translated into more positive reviews. Seller's past trustworthiness is reflected in better ratings and reviews, which provide a signal to other buyers that this seller delivers high-quality products. If buyers are more likely to trust sellers with a better reputation, sellers have an incentive to maintain that good reputation. Numerous experiments indeed show that reputation systems positively affect trust in homogenous markets (Boero et al., 2009; Bolton et al., 2004a; Charness et al., 2011; Duffy et al., 2013; Fehrler & Przepiorka, 2013).

A few studies have hinted at a different mechanism through which reputation systems may increase trust, suggesting that reputation systems may help buyers determine how accurate claims about product quality made by sellers are (Ba & Pavlou, 2002; Houser & Wooders, 2006; Lewis, 2011; Neto et al., 2016; Snijders & Zijdeman, 2004). However, we are not aware of any paper that elaborates on the mechanisms underlying this effect. In this chapter we therefore develop a theory that explains how reputation systems may help overcome Akerlof's adverse selection problem when there is heterogeneity in product quality between and within sellers, and derive testable hypotheses to assess this theory empirically.

Contrary to what the dominant theory of reputation systems assumes, reputation systems contain information about both the absolute quality *and* the accuracy of product descriptions as illustrated by the examples in the introduction. Reputation systems thus do not only allow buyers to select sellers on the basis of the quality of the products they have delivered in the past. They also allow buyers to assess the extent to which sellers have been honest in the past. Sellers who have lied in the past about the quality of their products are

likely to have worse ratings and reviews than sellers who have not lied in the past. When assessing the accuracy of a product description, buyers can thus use the information in the reputation system. Sellers who have lied more often in the past are more likely to do so in the future. Honest sellers that promise to deliver a high-quality product are therefore expected to receive higher prices than dishonest sellers that make the same promise. We thus expect that buyers are more likely to believe information if that information is provided by more honest sellers.

> **H1:** The better the reputation of a seller for being honest is, the more likely buyers are to trust claims about quality from this seller.

If we assume that reputation systems also contain information about the accuracy of past product descriptions, reputation systems do not only provide an incentive to deliver high-quality products, but also an incentive to be honest in their descriptions of products when they are of low quality. Sellers may sometimes have high-quality products and sometimes only low-quality products. When they only have products of inferior quality, it may be tempting to lie about the quality of the products in order to attract higher prices. However, we argue here that in the presence of a reputation system the long-term benefits of communicating honestly may outweigh the short-term costs. In the long run, sellers may benefit from the reputational gains from honest communication in the past. If the expected future reputation gains are larger than the short-term costs, sellers are expected to provide a high quality good if they can, and to communicate honestly when they only have products of inferior quality. The reputation system thus provides an incentive for honest communication, which makes that communication more credible. While sellers who promise to deliver a high-quality product in the absence of a reputation should not be trusted, the reputation system creates the conditions for these messages to be credible. Combining communication with reputation may thus "make cheap talk costly."

According to this new theory, sellers who want to keep up a good reputation need to be honest about the quality of their products to buyers. We therefore hypothesize that in the presence of a reputation system, sellers are more likely to do as they promised.

> **H2**: The correlation between promises made by sellers and their actual behavior is stronger in the presence of a reputation system.

Because the information that is provided by sellers about the quality of their products is expected to be more truthful when there is a reputation system, we expect that buyers rely more on that information when there is a reputation system.

> **H3**: The effect of promises to honor trust on trust is stronger in the presence of a reputation system.

When sellers and buyers interact only once and when there is no reputation system, noncooperative behavior has no direct disadvantages for individual sellers, but it leads to negative outcomes for the entire market. We expect that in the presence of a reputation system, and with the possibility to communicate, sellers can credibly communicate the quality of their products, and buyers can adjust the prices they offer to the quality of the products. They are expected to pay higher prices to goods of better quality, and lower prices to products of inferior quality. Without the reputation system, communication between buyers and sellers is cheap talk, and buyers are not expected to pay higher prices when sellers promise to deliver better quality products.

When the heterogeneity of product quality within sellers is large, reputation alone may not be enough to solve the trust problem. When there is only a reputation system, but no communication between buyers and sellers, the reputation system only contains information about the absolute quality of the product, not about the accuracy of the product description. This information is found to increase trust in experiments in which there is no heterogeneity in product quality within sellers (Boero et al., 2009; Bolton et al., 2004; Charness et al., 2011; Duffy et al., 2013; Fehrler & Przepiorka, 2013; Frey & Van De Rijt, 2016). In those homogenous markets, all sellers could potentially deliver the same quality, and only when they choose to not do that the buyer will end up with an inferior product. However, when the quality of the products offered by the same seller varies greatly over time, the quality of products delivered in the past are no guarantee for the future. Consequently, in markets with heterogeneity in product quality within and between sellers neither communication nor a reputation system are expected to be a sufficient condition for better market outcomes. Strictly speaking, the theory predicts that

there is no difference in trust between the cases when only a reputation building or communication are possible or neither, and that only when both are present trust increases. We thus expect that both communication and a reputation system are necessary to increase trust.

> **H4**: Buyers place more trust in sellers when there is a reputation system and communication than when there is a) only communication, b) only reputation or c) neither of them.

Empirical approach: Two studies

The notion that communication between buyers and sellers affects buyers' decisions because there is a reputation system is implicit or explicit in many studies using empirical data from real auction platforms, and especially data from eBay (Anderson et al., 2007; Ba & Pavlou, 2002; Bolton et al., 2004a; Houser & Wooders, 2006; Lewis, 2011; Neto et al., 2016; Rawlins & Johnson, 2007; Snijders & Zijdeman, 2004). However, none of these papers tests this hypothesis. We are aware of one study on the interaction between seller reputation and claims made by sellers: studying of baseball cards on eBay, Jin and Kato (2006) do not find evidence that buyers rely more on claims made by more reputable sellers. However, their dataset contained both auctions of graded and ungraded cards. While the quality of ungraded cards is very difficult to detect online, this is not true for graded cards. For the latter category, claims and ratings of sellers have a different effect (Jin & Kato, 2006). Moreover, the rating of the seller and the claims about the quality of the cards made by the seller were found to be correlated with the price of the cards, as well as with whether the card was graded or not. Based on the pooled results for graded and ungraded cards we cannot conclude that the effect of claims is equally strong for reputable and for unreputable sellers of ungraded cards.

We test the hypothesis that buyers trust claims more when they are made by sellers with a better reputation (H1) by using data from eBay auctions of second-hand cars collected by Lewis (2019). When using data from real platforms one is bound to the institutional design of that platform. eBay allows its sellers to communicate with buyers, and to build a reputation. However, to test hypotheses 2 through 4, we will need to systematically vary the presence of

the reputation system and the opportunity for sellers to communicate with buyers. We are aware of a few studies in which the combined effect of reputation systems and communication systems is tested (Brosig-Koch & Heinrich, 2018; Cason & Mui, 2014; Denant-Boemont, Masclet, & Noussair, 2011; Duffy & Feltovich, 2006; Wilson & Sell, 1997). However, none of these studies implement the Akerlof problem of quality heterogeneity and are therefore of little use for understanding the problem we address in this chapter. In Study 2 we use an online experiment with heterogeneity within and between sellers in which we test if the presence of a reputation system changes the effect of communication on trust (H2, 3 and 4). We also provide a second test of H1.

Methods study 1: field study of an online market

To provide a test of H1 that buyers trust claims made by sellers with a better reputation for honesty more, we made use of data that are collected by Lewis (2011, 2019) that are publicly available[14]. The dataset contains information on 146,734 auctions of cars offered between March and October 2006 on eBay. The data were collected using a web scraper. Trust was operationalized as the number of bidders, the highest bid and the probability that the car was sold. We assumed that more buyers are willing to pay for a car if the seller was considered to be more trustworthy, and that they were willing to pay higher prices. Paying for a car involves a risk, as the buyer is never sure that the car is of good quality. We analyzed how the seller's reputation and claims about the quality of the car affected these dependent variables. Following Lewis (2011), we excluded auctions of new cars, cars that were built before 1950, incomplete or incorrect observations (e.g. with a negative mileage) and salvaged cars. We also excluded cars that could be bought immediately for a specified price and cars with a warranty, because the trust problem is less severe in those cases. One outlier with a bidding price of $860,100 was excluded.

Positive seller reputation was operationalized as the number of positive ratings the seller obtained, and negative reputation was defined as the number of negative ratings. Because these distributions were highly skewed, with many sellers having no or a few ratings and only very few sellers with many ratings,

[14] Data can be retrieved from: https://doi.org/10.1257/aer.101.4.1535 .

we calculated the natural logarithm of these numbers, after first adding 1 to avoid ln(0), before including them in the analyses. 129 auctions by 78 sellers that had mode than 5000 positive reviews are removed. 9 outliers with more than 100 negative ratings are removed. The remaining dataset contains 65,113 observations.

In our measurement of claims about product quality made by sellers we follow Lewis (2019) who identified words in product descriptions that were relevant for the value of the car (e.g. rust, dent and scratch) and were frequently used. He then identified for each car if the seller mentioned the word and in which context. Sellers either did not mention a problem, informed the buyer that the car had the problem ('car has rust'), or explicitly mentioned that the car did not have the problem ('car is rust free'). We used this information to construct two variables counting the number of problems acknowledged and denied by the seller.

To test H1, we regressed measures of buyer trust on measures of seller reputation. We used the reputation of the seller and claims made by the seller as independent variables. We used three different dependent variables. First, we took the natural logarithm of the price offered by the highest bidder. Listings that did not receive any bids are excluded from the analyses. We used linear regressions for this variable. For the second dependent variable (whether the car was sold or not) we used logistic regressions. The third dependent variable is the number of bidders, which is a count variable. We used negative binomial regression for this variable.

For every dependent variable, we first constructed a model with the independent variables and all control variables, as well as fixed effects for car model, year of production and week of the auction. Some cars were auctioned by the same seller, so we used multilevel models with random intercepts for sellers. To the second model we added the interaction between the reputation variables and the number of problems denied and acknowledged by the seller to test if bidders rely more on claims made by sellers with a better reputation (H1).

We controlled for several car characteristics: logged mileage, logged length of the product description (in bytes), whether the car was a featured listing, the number of options a car had (none, one or two, or more than two). We also

controlled for the duration of the auction (in days) and for several seller characteristics: whether the seller provided a phone number and address and whether the seller was a power seller. The histogram of mileage shows a relatively normal distribution, except for a high peak at the maximum of 500,000 miles. The results do not change when these cars are excluded from the analyses.

Table 4.1 contains the descriptive statistics for the variables used in the analyses. The average seller has many more positive reviews than negative reviews. This skewness of rating scores has been reported for all kinds of platforms (Chapter 2; Bridges & Vásquez, 2018; Teubner & Glaser, 2018). The descriptive statistics also show that on average, sellers acknowledge more problems than they deny.

Results of study 1

The results of the regressions for the main variables of interest are displayed in Table 4.2. The full table including the results for the control variables can be found in the Appendix, Table B4.1. In line with Lewis (2011) our regression results show that sellers with a larger number of positive reviews received slightly lower prices (model 1). A potential explanation for this may be that sellers who set lower (reserve) prices generally receive more positive reviews. The finding that sellers with more positive reviews were more likely to sell their car (model 3) while they received lower prices supports this idea. Alternatively, the quality of the products sold by sellers may decrease over time. Sellers with more positive reviews also attracted more bids than sellers with fewer positive reviews (model 5).

Having more negative ratings is negatively related to price (model 1), and positively related to the probability that the car was sold (model 3). Sellers with more negative reviews may realize that their reputation negatively affects their chances to sell cars at a high price and may therefore set lower reserve prices, which might explain the larger sales probability. Negative reviews did not affect the number of bidders (model 5).

Table 4.1: Descriptive statistics

	N	Mean	SD	Min	Max
Price	56,809	8,693	9,909	0.99	195,302
Logged price	56,809	8.55	1.11	-0.01	12.18
Sold	65,113	0.43	-	0	1
Number of bidders	65,113	5.97	4.52	0	22
Number of positive reviews	65,113	122	320	0	4,993
Logged positive reviews	65,113	3.11	2.00	0	8.52
Number of negative reviews	65,113	1.16	3.71	0	100
Logged negative reviews	65,113	0.39	0.68	0	4.62
# of problems denied	65,113	0.42	0.76	0	6
# of problems acknowledged	65,113	0.55	0.87	0	7
Miles	65,113	107,596	96,904	1	500,000
Text size (in bytes)	65,113	2,867	3,983	29	77,163
Featured listing	65,113	0.06	-	0	1
Phone number provided	65,113	0.36	-	0	1
Address provided	65,113	0.16	-	0	1
Powerseller	65,113	0.08	-	0	1
No options	65,113	0.28	-	0	1
One or two options	65,113	0.21	-	0	1
More than two options	65,113	0.51	-	0	1
Age car (years)	65,113	19.60	13.36	1	57
Duration (days)	65,113	6.85	1.58	2	12
N = 65,113					

On average, denying problems did not affect prices (model 1), but had a positive effect on the probability that the car was sold (model 3) and on the number of bidders (model 5). Lewis' (2011) results on the effect of denying problems were mixed: he found that denying some problems led to higher prices, while denying other problems did not affect prices. The significant interactions between the number of problems denied and the number of positive reviews on the probability that the car was sold (model 4) and on the number of bidders (model 6) suggest that sellers with more positive reviews benefit less from denying problems than sellers with fewer positive reviews.

The effect of positive quality claims on prices did not differ across sellers with more or fewer positive reviews. These results suggest that quality claims are not more likely to be believed when sellers have a better reputation.

In line with Lewis' (2011) findings, we find that acknowledging problems had a negative effect on prices (model 1). Acknowledging problems had a positive effect on the number of bidders (model 5) and the probability that the car was sold (model 3). Sellers who acknowledge problems may have set lower secret reserve prices, which may explain why they were more likely to sell their car despite lower prices. If lower prices are bid, the auction may attract more bidders than auctions with higher prices. The positive effect of acknowledging problems on sales probability and number of bidders was smaller for sellers with more positieve reviews. This finding may be explained by ceiling effects: sellers with positive reviews attract more bidders and are more likely to sell their cars. The additional effect of acknowledging problems is therefore smaller. This may also explain why sellers with more negative reviews suffer less from acknowleding problems.

According to our theory, both having a good reputation and promising to deliver a high-quality product are necessary for trust. When one of the two is lacking, no trust should be placed. We hypothesized that buyers were more likely to believe sellers who claimed that their car had no problems if that seller had a better reputation (H1). The results do not provide support for this hypothesis: sellers who have more positieve ratings benefit less from denying problems. Making positive quality claims generally positively affects the outcomes of the sellers. This suggests that promising to deliver a high-quality product is enough to convince buyers to place trust, while reputation alone does not suffice: acknowledging problems has a more negative effect for sellers with a better reputation. This may imply that negative quality claims undermine the positive effect of reputation.

Table 4.2: Multilevel regression results (control variables omitted)

	Logged price		Sold		Number of bidders	
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Main effects						
Log positive ratings	-0.01***	-0.01***	0.33***	0.36***	0.03***	0.04***
	(-0.02, -0.01)	(-0.02, -0.01)	(0.31, 0.35)	(0.33, 0.38)	(0.02, 0.03)	(0.03, 0.04)
Log negative ratings	-0.02**	-0.04***	0.06*	0.05	-0.02	-0.02*
	(-0.04, -0.01)	(-0.06, -0.03)	(0.00, 0.12)	(-0.02, 0.13)	(-0.03, 0.00)	(-0.04, -0.00)
Denied	-0.00	-0.02**	0.19***	0.27***	0.02***	0.07***
	(-0.01, 0.00)	(-0.04, -0.01)	(0.15, 0.23)	(0.19, 0.35)	(0.02, 0.03)	(0.05, 0.09)
Acknowledged	-0.10***	-0.11***	0.26***	0.40***	0.02***	0.07***
	(-0.10, -0.09)	(-0.13, -0.10)	(0.23, 0.29)	(0.33, 0.47)	(0.01, 0.02)	(0.05, 0.08)
Interactions						
Positive * denied	-	0.00	-	-0.03*	-	-0.01***
		(-0.00, 0.01)		(-0.05, -0.00)		(-0.02, -0.01)
Negative * denied	-	0.01	-	0.01	-	0.00
		(-0.01, 0.02)		(-0.05, 0.08)		(-0.01, 0.02)
Positive *	-	0.00	-	-0.04***	-	-0.02***
Acknowledged		(-0.00, 0.00)		(-0.06, -0.02)		(-0.02, -0.01)
Negative *	-	0.03***	-	0.01	-	0.01*
Acknowledged		(0.02, 0.03)		(-0.05, 0.06)		(0.00, 0.02)

(Continued on next page)

Table 4.2: Continued

	Logged price		Sold		Number of bidders	
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Constant	9.70***	9.70***	-1.24	-1.32	1.98***	1.95***
	(9.41, 9.98)	(9.41, 9.98)	(-2.69, 0.20)	(-2.77, 0.13)	(1.63, 2.33)	(1.60, 2.30)
Control variables	Yes	Yes	Yes	Yes	Yes	Yes
FE model/year/week	Yes	Yes	Yes	Yes	Yes	Yes
Seller intercept	Yes	Yes	Yes	Yes	Yes	Yes
Observations	56,809	56,809	65,113	65,113	65,113	65,113

Note: 95% Confidence Interval in parentheses. *** p<0.001, ** p<0.01, * p< 0.05. Random intercepts for sellers. Models 1 and 2: Multilevel linear regression. Models 3 and 4: Multilevel logistic regression. Model 5 and 6: Multilevel negative binomial regression.

Discussion of study 1

The results from study 1 suggest that bidders do not attach more value to positive quality claims when they are made by sellers with a better reputation. Having more positive reviews does not increase the extent to which denying problems increases prices or the probability that the car is sold. Having more negative reviews does not affect the effect of denying problems on prices, sales probability and the number of buyers. Based on these results we reject Hypothesis 1.

Our results provide some evidence that sellers with a better reputation suffer more from making negative quality claims. This suggests that a seller's reputation alone is not enough to convince buyers to trust them. Based on these results we conclude that quality claims are necessary to convince buyers to place trust, while having a good reputation is not enough.

One potential explanation for why we did not find the expected effect is that we could not distinguish between having a reputation for honesty and a reputation for providing high-quality products. Reputation was operationalized by the numeric ratings provided by other users on the platform. These ratings were the results of a (mental) aggregation of the evaluation of different aspects of the interactions. The data do not provide any details on what these aspects were and their weighing, so we do not know if these ratings were the result of providing a high-quality product or of honest communication or both. Whereas rating scores tend to be overly positive, written reviews are found to be more detailed and nuanced (Bridges & Vásquez, 2018). The written reviews could thus provide more insight here, but these were not available.

Another potential explanation for why we did not find the expected effect at the individual level is that the majority (87%) of sellers has at least one positive review and that sellers on average only have a few negative reviews. This skewness in reputation scores is predicted by our theory: The reputation system provides an incentive to speak the truth. Hypothetically, this effect may be so strong that differences between individuals with different intrinsic motivations disappear. Since the reputation system remained unchanged during the data collection period, we cannot rule out this explanation with the current dataset.

A related limitation of using observational data is that we cannot rule out confounding effects of omitted variables. Both reputation scores and trust in promises made by sellers may be caused by other, unobserved variables. For example, earlier studies on discrimination in the platform economy show that demographic characteristics, such as ethnicity, affect an individual's chances of being trusted (Edelman & Luca, 2014; Tjaden et al., 2018). Given that on most platforms reviews can only be written after a completed transaction, these initial differences in the chances of being trusted may be reflected in the number of reviews an individual collected (see chapters 2 and 3). This would imply that the relationship between a seller's reputation and the extent to which buyers rely on claims made by these sellers is at least partly spurious.

To address these issues we conducted an online experiment in which we systematically varied the presence of a reputation system and the possibility to communicate between buyers and sellers. This allowed us to test H1 in a controlled environment, eliminating the influence of ommitted variables through randomization, and to test H2, H3 and H4 regarding the presence of a reputation system. It also allowed us to look in more detail into the effect of the different types of reputation information.

Methods study 2: an online experiment

Design

The experiment consisted of 36 rounds of an adaptation of the widely used Trust Game (Dasgupta, 1988; Kreps, 1996). Subjects in a session were divided into groups of six or eight participants. At the beginning of each round, half of the participants in a group were assigned to role A (buyer), the other half was assigned to role B (seller). The game was divided into three blocks of twelve rounds. In each of these blocks, every subject played six times in role A and six times in role B, but the order in which they played in the different roles was determined randomly (cf. Charness et al., 2011)[15]. In every round, each player A was randomly and anonymously matched with a player B.

[15] This procedure ensured that the number of times all subjects acted in the roles of trustor and trustee did not vary too much at any point in time, but at the same time that subjects

The pairs played one of two versions of a Trust Game. The left part of Figure 4.1 shows the first version of the game, which was a regular Trust Game. In that game, player A could choose to place trust ('pay a high price') or to not place trust ('pay a low price') in player B. If player A did not place trust, the game ended and both players received a payoff of 2 points. If player A chose to place trust, player B could choose between honoring ('deliver a high-quality product') or abusing trust ('deliver a low-quality product'). If player B honored trust, both players received a payoff of 3 points. If player B abused trust, player A received 0 points and player B received 4 points. The second version of the game was similar to the first version. The only difference was that player B did not have the option to honor trust if player A placed trust. The right panel of Figure 4.1 shows the structure of one round of this game without a reputation system. In this figure P, R, S and T represent the material payoffs of the actors. Figure 4.2 contains the numbers used in the current experiment.

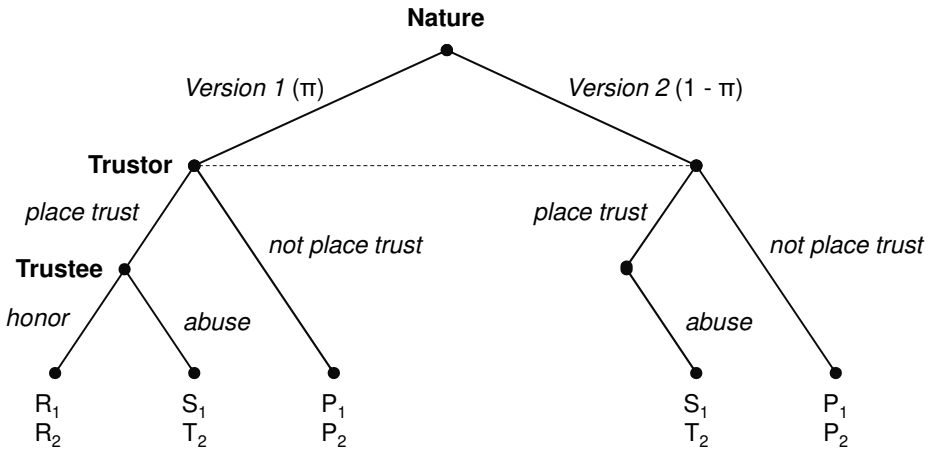

Figure 4.1: One round of the game used in the experiment, where $T_2 > R_i > P_i > S_1$.

did not know for sure what role they would play in the next round. Only when a subject quit the experiment in the middle of a session it was sometimes necessary for the other subjects to play more than six times in each role in a block of twelve rounds.

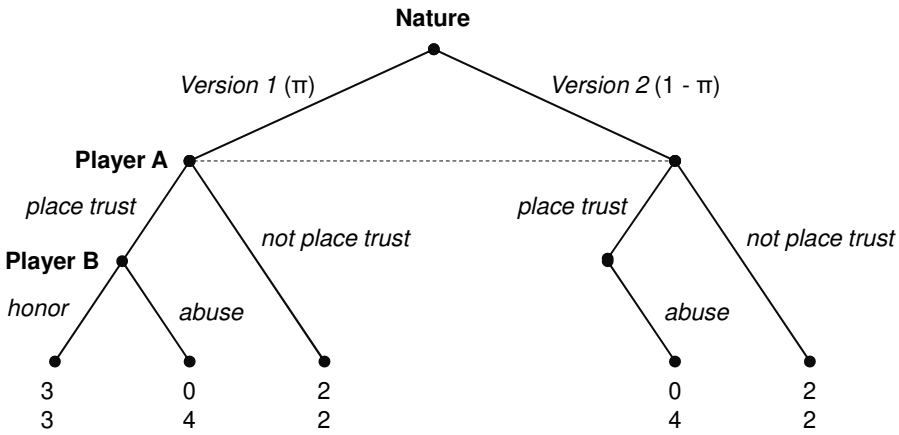

Figure 4.2: Numerical example of the game used in the experiment.

In each round, the pairs played version 1 of the game with probability 0.5 and version 2 with the remaining probability. Players B were informed about which version of the game they were assigned to at the beginning of the round. Players A never knew which version they were playing. This heterogeneity introduced uncertainty for players A, since they did not know when a player B had the opportunity to honor trust and when not. Even when a player B was very trustworthy and always chose to honor trust when they had the opportunity to do so, they may frequently have no choice but to abuse trust.

The experiment had a two (No communication versus Communication) by two (No reputation versus Reputation) between-subjects design. Each session was assigned to one of four conditions. To avoid confusion, we refer to the treatment with both communication and reputation as the 'Combined' condition. In the Communication treatments, sellers could send a binary, non-binding message at no cost to the buyers in which they revealed what their choice would be if player A chose to place trust. In the Reputation conditions, reputation was operationalized as information about previous decisions made by players B. Before choosing to place trust or not, players A were informed about the actual decisions made by player B in the rounds in which player B acted in role B. In the Combined condition players A were also informed about promises made by players B in the past. Note that B's did not receive any information about A's previous behavior in any of the conditions. The complete instructions for the participants can be found in the Appendix.

Procedures

260 subjects participated the experiment in May 2020[16]. Nine subjects quit the experiment before the end and were excluded from the analyses[17]. Ten sessions that each lasted for about 75 minutes were conducted. Subjects were recruited via Prolific (2019)[18]. Only subjects with the British nationality who were 18 year and older could participate. Subjects could participate only once and were paid according to their decisions and decisions made by others in the experiment. For every 15 points earned in the game, the subjects received £1.00. 73% of the subjects was female, 20% male, one subject indicated 'other gender' and the remaining subjects did not indicate their gender. The age of the subjects ranged from 18 to 79 (mean = 39). 12% of them indicated to be a student. 8% of the subjects had no degree, or only completed elementary school. About one third completed high school, another third obtained an undergraduate degree and 15% held a graduate degree.

The subjects received on-screen instructions. After reading the instructions they completed some comprehension questions. They could only continue with

[16] Each session was scheduled at a fixed date and time. Subjects could sign up for a specific session through a link that was posted on Prolific. Subjects who finished the instructions, but who could not be assigned to a group received a show-up fee of £2.50.

[17] A total of 260 subjects started the experiment. Eight of them chose to quit the experiment before the end and nine lost the connection with the server. If a participant quit the experiment, the participant they were matched with also had to stop, because the experiment is conducted in groups with an even number of subjects. All data from subjects who decided to quit are excluded from the experiment. For the other subjects the data from the rounds that they completed were included. If the remainder of the group would consist of only two members, the data of the remainder of the rounds are excluded from the analyses.

[18] The experiment was initially designed to be conducted in a laboratory, but due to the Covid19-crisis we diverted to an online solution. While laboratory experiments allow for more control over the behavior of the subjects during the experiment, the online setting provided a higher level of anonymity, since subjects did not see each other before and after the experiment. Since the theory assumes that buyers and sellers are strangers, the online setting provides a cleaner test of the theory. Moreover, the subject pool of Prolific is more varied in terms of a number of demographic characteristics compared to the average laboratory student pool.

the experiment after they answered all questions correctly. The experiment was computerized using oTree (Chen et al., 2016). Figure A4.2 in the Appendix contains an example of the screen the participants saw during the experiment. Participation was anonymous and no deception was used[19]. After playing the game, the subjects completed a short survey asking about some demographic characteristics as well as about their experience with game theory and their idea about the goal of the study.

Variables

Two treatment dummies were constructed that indicated whether a subject participated in a game with or without a reputation system, and whether player B could communicate with player A. The main dependent variable in most of the analyses was the decision of player A to place trust in player B.

The message sent by player B was operationalized by two dummy variables. These dummies indicated whether player B promised to honor trust or informed player A that they would abuse trust. The reference category was not sending a message. Honest communication was defined as promising to abuse trust, or as promising to honor trust and actually honoring trust. Dishonest communication entailed promising to honor trust, but then abusing trust. When player B promised to honor trust, but did not get the opportunity to prove that they were honest because player A did not place trust, players A could not judge whether that promise was honest or dishonest. Rounds in which the promise could not be classified as honest or dishonest have missing values for this variable.

Player B's reputation was measured by four variables that were constructed in two different ways. The first method included all previous decisions made by a player B in role B in the reputation variables. The first pair of variables counted the number of times player B honored and abused trust in the past. The second pair of reputation variables was only constructed for subjects in both conditions with communication and counted the number of times player B communicated honestly or dishonestly in the past. The second method

[19] In order to pay the subjects according to their earnings in the game, we had to collect their Prolific participant IDs. After completing the cleaning of the data, these IDs are stored separately from the experimental data.

accounts for the fact that although players A had access to all previous decisions of player B, they had to scroll down in the history window when they wanted to see more than the last six rounds (see Figure A4.1 in the Appendix). According to this second method, the four reputation variables were constructed in the way we just described, except that only the last six rounds in which player B acted in role B were included. In every analysis, only one pair of reputation variables is included, to avoid multicollinearity problems.

Results of study 2

Descriptive statistics

Table 4.3 contains the descriptive statistics.

Table 4.3: Descriptive statistics

	N	Mean	SD	Min	Max
A placed trust	3849	0.28	-	0	1
B honored	1088	0.40	-	0	1
Both communication conditions					
B promised to honor trust	1849	0.60	-	0	1
B promised to abuse trust	1849	0.18	-	0	1
Honest message (if defined)	863	0.73	-	0	1
Both reputation conditions					
Reputation B: # honored	1953	1.06	1.45	0	11
Reputation B: # abused	1953	1.70	1.69	0	9
Combined condition					
Reputation B: # honest	1000	2.66	2.85	0	16
Reputation B: # dishonest	1000	1.34	1.76	0	8
Number of decisions	3849				
Number of subjects	242				
Number of groups	36				

Trust was placed in 28% of the rounds. 58% of players B in the communication conditions promised to honor trust, while 40% actually honored trust. Most players B only accumulated a few reputation scores in a typical game.

Figure 4.3 shows the distributions of the outcomes per condition. The y-axis shows the percentage of rounds in which an outcome occurred. The x-axis shows whether player A placed trust or not, and if so, whether player B honored or abused trust. The different colors indicate the message sent by player B.

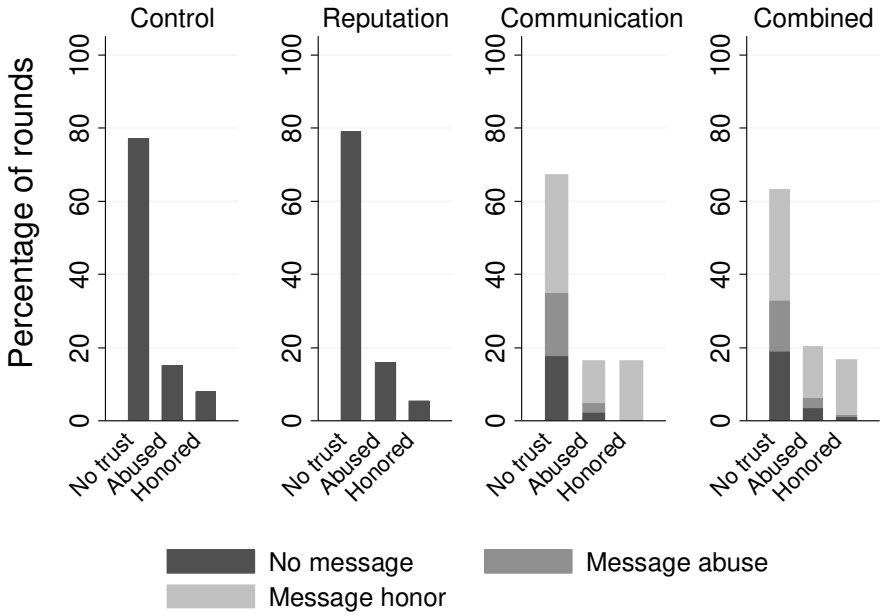

Figure 4.3: Distribution of outcomes across the conditions.

Group-level tests of Hypotheses 2 and 4

Hypotheses 2 and 4 are about differences between conditions, while Hypotheses 1 and 3 are about individual differences. First, we test hypotheses 2 and 4 using group level tests.

Hypothesis 4 predicts that trust will be placed more often if both a reputation system and the opportunity to communicate are present. The size of the leftmost bar for every condition in Figure 4.3 shows that more trust was placed in both communication conditions than in the two no-communication conditions. To test Hypothesis 4, we compared the trust level of the Combined

condition with each of the other conditions individually, using three Wilcoxon rank-sum tests at the group level in which we compare the average level of trust across the conditions. These three tests show that players A placed more trust in the Combined condition than in the Control condition ($z = 3.136$; $p = 0.002$, $N = 19$) and reputation condition ($z = 2.887$, $p = 0.004$, $N = 18$), supporting Hypothesis 4b and 4c, but not than in the Communication condition ($z = 0.572$, $p = 0.567$, $N = 19$), contrary to Hypothesis 4a. The evidence thus supports the idea that communication enables trust, but not that this effect is conditional on there being a reputation system. We reject Hypothesis 4a, and accept Hypotheses 4b and 4c.

To test if the limited effect of the presence of the reputation system is due to the small average number of reviews per player B, we repeated the Wilcoxon rank-sum tests of differences between the level of trust between the conditions, but only included the second half of each game. In takes time for trustees to build a reputation, so the hypothesized effects are expected to be stronger in the second half of the game. The results of these tests are not different from the main tests, suggesting that the lack of an effect of the presence of a reputation system cannot be explained by the limited number of reviews in the first half of the game.

According to Hypothesis 2, sellers are more likely to be honest in the presence of a reputation system. Figure 4.3 does not show any clear differences in the accuracy of promises between the Communication condition and the Combined condition. A Wilcoxon rank-sum test confirms that there was no significant difference in the fraction of honest messages between the two conditions ($M_{\text{Communication}} = 0.73$, $M_{\text{Combined}} = 0.71$, $z = 0.089$, $p = 0.929$, $N = 18$). We reject Hypothesis 2.

In all conditions, trust decreased over the rounds, from about 66% in the communication conditions and 35% in the other conditions in the first rounds, to about 20% in the last round. Figure A4.3 in the Appendix shows the change in trust over the rounds in each condition.

Individual-level tests

Hypotheses 1 and 3 are tested at the level of individual decisions made in a single round. As a robustness check we also evaluated hypotheses 2 and 4 at the individual level. To account for interdependencies between observations, we included random intercepts for individual subjects and for the group they played in. 18% of the variance was explained at the level of the individual and 2% of the variance was explained at the group level. In all analyses we controlled for round number and group size.

Table B4.2 in the Appendix contains the results of a logistic regression per condition of the probability that trust was placed on the reputation and communication variables. In all conditions, players A were more likely to place trust in players B who promised to honor trust. As expected, player B's reputation did not affect player A's decisions when this information was not visible. In the Reputation condition, players A were informed about the total number of times a player B honored and abused trust, but this information did not affect trust either. This was also predicted by the theory: honoring trust in the past is no guarantee for future behavior, because in half of the rounds player B did not have the option to honor trust. We expected that in the Combined condition more trust would be placed if player B had a better reputation. We indeed find that trustors were more likely to place trust in trustees who had honored more trust in the past and who had been less dishonest in the past. We did not find that the number of times a player B abused trust and had been dishonest in the past affects player A's decision.

As a robustness check, we repeated these analyses with a slightly different calculation of the reputation scores. Each player A viewed player B's reputation in a history box on their screen. The last six rounds in which player B acted in role B were always visible, but to see earlier rounds, player A had to scroll down. When only including the last six rounds in the role B in the reputation variables, the number of times a player B honored trust positively affected A's decision to place trust (odds ratio = 1.33, $z = 2.18$, $p = 0.029$). The results for the Combined condition did not change when reputation was operationalized as the last six decisions of player B.

To test Hypothesis 1, we ran two logistic regressions with the decision to place trust in the Combined treatment as the dependent variable. The two honesty-

reputation variables that are based on all previous decisions of player B in role B were centered and interacted with two dummy variable that indicated if player B promised to honor trust. In each model, we included one of the interactions between the B's reputation for honesty and the message sent. The results do not change when including all interactions at the same time, so model 1 in Table 4.4 contains the results of a logistic regression model in which both interactions are included. The large and significant coefficient of the variable 'Promised to place trust' shows that players A were more likely to place trust in players B who promised to honor trust. The interactions between the promises made by player B and B's reputation for honesty are not significant. For models that include interaction effects, differences between logistic regression models and linear probability models may arise (Ganzach, Saporta, & Weber, 2000). In our case the results do not change when running a linear probability model. The results also remain the same when transforming the reputation scores of the trustees to dummy variables that indicate whether the trustee has at least once been (dis)honest in the past. These results do not provide support for Hypothesis 1.

When using the reputation variables that are only based on the last six rounds in which the player B acted as player B the results slightly change. The interaction between the number of times player B had been honest and the promise to honor trust turns positive and significant (odds ratio = 1.254, z = 2.04, p = 0.041). This means that players A were more likely to place trust in a player B who promised to honor trust when that player B had been more honest in the last six rounds. In sum, we find only weak support for Hypothesis 1.

The results at the individual level corroborate the findings from the group-level rank-sum test of Hypothesis 2. Model 2 in Table 4.4 shows the results of a logistic regression of both communication conditions with the honesty of player B as the dependent variable. The dummy variable that indicated if player B promised to honor trust was interacted with the dummy variable for the reputation treatment. We find that players B were not more likely to make an honest promise in the presence of a reputation system than when there was no reputation system. The results did not change when fitting a linear probability model.

Model 3 in Table 4.4 contains the results of a logistic regression of both communication conditions with the decision of player A as the dependent variable. We interacted the variable that indicated if player B promised to honor trust with the reputation treatment variable to test Hypothesis 3. Contrary to our expectations, players A attached less value to B's promise to honor trust when there was a reputation system than when there was no reputation system. This suggests that the reputation information partly substituted communication between As and Bs. When running a linear probability model, we find that the interaction turned insignificant. Neither model thus supports the hypothesis that the effect of promises to honor trust is stronger in the presence of a reputation system. We therefore reject Hypothesis 3.

To evaluate Hypotheses 4a, b and c at the individual level, we ran a logistic regression with the variables that indicated to which treatment the subjects were assigned as the independent variables and with the decision of player A as the dependent variable. Model 4 contains the results of this regression. Corroborating the finding from the Wilcoxon rank-sum test at the group level, we find that more trust was placed when players B could make promises to players A than when they could not. The presence of the reputation system did not affect the probability with which player A placed trust. Running a linear probability model does not affect the results.

To summarize, the results from the experiment do not provide support for the theory. We found that players A relied more on promises made by players B when that player B had been more honest in the past. However, this effect was only significant when including the last six rounds in the reputation scores, not when looking at all rounds. We neither found that players B were more honest, nor that the effect of these promises on player A's decision was stronger in the presence of a reputation system. Finally, the presence of a reputation system did not increase overall trust.

Table 4.4: Multilevel logistic regression results (hypothesis tests)

	Model 1 DV = Trust	**Model 2** DV = Honest	**Model 3** DV = Trust	**Model 4** DV = Trust
Treatment variables				
Reputation condition	-	0.73	2.19*	0.89
(ref. cat. = Control condition)		(0.40, 1.37)	(1.02, 4.74)	(0.53, 1.49)
Communication condition	-	-	-	2.10**
				(1.25, 3.50)
Combined condition	-	-	-	1.41
				(0.69, 2.89)
Promises made by player B				
Promised to honor trust	7.01***	-	14.85***	-
	(4.45 - 11.06)		(8.63, 25.54)	
Promised to abuse trust	0.91	-	1.01	-
	(0.46 - 1.81)		(0.63, 1.60)	
Reputation player B				
# Honest	1.09	-	-	-
	(0.87 - 1.35)			
# Dishonest	0.78	-	-	-
	(0.49 - 1.23)			

(Continued on next page)

Table 4.4: continued

	Model 1 DV = Trust	**Model 2** DV = Honest	**Model 3** DV = Trust	**Model 4** DV = Trust
Interaction variables				
Reputation honest	0.98	-	-	-
* promise to honor	(0.79 - 1.21)			
Reputation dishonest	1.14	-	-	-
* promise to honor	(0.72 - 1.82)			
Reputation honest	0.95	-	-	-
* promise to abuse	(0.74 - 1.21)			
Reputation dishonest	1.18	-	-	-
* promise to abuse	(0.67 - 2.09)			
Reputation treatment	-	-	0.48*	-
* promise to honor			(0.26, 0.86)	
Control variables				
Round number	0.95***	1.01	0.95***	0.96***
	(0.92 - 0.97)	(1.00 - 1.03)	(0.94 - 0.96)	(0.95 - 0.97)
Group size	0.92	1.11	1.01	1.04
	(0.76 - 1.11)	(0.93 - 1.32)	(0.87 - 1.18)	(0.94 - 1.16)
Constant	0.69	1.34	0.15**	0.29**
	(0.15 - 3.07)	(0.34, 5.26)	(0.04, 0.56)	(0.12, 0.73)
Number of decisions	1,000	863	1,849	3,849
Number of subjects	61	117	119	238
Number of groups	10	18	19	36

Note: Odds ratios. 95% Confidence Interval in parentheses. *** $p<0.001$, ** $p<0.01$, * $p< 0.05$

Additional analyses

To better understand the findings, we explored which decision would yield players A the best outcome, conditional on the condition they were in, any promises made by player B and the reputation of B. In the Control condition, placing trust on average yielded player A a payoff of 1.03 points per round, while not placing trust would result in a payoff of 2 points. In this condition, the expected payoff of not placing trust was thus higher than the expected payoff of placing trust. A linear regression with the payoff of player A as the dependent variable and the decision of the player A to place trust or not and while controlling for the round number and group size confirms that the expected payoff of not placing trust was higher than the expected payoff of placing trust (t = -19.31, p <0.001).

In the Communication condition, three regression analyses show that the average payoff of placing trust was lower than the payoff of not placing trust: the average payoff of player A of placing trust was 0.29 points if player B did not send a message (t = -23.28, p < 0.001), 0 points if player B promised to abuse trust and 1.75 points if player B promised to honor trust (t = -2.78, p = 0.006).

The left panel of Figure 4.4 shows the average payoff for players A who placed trust, conditional on the reputation of player B in the Reputation condition. The left panel of the figure shows that in the Reputation condition the payoff of placing trust was always lower than the expected payoff of not placing trust (2 points). A regression analysis shows that the expected payoff decreased if a player B had abused more trust (b = -0.110, t = -3.11, p = 0.002) and that the expected payoff of placing trust increased with the number of times player B had honored trust in this condition (b = 0.135, t = 2.77, p = 0.006). However, the expected payoff of placing trust only exceeded the payoff of not placing trust when player B had honored trust more than nine times in the past, while B's honored trust a maximum of five times in the current data.

The right panel of Figure 4.4 shows the average payoff of placing trust in the Combined condition, conditional on player B's reputation and only if the player B promised to honor trust. Two regression analyses confirm that placing trust in players B who had honored more trust in the past (b = 0.225, t = 4.64, p <0.001) and who had been more honest (b = 0.112, t = 3.93. p < 0.001) yielded

a higher payoff. The opposite was true for players B who had abused more trust in the past (b = -0.245, p < 0.001). The expected payoff of placing trust exceeded the payoff of not placing trust only if player B had honored trust at least twice in the past, or if player B had been honest at least three times in the past, which was within the range of the data. Together these results suggest that player B's behavior followed the patterns we predicted: only in the Combined condition B's reputation and promise could help improve player A's decision to place trust. These results also support the claim that reputation is type-separating: players B with a better reputation were more likely to be trustworthy and honest.

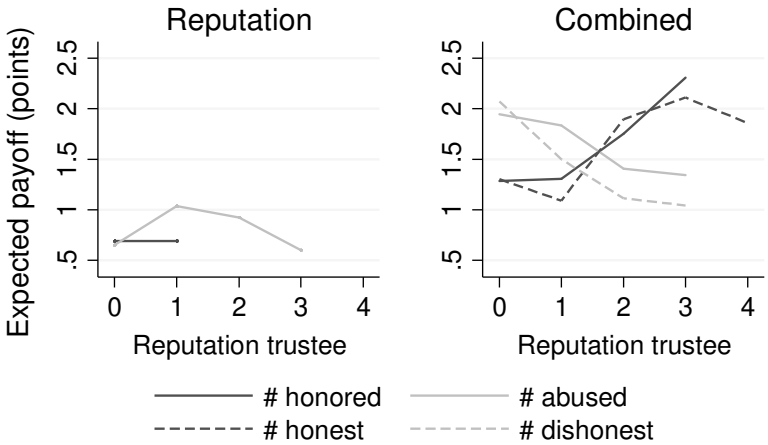

Figure 4.4: Average payoff of player A per condition conditional on the player B's reputation. Left panel: Reputation condition. Right panel: Combined condition. In the right panel only players B who promised to honor trust are included. Points representing fewer than 20 observations are excluded.

Discussion of study 2

We hypothesized that in heterogeneous markets, the possibility for sellers (players B in the experiment) to send messages to buyers (players A in the experiment) would only have an effect on trust in the presence of a reputation system. The results from the experiment show that buyers always place more trust when sellers communicate that their product is of good quality, regardless

of the presence of a reputation system and regardless of the reputation of the seller. Buyer behavior thus does not follow the patterns we predicted, namely that pre-play communication should only affect trust in the presence of a reputation system.

Interestingly, while buyer behavior did not follow our main predictions, sellers, and consequently the optimal strategies for buyers, did to some extent. The optimal strategy of buyers in the Combined condition depended on the reputation information available and on the promise made by the seller. Buyers who wished to maximize their payoff should not solely rely on the promises made by sellers, or on the reputation of the sellers, but on both. While we do not find that sellers are more honest in the presence of a reputation system, the combination of reputation information and communication could improve buyer's decision making. Reputation information about the extent to which sellers had been honest in the past could help player A assess the probability that a seller is lying in the current round. Given that buyers did not combine the historic information with current promises, it is not surprising that sellers did not respond by communicating more honestly.

The question remains why buyers did not follow the payoff-maximizing strategy. We assumed that buyers and sellers were rational and selfish. Buyers may have strongly relied on claims made by sellers, because they believed that most sellers would speak the truth. Alternatively, they may have found it difficult to decide what strategy would yield them the best outcome in the abstract decision situation. The steady decrease in trust over the rounds may indicate that this is indeed the case. Buyers may have learned through experience that placing trust does not increase their payoff in most cases.

General discussion

Messages sent by sellers are often found to affect prices that buyers are willing to pay in markets with heterogeneous products where prices do not signal quality (Anderson et al., 2007; Lewis, 2011; Rawlins & Johnson, 2007). However, when there is no way to find out if sellers speak the truth or not, this information can be considered cheap talk and should not affect buyer's decisions under the assumptions of rationality and selfishness. In this chapter

we have theoretically elaborated an often-suggested mechanism through which these messages may affect buyer decisions, even when all buyers and sellers are rational and selfish. We argued that the presence of a reputation system provides an incentive for sellers to speak the truth and a tool for buyers to find out who speaks the truth and who does not.

We first used data from eBay auctions (Lewis 2011, 2019) to test if buyers trust quality claims more if the seller has a better reputation. The results do not support this hypothesis. Using a second, experimental study we provide a second test of the same hypothesis, and we test a number of hypotheses about the influence of the mere presence of a reputation system on the decisions of buyers and sellers. The results from the second study corroborate the results from the first study. Although the combination of promises made and historical reputation information could help buyers improve their decision, they did not use this information to maximize their payoff.

Although the hypotheses are not supported, the results support our idea that reputation information is not enough to convince buyers to pay a high price in markets with heterogeneous product quality. In these markets, current product quality may differ from quality delivered in the past, so information about the past is not enough to make an accurate estimation of current quality. Supporting this claim, we found that sellers on eBay with a better reputation suffered more from acknowledging problems. Moreover, the experimental results show that information about the number of times a seller provided a good and a bad product in the past did not affect buyers' behavior when there was no communication possible.

In both studies buyers relied on the messages sent by sellers, even when the seller did not have a good reputation. This means that the puzzle why buyers trust promises made by sellers remains. One approach towards resolving this puzzle is to assume forms of irrational or boundely rational behavior. One possibility is that buyers and sellers do not realize what the true meaning of these promises is (Jin & Kato, 2006). This idea is supported by the experimental finding that the extent to which a seller had been honest in the past predicted the probability that they would be honest in the future. This finding implies that reputation information about past honesty is type-separating: only sellers who are intrinsically honest will communicate honestly.

The finding that buyers did not combine reputation information and promises made by sellers when making decisions might indicate limits to strategic rationality: in order to adapt their decision to the anticipated behavior of sellers, buyers need to reason a step ahead (Buskens et al., 2010). Similarly, sellers may not have realized that lying and speaking the truth in the present would affect their reputation. Alternatively, when buyers do not condition their behavior on the seller's reputation for honesty and their promise, there is no reason for sellers to communicate more honestly, which may explain why we did not find a difference in the communication of the sellers between the conditions. More qualitative research into people's understanding of the strategic setting may provide new insights here. Another idea would be to design a study that provides buyers with information that makes it easier for them to determine what their payoff-maximizing strategy would be, to approximate a situation with rational buyers.

Our theory relied on assumptions of rationality and selfishness. An alternative explanation for the effectiveness of communication is that participants in one-shot interactions often behave as if they expect repeated interactions (Hoffman, McCabe, & Smith, 1996; Kiyonari, Tanida, & Yamagishi, 2000). According to such arguments, communication may improve trust and cooperation by facilitating coordination, decreasing social distance, raising solidarity, and providing the cues of familiarity that are normally associated with durable relationships (Balliet, 2010; Kerr, Garst, Lewandowski, & Harris, 1997; Schniter, Sheremeta, & Sznycer, 2013). To explain why buyers rely on claims made by sellers, explanations that rely on different psychological mechanisms than rationality and selfishness appear to be necessary.

CHAPTER 5

Reputations in mixed-role markets: A theory and an experimental test

This chapter has been co-authored by Rense Corten and Arnout van de Rijt.

A slightly different version of this chapter is published as: Kas, J., Corten, R., & van de Rijt, A. (2020). Reputations in mixed-role markets: A theory and an experimental test. Social Science Research, 85, 102366.

Abstract

The traditional understanding of reputation systems is that they secure trust between strangers by publicly calling out cheaters. In modern, online markets, it is increasingly common for providers of a good to also act as consumers, and vice versa. We argue that in such mixed-role markets, reputation systems serve a second important function: They allow providers who lend out their possessions (such as their house, car or tools) to earn reputational credits that can be spent on future borrowing, especially when lending lacks monetary compensation. In an experiment that introduces a new game, "the Lending Game", we show that, consistent with our argument, information on past lending leads subjects to lend to those who have themselves lent before, increasing overall lending. However, when lending is financially compensated, this mechanism of reciprocal lending ceases to operate.

Introduction

The Internet opens up the possibility for interaction and economic exchange between strangers all around the world. Online exchange is characterized by anonymity, and strangers often interact without the prospect of future interactions (Kuwabara, 2015; Parigi et al., 2017). Individuals may benefit from exchange mediated by online environments, but with limited information and limited control over the decisions made by others, there is always the threat that one's trust is abused. A seller may not deliver after payment, a buyer may not pay after delivery, and a borrower may not return a lent good in time or intact.

Reputation systems are widely used to overcome trust problems in exchange between strangers (Cook, 2005). These systems inform individuals about past behavior of potential exchange partners. The positive effect on trust of the classic, single-role reputation system in which one actor only acts in one role (e.g. either as provider or as user) is well established in the existing literature (Boero et al., 2009; Bolton et al., 2004a; Charness et al., 2011; Duffy et al., 2013; Fehrler & Przepiorka, 2013). However, individuals increasingly act in multiple roles in modern-day exchange, acting as supplier of a good or service in one transaction while being on the receiving end in another (Cova et al., 2011; Ritzer & Jurgenson, 2010). For example, consumer-to-consumer marketplace eBay allows individuals to both buy and sell goods, peer-to-peer hospitality services Couchsurfing and Airbnb allow home owners to open up their houses to strangers, while staying at others' homes when traveling themselves (Lauterbach et al., 2009), and crowdfunding platform Kickstarter allows individuals to raise funds for their own projects and to fund projects of others.

In such cases where the lines between actors in different roles are blurred, a new type of reputation system emerges. This mixed-role reputation system allows individuals to not only examine another individual's past behavior in his or her current role, but also past behavior in other roles. For example, Couchsurfing hosts cannot only view ratings potential guests received from other hosts, but also reviews these potential guests received from other guests when they were host themselves.

Mainstream theory of reputation systems suggests this information is, at best, of secondary importance, because the best indication of whether someone can be trusted is that they have honored the trust they received in the past, not whether they themselves have trusted other people (e.g., Buskens & Raub, 2002). Yet anecdotal evidence suggests that this extra information is actually a very important criterion for platform users when making decisions. For example, on online forums Couchsurfing hosts said: *'One thing that may help is to host, if you're not already. That's the first thing I look for whenever I get a request. If they've hosted before and the review was positive, I almost always accept'* [them as a guest]*'* (Reddit, 2015) and *'One thing I have considered is only hosting people who are themselves active hosts'* (Quora, 2016). Research among Couchsurfing hosts suggests that hosts care about potential guests' past hosting behavior and that it is considered a general principle that one should alternate between surfing and hosting (Geiger, 2015; Geiger & Germelmann, 2015). On Couchsurfing, the number of times a member has surfed and hosted are highly correlated (Lauterbach et al., 2009), consistent with the notion that Couchsurfing users care about a balance between hosting and surfing experiences. The purpose of this chapter is to develop and experimentally test an explanation for the apparent importance of other-role reputations in modern mixed-role markets.

Not only free platforms such as Couchsurfing use this mixed-role reputation system. Car sharing platform Turo allows car owners to rent out their car to others and shows both reviews from travelers and reviews from car owners in the user profiles of car owners. The distinction between monetized exchange that includes financial remuneration (selling, renting etc.) and non-monetized, free exchange (e.g. lending) is relevant here, because the motivation of the users to participate in exchange is different. In the case of monetized exchange, individuals may decide to take the risk because they believe the monetary benefits outweigh the risk. In free exchange, there are no direct material benefits related to placing trust, so other motivations may play a larger role. Placement of trust, i.e., lending something valuable without financial compensation, is then easily viewed as an act of generosity or kindness. We argue that this difference between paid and free platforms not only leads to an overall difference in trust levels, but also to a difference in the effectiveness of implementing a mixed-role reputation system.

Thus, although mixed-role reputation systems are increasingly common, especially in online exchange, existing theories about single-role reputation systems do not provide an explanation for why and when this extra reputation information affects trust. In the current chapter we study markets in which users take multiple roles. We develop a theoretical framework explaining why information about the other role is important when making decisions and how that affects trust at the platform level. We argue that the effect of mixed-role reputation systems is stronger for platforms where goods or services are provided for free than on platforms where users have to pay. We perform a test of the theory in a controlled laboratory experiment in which we compare the effect of the single-role reputation system with the mixed-role reputation system in both paid and free exchange.

Theory

In the scenarios we study, individuals alternate between the role of provider and user.

Game structures of paid and free exchange systems

The classic understanding of how reputation systems help achieve trust in markets can be well explained using the analytical tools of game theory (Kreps, 1996). We use games as abstract representations of markets in which we quantify the incentives for the actors.

Paid exchange systems as a Trust Game

The situation in the direct (paid) exchange is typically represented as a Trust Game (Coleman, 1994; Dasgupta, 1988; Kreps, 1996). The left panel of Figure 5.1 shows the structure of a one-shot Trust Game without a reputation system. In this figure, T, R, P and S represent the material payoffs of the actors. In the Trust Game, a trustor first chooses whether to place trust or not. If the trustor decides not to place trust, the game is over and both players receive payoff P_i. For example, if a car owner on Turo decides not to rent out the car, the user cannot make use of the car and the car owner will not receive the payment from the guest. If the trustor decides to place trust (i.e., decides to rent out the car), the trustee (the car user) can choose to honor the trust he received (i.e.,

return the good in a good state) or abuse (i.e., not return the good, or not return the good in a good state)[20]. If the user honors, both car owner and user are better off than when the request is not accepted, because the car owner receives the payment and the user can drive a car. In that case, both players receive a payoff of R_i that is higher than the payoff P_i that they would receive if no trust was placed. If the user however decides to use the car in a way that induces costs for the owner, e.g. by smoking in the car or by using the car for moving houses, the user benefits more than when he or she would have refrained from smoking or using the car for moving, while the host would have been better off when the user would not have rented the car at all. In the Trust Game, this incentive structure is represented by the car owner receiving a payoff of S_1 that is smaller than P_1 and the user receiving a payoff T_2 that is larger than R_2[21,22]. A numerical example of the Trust Game is shown in the right panel of Figure 5.1.

Under the standard game-theoretic assumptions of rationality and selfishness, trustors in a one-shot Trust Game would anticipate that trustees will always abuse when they get the chance. Trustors would therefore not place trust. Yet this is not a Pareto-optimal outcome: both trustor and trustee would have been better off had the trustor placed trust and the trustee honored.

[20] To facilitate the distinction between trustors and trustees, we refer to the trustors with female pronouns and to the trustees with male pronouns.

[21] We further assume that $P_1 + P_2 > S_1 + T_2$ which avoids that alternating between being abused and abusing results in a higher payoff than not placing or receiving trust at all. Namely, it is implausible that the gains from abusing someone outweigh the material loss and immaterial harm suffered when abused by someone else.

[22] Most platforms for online exchange provide insurance that covers damage that can be proved to be caused by the user of the good. However, when the cause of the damage is unclear (e.g. in case of a small scratch on a car) or when it is hard to immediately detect the damage (e.g. damage to the gears or breaks of a car due to wear and tear) or when it is unclear who is responsible for the repair costs, insurances may not cover the damage and trust plays a role. Then, even with insurance in place, the Trust Game would still accurately represent material incentives on paid platforms with single-role reputation systems.

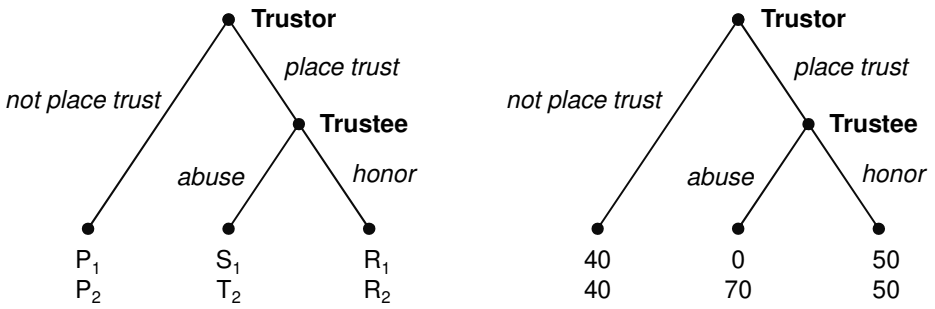

Figure 5.1: left panel: one round of the Trust Game without reputation system, where $R_1 > P_1 > S_1$, $T_2 > R_2 > P_2$ and $P_1 + P_2 > S_1 + T_2$. Right panel: numerical example used in the current experiment.

Free exchange systems as a Lending Game

To characterize the indirect (free) exchange situation, we introduce a new game, the "Lending Game" (Figure 5.2 shows a one-shot game without a reputation system). In the Lending Game, just as in the Trust Game, a trustor decides whether to place trust or not. For example, a Peerby user can decide to lend out a good for free to a neighbor or not. Lending out goods is beneficial to the borrower but comes at a cost to the owner of the good[23]. Even if the borrower honors the trust he received (i.e., returns the good in a good state), the owner needs to spend some time and effort arranging a meeting with the borrower and the owner cannot use the good for the time the user is borrowing it. If the borrower abuses the trust he receives (i.e., does not return the good, or does not return the good in a good state) the owner suffers a considerable material loss. In this situation, the trustor is always worse off when she places trust than when she withholds trust, even when the trustee honors the trust he received ($R_1 < P_1$).

[23] Online interactions on two-sided markets are often followed by offline interactions which in some cases induce additional (non-material) costs and benefits and informal gift exchange. For exchange, for Couchsurfing guests it is common practice to bring a small gift for the host and to spend some time together. However, since the extent to which provider and consumer spend time together differs to a large extent between platforms and is decreasing with technological advantages (e.g. smart locks), we focus on the material costs and benefits that are mediated by the platform.

The unique Nash equilibrium for one-shot play of this game is the same as in the Trust Game: the trustor does not place trust. In contrast to the Trust Game, however, this equilibrium is Pareto-optimal: Neither party could have been better off, without the other being worse off. Nonetheless, it does not maximize overall welfare, as the borrower (trustee) could have been made a lot happier at a small price to the owner (trustor).

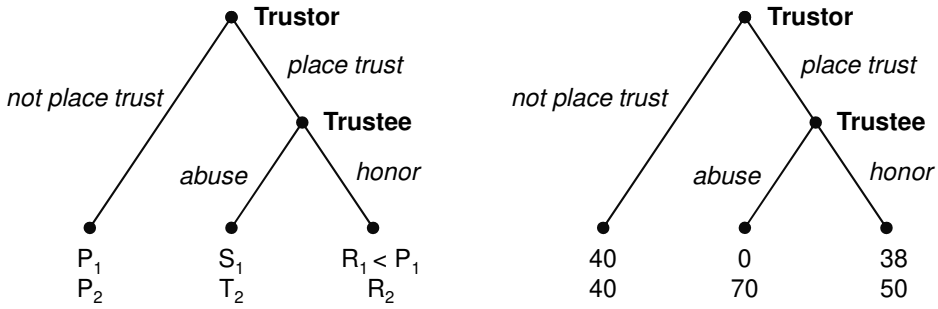

Figure 5.2: left panel: one round of the Lending Game without reputation system, where $P_1 > R_1 > S_1$ and $P_1 + P_2 > S_1 + T_2$. Right panel: numerical example used in the current experiment.

Reputation systems

Reputation systems are often mentioned as a very effective means for solving the social dilemmas described above if the trustee expects to face other trustors in the future. These systems collect, aggregate and distribute feedback about trustees' past decisions to trustors (Resnick et al., 2000)[24]. Reputation systems are thought to increase trust because they allow trustors to assess the trustworthiness of trustees (*learning*) and because they allow for rewards and punishments (*reciprocity*), thereby providing an incentive for trustees to show trustworthy behavior, and because they allow trustors to conditionally place trust (*control*). We will proceed with describing the learning, reciprocity and control mechanisms in more detail.

[24] In many real-life situations, provision of reviews is a social dilemma in itself, because review-writers do not benefit materially from writing the review, while everybody can access them (Diekmann et al., 2014). In the current paper we abstract away from this second dilemma, assuming that reputation information is provided automatically and without costs.

First, if we assume that there is variation in the extent to which individuals are willing to give up their own payoff to increase other people's payoffs, trustors can learn through reputation about the strategies and incentives of the trustees. Completely selfish individuals are only driven by the motivation to maximize their own payoff, while completely altruistic individuals only care about other people's payoffs. Most individuals will be somewhere between these two extremes. Trustees will thus differ in the extent to which they are willing to honor the trust they receive, but without any information about the incentives of trustees, trustors do not know which trustees will probably honor and which will likely not (Macy & Skvoretz, 1998). A trustee's reputation may help here. Trustors may use the available information on past trustee behavior to update their beliefs about the probability that the trustee is trustworthy (Buskens, 2003; Cook et al., 2005; Weigelt & Camerer, 1988), meaning that that trustee has only weak incentives for abusing trust in a given round. Trustees who have weak incentives for abusing trust are less likely to abuse trust in the future. If the probability that a trustee is trustworthy is high enough, the expected payoff of placing trust is higher than the expected payoff of withholding trust and the trustor will place trust.

An alternative explanation of why reputation information affects trustor behavior is indirect reciprocity. Trustors may be motivated to return helpful and harmful acts in kind, even if this is costly for the reciprocator (Stanca, 2009). In the case of online platforms, individuals often act only a single time with another buyer or seller, so it is often not possible to directly reciprocate a favor (Cabral & Hortaçsu, 2010), but reciprocity may also be indirect (Molm, Collett, & Schaefer, 2007). Evidence of such reciprocal behavior comes, for example, from public goods experiments where people were willing to punish free-riders, even when punishing is costly (Fehr & Gächter, 2000; Kuwabara, 2015). The meta-analysis by Balliet, Mulder and Van Lange (2011) shows that giving individuals the possibility to punish and reward others has a positive effect on overall cooperation.

Under the assumption of rationality all players know that trust will only be placed in trustees with a good reputation. Using backward induction, the players will identify the conditions under which cooperating (i.e. place trust and honor) results in a higher payoff than defecting (i.e. withholding and abusing trust). Choosing the cooperative option may be beneficial in the long term,

even for selfish individuals that are not intrinsically motivated to be cooperative if having a good reputation is necessary for being trusted in the future. Buskens and Raub (2002) refer to this as the 'control' mechanism. This mechanism predicts that when a reputation system is in place trustors are more willing to trust even trustees lacking any reputation. The prediction that individuals anticipate reputation effects is borne out in numerous studies (Benard, 2013; Buskens & Raub, 2002; Buskens & Weesie, 2000; Charness et al., 2011; Cheshire, 2007; Kroher & Wolbring, 2015; Rooks, Raub, & Tazelaar, 2006).

Single-role reputation systems

Single-role reputation systems only display information about the trustee's behavior in the trustee role. When there is no reputation system, the best strategy for egoistic trustees is to always abuse and for egoistic trustors to never place trust. However, a single-role reputation system changes the incentives of trustees in a way that enables the emergence of trust in paid exchange. If egoistic trustors assess the probability that a trustee is trustworthy as high enough, they will place trust in them, because the expected payoff of placing trust is higher than the expected payoff of not placing trust. The reciprocity argument leads to the same prediction, as rational trustors with a preference for reciprocity may perceive trustees who have honored a lot in the past as deserving more trust than an untrustworthy trustee. Likewise, these trustors might want to punish by not placing trust in trustees who have abused in the past. Consequently, in paid exchange, trustors are expected to place more trust in trustees who have honored more in the past. Knowing how trustors will respond to their reputation, trustees will select the option that gives them the highest payoff in the long term. Since their current decisions affect their future outcomes, they may strategically honor to maintain a good reputation (Robbins, 2016), which in turn increases the probability that a trustee will be trustworthy, as assessed by the trustor.

While single-role reputation is expected to affect trustor behavior in paid exchange, this reasoning does not lead to the same predictions for free exchange (or generalized exchange systems). Since providers in this case do not receive financial remuneration, they cannot gain anything by lending out their goods, so, under the standard assumptions of rationality and selfishness, they

are not expected to ever do so, even when they are confident that the user would act honorably.

Extensive research shows that single-role reputation systems have a strong, positive impact on trust in the Trust Game (Boero et al., 2009; Bolton et al., 2004a; Charness et al., 2011; Duffy et al., 2013; Fehrler & Przepiorka, 2013) and that general exchange systems can emerge when information about others' past behavior is available (Takahashi, 2000). A large number of studies on reputation systems show that trustors indeed use same-role reputation information as predicted: the more a trustee has honored in the past, the higher the probability that a trustor places trust (Buskens et al., 2010; Charness et al., 2011; Duffy et al., 2013), and the higher the prices they can charge (Ba & Pavlou, 2002; Dewally & Ederington, 2006; Dewan, Hsu, Dewanw, Hsuz, & Rogers, 2004; Resnick, Zeckhauser, Swanson, & Lockwood, 2006; Snijders & Zijdeman, 2004; see also the literature review in Diekmann, Jann, Przepiorka, & Wehrli, 2014).

Mixed-role reputation systems

Whereas single-role reputation systems only include previous decisions of the trustee in the trustee role, mixed-role reputation systems also convey information about decisions of those trustees made in the trustor role. We argue that trustors will use this extra information when deciding whether to trust a trustee, and more so in the Lending Game. Namely, the decisions individuals make in one role can be assumed to be driven by the same motivations as the decisions that they make in the other role in free exchange, but not in paid exchange. In free exchange both placing trust and honoring lead to a decrease in one's own payoffs, while it has a positive effect on the other player's payoff. Therefore, individuals who are more inclined to give up their own payoff for an increase in the payoff of others should, all else equal, be more likely to both place trust and to honor, because both behaviors entail a cost to oneself, while others benefit. More altruistic trustors may be more willing to take the loss associated with placing trust when the trustee honors. For example, Peerby users who are more likely to treat someone else's goods well may also be more likely to lend their goods to other people.

Experimental evidence indeed suggests that when subjects play both roles, behavior in the trustee role is often related to behavior in the trustor role

(Charness et al., 2011; Seinen & Schram, 2006). Other studies have found that an individual's cooperative behavior in one situation is related to that individual's trustworthiness in other situations (Chaudhuri & Gangadharan, 2007; Fehrler & Przepiorka, 2013).

We expect there to be a difference in the effect of the extra reputation information on trust in the Trust Game as compared to the Lending Game. In the multi-round Lending Game with stranger matching, in which there is no material incentive to place trust other than building up credit that may or may not be repaid one day, a history of placing trust in others (while playing in the trustor role) is a particularly strong signal of altruism and an act of generosity. The increased uncertainty in generalized exchange systems (i.e., the Lending Game) makes being generous more risky than in direct exchange systems (i.e., the Trust Game) (Molm, 2010; Molm, Schaefer, & Collett, 2009). Past trustful behavior (while playing in the trustor role) can therefore be considered a stronger signal of future trustworthiness or may evoke stronger reciprocity in the Lending Game than in the Trust Game, and we therefore expect that trustors rely more on this information in the Lending Game than in the Trust Game.

> **H1a**: Trustors are more likely to place trust in trustees who have placed more trust in previous rounds (while playing in the trustor role).
>
> **H1b**: The positive effect of placing trust in previous rounds (while playing in the trustor role) on being trusted in the present (while playing in the trustee role) is stronger in the Lending Game than in the Trust Game.

In a laboratory experiment by Charness et al. (2011), trustors used the information they received about the trustees' decisions to place or not place trust (while playing in the trustor role). However, the trustors in that experiment did not also receive information about the trustee's behavior in the role of trustee, so it is unclear whether trustors would still have used the other-role information had they also been given same-role information. Other studies have studied the effect of other-role reputation in scenarios outside the realm of trust problems. In the Helping Game or Image Scoring Game, where a donor can choose to give a gift to a receiver or not, after which the game ends without a choice for the recipient, donors were found to give more to receivers

who had helped more in the past and helping rates were higher when donors were able to learn about past behavior of receivers than when no reputation system was in place (Bolton, Katok, & Ockenfels, 2005; Engelmann & Fischbacher, 2009; Seinen & Schram, 2006). Other studies have described the emergence and importance of reciprocity in generalized exchange situations in the field, such as the Kula Ring (Malinowski, 1922) and in simulation studies (Takahashi & Mashima, 2003).

A mixed-role reputation system allows players to invest in their reputation by placing more trust. If opportunistic trustors believe that their trust will be reciprocated in the future if they place more trust in the present, they may place trust with the intention to evoke this reciprocity. Consequently, when information about past trustor behavior is public, trustors will be more inclined to place trust, even when they are not sure about the motives of the trustee.

> **H2a**: Both in the Trust Game and in the Lending Game, more trust is placed under a mixed-role reputation system than under a single-role reputation system.

Because past trustful behavior (in the trustor role) can be considered a stronger signal of future trustworthiness (in the trustee role) in the Lending Game than in the Trust Game, we expect that mixed-role reputation systems should increase trust levels more in the former than in the latter. Previous research on control mechanisms has shown that trustees may indeed seek to increase the probability of receiving trust by making generous choices in other, unrelated situations (Elfenbein, Fisman, & Mcmanus, 2012; Fehrler & Przepiorka, 2013).

> **H2b**: Mixed-role reputation systems increase trust (i.e. the decision of the trustor to place trust) more in the Lending Game than in the Trust Game.

By using games as abstract representations of markets, we reduce the differences between free and paid exchange to monetary differences and we assume that the incentives for trustors to place trust in paid exchange are larger than in free exchange. We acknowledge that this dichotomy may be too simplistic in real-life situations, because other, non-monetary incentives have been found to play a larger role in free exchange than in paid exchange (Frey & Oberholzer-Gee, 1997). Nonetheless, our study is general in the sense that

it tests the differential effect of reputation systems across settings that differ in the extent to which the benefits of trusting offset the costs, regardless of the monetary or nonmonetary origins of these costs and benefits. We do not deny other differences, but merely seek to theoretically and experimentally isolate the lack of immediate gains from lending in free exchange systems as an important trigger of reciprocal behavior.

Methods

The experiment was conducted at Utrecht University between November 29, 2017 and December 12, 2017[25]. Four pilot sessions were organized on November 15 and 16, 2017. Using a laboratory experiment allows us to study reputation in isolation and to make inferences about causality. Subjects were students of Utrecht University and were recruited using the internet recruitment system ORSEE (Greiner, 2015). All laboratory sessions were computerized using Z-tree 3.6.7 (Fischbacher, 2007). The subjects received printed instructions in English. There were 228 subjects in total, divided over twelve sessions. Each session consisted of two or three matching groups of eight, ten or twelve subjects each. After reading the instructions and before the start of the experiment, subjects had to answer some questions to test their comprehension of the instructions. At the end of the session subjects completed a short survey that included demographic questions, a question on their experience with game theory and questions on subjects' motivations and beliefs about other participants.

Each session took about 60 minutes (including the time for reading instructions). The average payment was 8.89 euro (100 points were equivalent to 0.60 euro). There were 156 (68.4%) subjects identifying as female, 69 (30.3%) identifying as male, and one subject indicating 'other gender'. Age ranged from 18 to 66 and was on average 23.4 (SD = 6.0).

The participants were informed that the session consisted of 36 rounds. At the beginning of each round, half of the participants were assigned the role of

[25] At the time of the experiment, submitting studies to the faculty's ethical review board was not required.

trustor (role A) and half of the participants were assigned the role of trustee (role B). We arranged the draws such that each person was a trustor six times and a trustee six times in each block of 12 rounds (cf. Charness et al., 2011). At the end of the session, each subject had played 18 times in the trustor role and 18 times in the trustee role[26]. This procedure ensured that the number of times all subjects acted in the roles of trustor and trustee did not vary too much at any point in time, but at the same time that subjects did not know for sure what role they would play in the next round. The participants were randomly and anonymously matched in pairs within their matching group at the beginning of each round. They knew they were divided in two or three different groups (depending on the number of participants in the session) and they knew they only interacted with the people in their group. No communication between the participants was allowed during the experiment. The full instructions can be found in the Appendix. At the end of the round, after the players made their decisions, both players were informed about their own choice and payoff of that round, the choice and payoff of their partner in that round and their own total payoff.

The experiment had a 2 (Trust Game versus Lending Game) by 2 (single-role reputation versus mixed-role reputation) design. There were three sessions in each of the conditions, for a total of 12 sessions. In the Trust Game, the trustor chose at the start of each round between RIGHT (not place trust) and DOWN (place trust). If the trustor chose RIGHT, the round was over and both players received a payoff of 40. If the trustor chose DOWN, the trustee had to make a choice between RIGHT (abuse) and DOWN (honor). If trustee chose to abuse, the trustor received 0 and the trustee received 70. If the trustee chose to honor, both players received a payoff of 50. The only difference between the Lending Game and the Trust Game was the payoff of the trustor if the trustee decided to honor. Instead of 50, the trustor received a payoff of 38 in the Lending Game when the trustee honored.

[26] Lotteries at the beginning of each round determined which subjects in a matching group were to play in the role of trustor. The probability that a subject's lot was drawn decreased with the number of times that subject acted in the trustor role in previous rounds in the same block and increased with the number of times the other subjects acted in the trustor role in the block.

In the single-role reputation conditions, before making a choice, the trustor was informed about the total number of decisions the trustee made in the trustee role and the number of times the trustee chose to honor in the past (while playing as a trustee). In the mixed-role reputation condition, the trustor was also informed about the number of past decisions trustee had made as trustor and the number of times he had placed trust (while playing as a trustor). Note that trustees received no information about the trustor's past behavior in any of the conditions.

Analysis and results

Dependent variable

The dependent variable in our analyses is trust, which is operationalized as the decision of the trustor to place trust or not. This variable does not consider behavior of the trustee, i.e. the decision to honor or abuse.

Independent variables

Our first independent variable is the treatment to which a subject was assigned: Trust Game with single-role reputation (TG, SRR) or mixed-role reputation (TG, MRR) or Lending Game with single-role reputation (LG, SRR) or mixed-role reputation (LG, MRR).

Next, we constructed reputation variables based on behavior in the trustor role and the reputation based on behavior in the trustee role. We constructed two variables per type of reputation information: reputation availability and reputation score. The former is a dichotomous variable indicating whether the trustee already made decisions in the relevant role ('Made trustee decision' and 'Made trustor decision'). This allows us to evaluate the effect of not having a reputation vs having a positive or a negative reputation. The second variable reflects the fraction of decisions in the trustee (trustor) role in which the player chose the prosocial option in the relevant roles: honor in the trustee role ('Honor rate history') and place in the trustor role ('Place rate history').

Control variables

We control for the personal experiences of the trustor. After every round and in every experimental condition, all subjects are informed about the decision in the current round of their current interaction partner. This information about the past experiences of the trustor is summarized in two variables. The first variable 'Experienced trustfulness', reflects how often a trustor was trusted by others, as a fraction of the total number of times she acted in the trustee role in earlier rounds. Likewise, the variable 'Experienced trustworthiness' reflected the number of times the partner of a trustor honored as a fraction of the total number of times she placed trust in earlier rounds. We also control for round number.

Analytical strategy

To test our hypothesis about individual-level behavior (Hypotheses 1a and 1b), we ran multilevel logistic regressions with trust as the dependent variable and with random intercepts for subjects and groups to accommodate the multi-level structure of the data.

To test Hypothesis 1a we ran the analysis separately for the four conditions (models 1 through 4 in Table 5.2). We included the relevant reputation variables as independent variables: the models for the single-role reputation conditions only included the reputation variables for past behavior in the trustee role, while the models for the mixed-role condition also included the variables for past behavior in the trustor role. We controlled for the personal experiences for the trustor. The model did not converge for the Trust Game with single-role reputation, so instead we performed a multilevel regression with only random intercepts for subjects in that condition.

To test Hypothesis 1b, we performed two multilevel logistic regressions on the mixed-role reputation condition only (models 5 and 6 in Table 5.2). Game type (TG or LG), round number, two personal experience indicators and the four reputation variables are included as predictor variables.

To test Hypothesis 2a about the differences in trust between the four conditions, we use Wilcoxon rank-sum tests that treat each matching group as one observation to compare the trust rates between the condition. The results

of these tests are reported in section 4.5. As a second test of Hypothesis 2a and to test Hypothesis 2b, we again use the multilevel logistic regression approach, which also allowed us to test the interaction between game type and reputation type. Game type (Trust Game versus Lending Game) and reputation type (single-role versus mixed-role reputation) are included as independent variables, as well as the interaction between the two. Round number is included as a main effect and as an interaction with reputation condition. The results of this regression are reported in model 7 in Table 5.2. The conclusions from the rank-sum tests are the same as those of the regressions.

We additionally ran a number of analyses with the decision of the trustee (honor or abuse) as the dependent variable. A description of these analyses and the results can be found in the Appendix.

Descriptive statistics and results rank-sum tests

Table 5.1 shows the average rates of trust (i.e. fraction of interactions in which the trustor placed trust), trustworthiness (i.e. fraction of interactions in which the trustee honored trust, if the trustor placed trust). and the 'success rate' per condition: the fraction of interactions in which trust is placed and honored.

Table 5.1*:* Trust and trustworthiness rate, by condition (standard deviations in parenthesis)

Condition		N	Trust rate	N	Trustwor-thiness	N	Rate of both, by pair
TG	*SRR*	1080	.433 (.496)	468	.767 (.423)	1080	.332 (.471)
	MRR	972	.416 (.493)	404	.599 (.491)	972	.249 (.250)
	Total	2052	.425 (.494)	872	.689 (.463)	2052	.293 (.455)
LG	*SRR*	972	.112 (.315)	109	.459 (.501)	972	.051 (.221)
	MRR	1080	.200 (.400)	216	.690 (.464)	1080	.138 (.345)
	Total	2052	.158 (.365)	325	.612 (.488)	2052	.097 (.296)
Total		4104	.291 (.455)	1197	.668 (.471)	4104	.195 (.396)

The Wilcoxon rank-sum test shows that, in line with our assumption, the trust rate is significantly higher in the Trust Game than in the Lending Game (z = 4.135, p < .001). There is no significant overall difference in the trust rate

between the single-role reputation condition and the mixed-role reputation condition (z = -0.626, p = 0.531), rejecting Hypothesis 2a. When analyzing the two game types separately, we find that there is no significant difference between the two reputation conditions when only looking at the Trust Game (z = .000, p = 1.000), but in the Lending Game, the trust rate is higher in the mixed-role reputation condition than in the single-role reputation condition (z = -2.074, p = .038), consistent with Hypothesis 2b.

Figure 5.3 shows how the trustee's previous decisions and the probability that the trustor places trust in the current round are related. We hypothesized that the more a trustee has placed trust in the past (while playing in the trustor role), the higher the likelihood that the trustor placed trust in the trustee (Hypothesis 1b). The figure suggests that this is indeed the case.

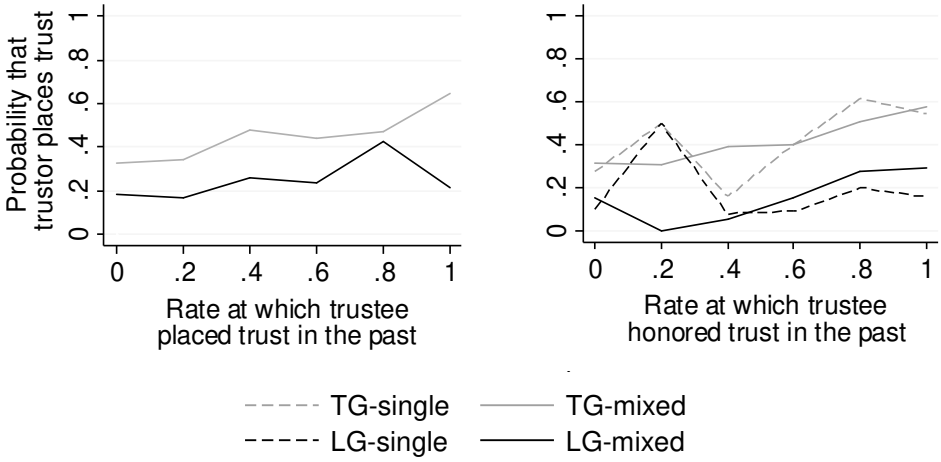

Figure 5.3: Relation between previous decisions made by the trustee and the probability that trust is placed by the trustor in the current round. Left panel: previous decisions made by the trustee in the trustee role (honor versus abuse) on the x-axis. Right panel: previous decisions made by the trustee in the trustor role (place versus not place trust) on the x-axis.

Figure 5.4 shows how the trust rate changes over time in the various conditions (left panel) and how the trustworthiness rate changes over time (right panel). Both panels show the "end game effect", the familiar tendency for trust

placement on honoring to decrease as the known end of the session is approached.

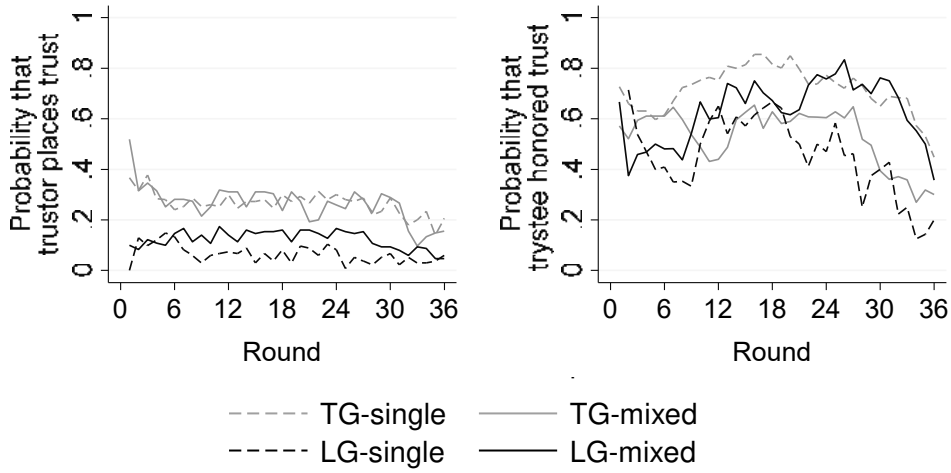

Figure 5.4: Left panel: probability that the trustor places trust over the rounds per condition. Right panel: probability that the trustee honors over the rounds per condition. Moving average over five rounds.

Results regressions: determinants of trust

Individual trustor decisions

Contrary to what we hypothesized (H1a), we find that the trustee's historical trust rate ('Place rate history') does not significantly affect the probability that a trustor placed trust in the mixed-role reputation conditions if we separately look at the Trust Game and the Lending Game (models 2 and 4 in Table 5.2). This means that trustors are not more likely to place more trust in trustees who have placed more trust in the past (while playing in the trustor role). However, these tests may lack the necessary power, because when we combine the two games in one analysis, we find that the historical trust rate significantly predicts trustor behavior (model 5 in Table 5.2), supporting Hypothesis 1a. A 0.1 increase in the fraction of trust placed on average leads to an increase of 0.01 in the probability that the trustor places trust in both the Trust Game and the Lending Game. The effect of the historical fraction of trust placed by the trustee (while playing in the trustor role) on the probability that a trustor places

trust is not stronger in the Lending Game than in the Trust Game ('Place rate history * LG', model 6 in Table 5.2), so we reject Hypothesis 1b.

The coefficients for availability of first-mover role reputation ('Made trustor decision') are not significant, also not when we combine the Trust Game and the Lending Game (b = -0.354, z = -1.18, p = 0.239). Because all trustees have played in the trustor role after the sixth round of the game, this coefficient is based on the first six rounds only, so we cannot conclude that in all rounds of the game having no reputation and having a bad reputation have a similar effect on the probability of receiving trust.

In all four conditions except the Lending Game with single-role reputation information, a trustee's historical trustworthiness rate ('Honor rate history') has a significant effect on the trustor's decision whether to place trust. Trustors more often place trust in trustees who have honored more in the past (while playing in the trustee role). If the fraction with which a trustee chose to honor in the past (while playing in the trustee role) increases with 0.1, the probability that a trustor placed trust on average increases with 0.04 and 0.03 in the Trust Game with single-role and mixed-role reputation respectively, and with 0.02 in the Lending Game with mixed-role reputation. While the coefficient of reputation score is positive, the coefficient of having a trustee reputation score ('Made trustor decision') is negative and significant in the same three conditions. This means that trustees who have abused trust relatively often (while playing in the trustee role) are trusted less often than trustees who have not had the chance to choose between honoring and abusing trust.

The rate at which a trustor has experienced trust ('Experienced trustfulness') does not significantly affect the probability of placing trust in any of the conditions, but the more trustworthiness the trustor has experienced ('Experienced trustworthiness'), the higher the probability of placing trust. In both Trust Games the probability that a trustor places trust increases with about 0.03 with every 0.1 increase in the rate of experienced trustworthiness. In the Lending Game with single-role reputation this probability increases with 0.01 with every 0.1 increase in the rate and in the Lending Game with mixed-role reputation the probability increases with 0.03.

Table 5.2: Determinants of trust

Independent variables	Model 1: TG, SRR	Model 2: TG, MRR	Model 3: LG, SRR	Model 4: LG, MRR	Model 5: MRR	Model 6: MRR	Model 7: All conditions
Treatment							
LG	-	-	-	-	-1.281***	-1.186***	-2.583***
					(0.266)	(0.338)	(0.301)
MRR	-	-	-	-	-	-	-0.085
							(0.304)
LG * MRR	-	-	-	-	-	-	1.198**
							(0.396)
Round number	-0.023**	-0.034***	-0.040**	-0.023**	-0.032***	-0.031***	-0.026***
	(0.007)	(0.008)	(0.013)	(0.009)	(0.006)	(0.006)	(0.006)
Reputation trustee							
Made trustee decision	-1.108***	-1.207***	0.491	-0.785**	-0.753***	-0.149	-
	(0.257)	(0.335)	(0.311)	(0.300)	0.203	(0.225)	
Honor rate history	2.134***	1.654***	0.048	1.684***	1.282***	0.839**	-
	(0.244)	(0.270)	(0.389)	(0.311)	(0.183)	(0.266)	
Made trustor decision	-	-0.194	-	-0.055	-0.354	0.298	-
		(0.442)		(0.393)	(0.300)	(0.364)	
Place rate history	-	0.534	-	0.766	0.668**	-0.429	-
		(0.327)		(0.406)	(0.247)	(0.514)	

(Continued on next page)

Table 5.2: continued

Independent variables	Model 1: TG, SRR	Model 2: TG, MRR	Model 3: LG, SRR	Model 4: LG, MRR	Model 5: MRR	Model 6: MRR	Model 7: All conditions
Trustor experience							
Experienced trustfulness	-2.288	2.225	-1.599	1.031	-1.514	-1.752	-
	(2.476)	(2.638)	(1.031)	(0.922)	(1.279)	(1.408)	
Experienced trustworthiness	1.449***	1.903***	3.410***	2.457***	2.122***	2.564***	-
	(0.290)	(0.355)	(0.566)	(0.329)	(0.235)	(0.252)	
Interactions							
MRR * Round number	-	-	-	-	-	-	-0.005
							(0.008)
Place rate history * LG	-	-	-	-	-	-1.039	-
						(0.665)	
Constant	-1.106	-0.652	-3.481	-2.992	-0.763	-1.533	0.107
	(0.339)	(0.417)	(0.462)	(0.430)	(0.329)	(0.363)	(0.220)
Observations	1080	972	972	1080	2052	2052	4104

Note: * $p<0.05$, ** $p<0.01$, *** $p<0.001$ (two-tailed tests). Standard errors in parentheses. TG refers to Trust Game, LG refers to Lending Game. SRR refers to single-role reputation system and MRR to mixed-role reputation system. Made trustee decision (Made trustor decision) indicates whether the trustee has made a decision in the trustee (trustor) role yet. Honor rate history (Place rate history) indicates the rate at which the trustee has honored (placed) trust in the past. Experienced trustfulness (Experienced trustworthiness) is the rate at which the trustor experienced trust placed (honored).

Overall trust rate

Confirming the results of the rank-sum test, the regression shows that the probability that a trustor placed trust is 0.35 higher in the Trust Game than in the Lending Game ('LG', model 7 in Table 5.2), and that there is no effect of type of reputation system in the Trust Game (rejecting Hypothesis 2a). The interaction between game type and reputation system type ('LG * MRR') is significant, providing further support for Hypothesis 2b. In a separate analysis for two game types, we find that after controlling for round number, in the Lending Game the probability that trust is placed is 0.10 higher in the mixed-role reputation condition than in the single-role reputation condition (b = 1.129, z = 2.81, p = 0.005), but that the effect of reputation system type is not significant in the Trust Game (b = -0.154, z = -0.62, p =0.537).

Discussion

The focus of this chapter has been on the increasingly common scenario in which individuals are sometimes provider and sometimes user of a borrowed good (e.g. eBay, Couchsurfing, Kickstarter). We have argued that in these settings, a new form of reputation building emerges: individuals lend valuable property to unknown others, risking exploitation, in the expectation that when they themselves call on help that requires another person's willingness to trust them, they will more likely receive it themselves. This argument was largely supported in a laboratory study. We found that, indeed, mixed-role reputation systems increased trust levels, and only for scenarios of lending, not letting. We also found evidence that providers are more likely to place trust in users who themselves have placed more trust in the past. We expected the effect of the mixed-role reputation system in paid exchange to be smaller, and in practice we found no effect at all.

This chapter thus provides an explanation for the numerous examples of free platforms that function well and attract both borrowers and lenders. Why would lenders suffer a sure loss if they do not get anything in return? Of course, the non-material rewards, other-regarding preferences and social aspects of these exchanges assumed away in narrowly formulated game-theoretic models may well be part of the explanation (e.g. Grybaitė & Stankevičienė, 2016;

Hamari, Sjöklint, & Ukkonen, 2015; Hars & Ou, 2002). Altruists just generously share their property with others, at the risk of being taken advantage of. We readily acknowledge that there are other contextual factors that affect willingness to rent or lend. The aim of this study was not to realistically model free and paid exchange systems, but to test the differential effect of reputation systems on trust in a setting in which the benefits of trusting offset the costs versus a one where they do not. We propose that in addition to sheer goodness, mixed-role reputation systems play an important role in facilitating free exchange. Namely, they do not only provide information on the provider's behavior in the provider role, but also on that provider's behavior in the user role. This extension of the classic, single-role reputation system allows providers to get something in return for their generosity, explaining why even purely egoistic providers would lend out their goods. Our findings confirm our theoretical claim that the norm of reciprocal lending emerges in settings where lending does not provide an immediate payout and not in those where it does.

While the results of the experiment are in line with our theory that other-role reputation information helps trustors to identify trustworthy trustees, we acknowledge that there may be alternative explanations that could explain this finding. One interesting alternative explanation was suggested by a reviewer and also points to the possibility for reciprocation in mixed-role reputation systems, but operates through a somewhat different mechanism. Namely, perhaps the extra reputation information may change the structure of the game in a way that makes it worthwhile even for purely egoistic trustors to play a tit-for-tat strategy (i.e. conditionally cooperate) if all others do the same. According to this argument, a single-role reputation system is a necessary and sufficient condition for the emergence of trust in paid exchange, but not in free exchange. Under the assumption that all trustors in the Lending Game are selfish, they can maximize their payoff by cooperating conditionally. While knowing that a trustee is probably trustworthy is enough motivation to place trust in paid exchange, this is not the case in free exchange. In free exchange, trust may emerge if trustors condition their decision on the past behavior of the trustee in the trustor role. If all trustors follow this strategy, being trustful will result in receiving more trust in the future and can thus be a better strategy than never placing trust. While this argument explains why trustors would place trust in free exchange when there is a mixed-role reputation system, it raises

the question why trustors would condition their decision on the trustee's decision in the trustor role. If placing trust is good for an individual's reputation, and if the reputational benefits are larger than the costs of cooperating, why would trustors not always do so, regardless of the reputation of the trustee. A game-theoretic analysis may formally specify whether and how this may exactly work, which we leave for future research. Further experimental research would be needed to test if trustors in practice adopt such a strategy.

Why did we find no effect of mixed-role reputation systems on trust in paid exchange, even if we predicted to find a positive effect, albeit smaller? This might have to do with the overall trust levels in our experiment, which were low compared to other experiments on the Trust Game and the Helping Game with reputation systems (Boero et al., 2009; Bolton et al., 2005; Buskens et al., 2010; Charness et al., 2011; Seinen & Schram, 2006). We predicted that the other-role information would only matter if the trustee placed enough trust in the past, and when the trustor assessed the probability that the trustee is trustworthy to be low. The relatively low trust rate suggests that there were very few occasions in which these conditions were applicable, meaning that the extra information was of little added value and thus did not change overall trust levels.

Two other questions that remain unanswered are why the trust rate was so low in this experiment, and why the reliance of the providers on the extra reputation information was limited. It may be that experimental subjects' bounded rationality limits their ability to fully comprehend the consequences of their choices (Buskens et al., 2010; Kahneman, 2002; Simon, 1955). Knowing why trustees who have honored more trust in the past will be more trustworthy is cognitively less demanding, as it requires understanding fewer steps than understanding why one should place more trust in trustful providers, or why placing trust in the present may have consequences for the rate at which one will be trusted in the past.

These open questions notwithstanding, altogether, our findings suggest that in modern, Internet-enabled marketplaces, reputation systems allow individuals to not only develop a classic reputation for someone who *can* be trusted, but also to become known as someone who has paid their dues to the community, and therefore *deserves* to be trusted.

APPENDICES

Appendix chapter 2

Simulations

Model

The goal of the simulation is to illustrate under what conditions inequality between renters with different ethnicity may increase due to the presence of a reputation system. We simulated a sequence of rounds *t* in each of which two renters ($i = 2$) who have a different ethnicity ($E_i = 0$ for a renter with majority ethnicity, $E_i = 1$ for a renter with minority ethnicity), both submit a rental request to different owners. We simulate how the outcome of that request depends on the renter's ethnicity and reputation.

The outcome of a request in a given round, A_{it}, equals 1 if the request is accepted and 0 otherwise. We assume that a renter receives a positive review after every accepted request. The number of reviews a renter has thus solely depends on the sum of previously accepted requests. The probability that the request is accepted in a given period is $0 \leq P_{it} \leq 1$. We use a logistic regression function to simulate the probability: $P_{it} = \frac{\exp(L_{it})}{1+\exp(L_{it})}$, where L depends on the ethnicity and reputation of the renter: $L_{it} = G - E_i * D * (1 + R * \sum_1^{t-1} A_{2t})^{-C} + R * \sum_1^{t-1} A_{2t}$.

G is the level of trust in ethnic majority renters. We assume that there is discrimination for actors without reputation: requests from renter 1 (belonging to the ethnic majority) are accepted more often than requests from renter 2 (who belongs to the ethnic minority): $D > 0$. The more reviews a renter has, the higher the probability that the request is accepted: $R > 0$. Following the argument that discrimination in the platform economy is mostly statistical discrimination, owners may pay less attention to ethnicity when the renters have more reviews. This 'compensation effect' is denoted as $C \geq 0$. The larger C is and the more reviews renter 2 has, the smaller the influence of discrimination. C is operationalized as a power term to avoid that minority renters will eventually have a higher acceptance probability than majority renters.

Table A2.1 shows the parameter settings used in the simulations. The initial level of trust in minority and majority renters can be derived from filling in different values of G and D in logistic regression function. The intial level of trust in renter 1 is between 12% and 88% the intial level of trust in renter 2 is between 5% and 73%. In total we have 5 (initial level of trust in the market) * 2 (discrimination) * 4 (importance of reputation) * 11 (compensation effect) = 440 conditions. We ran 1000 iterations per condition, resulting in 440,000 datapoints.

Table A2.1: Parameter settings used in the simulation

Parameter	Values tested
Number of rounds	$t = 10$
Number of renters	$i = 2$
Initial level of trust in majority	$G \in \{-2, -1, 0, 1, 2\}$
Discrimination	$D \in \{1, 2\}$
Importance of reputation	$R \in \{0.5, 1, 1.5, 2\}$
Compensation effect	$C \begin{Bmatrix} 0.0, 0.1, 0.2, 0.3, 0.4, 0.5, \\ 0.6, 0.7, 0.8, 0.9, 1.0 \end{Bmatrix}$

We calculated the initial level of trust in minority renters $G - D$ (ranging from -4 to 1) and we standardized the four parameters to be able to compare their coefficients. Using a logistic regression with the outcome (place trust or not) as the dependent variable, and all interactions between the ethnicity of the renter, the round number, and all parameters we predicted the probability that trust was placed for each parameter combination, round and renter. We then calculated the size of the ethnic gap, which is defined as the difference in the probability that majority and minority renters are trusted in a given round and parameter setting.

Simulation results

Figure A2.1 shows the relation between the ethnic gap and the input parameters over the rounds. Especially the lines in the graph of the initial level of trust in the ethnic minority are interesting. At low levels of trust, inequality increases over the rounds, while at higher levels of trust, the ethnic gap becomes smaller over the rounds. Discrimination has a positive effect on the ethnic gap, but does not affect the change in the ethnic gap over the rounds. The lines of the

importance of reputation and the compensation effect are more or less horizontal and at the same level, which means that they do not affect the (change in) the ethnic gap.

The results of a regression analyses with the ethnic gap as the dependent variable and the parameters and round number as independent variables are in Table A2.2. We added the interaction between round number and the four parameters one-by-one to investigate which parameters have the strongest influence on the change in the ethnic gap. The results of these regressions are in Table A2.2. As we would expect, the coefficients of the initial level of trust in minority renters and compensation effect are negative, which means that the ethnic gap is smaller when these variables have higher values. More discrimination and a higher importance of reputation lead to a smaller ethnic gap.

The coefficient of the round number is negative, which indicates that on average the ethnic gap decreases over time. However this change in the ethnic gap between round 1 and round 10 is negligible because the size of the change is very small. In the first round the average probability that ethnic majority renters are trusted is 22.3 percent point higher than the probability that ethnic minority renters are trusted. In the last round this gap is 22.2 percent point.

The negative coefficient of the interaction between the round number and the initial level of trust in ethnic minority renters indicates that for low levels of trust the ethnic gap increases over the rounds. For higher initial trust levels this effect is reversed: the ethnic gap decreases over time. This can be explained by a combination of factors. When intial trust increases, it becomes easier for ethnic minority renters to acquire reviews and to profit from their reputation. However, a higher initial trust rate also implies that there is less room for and increase in inequality as the maximum probability that a request accepted is fixed. The coefficients of the interactions between the round number and the other three parameters are very small, so they have a minimal effect on the change in the ethnic gap over time.

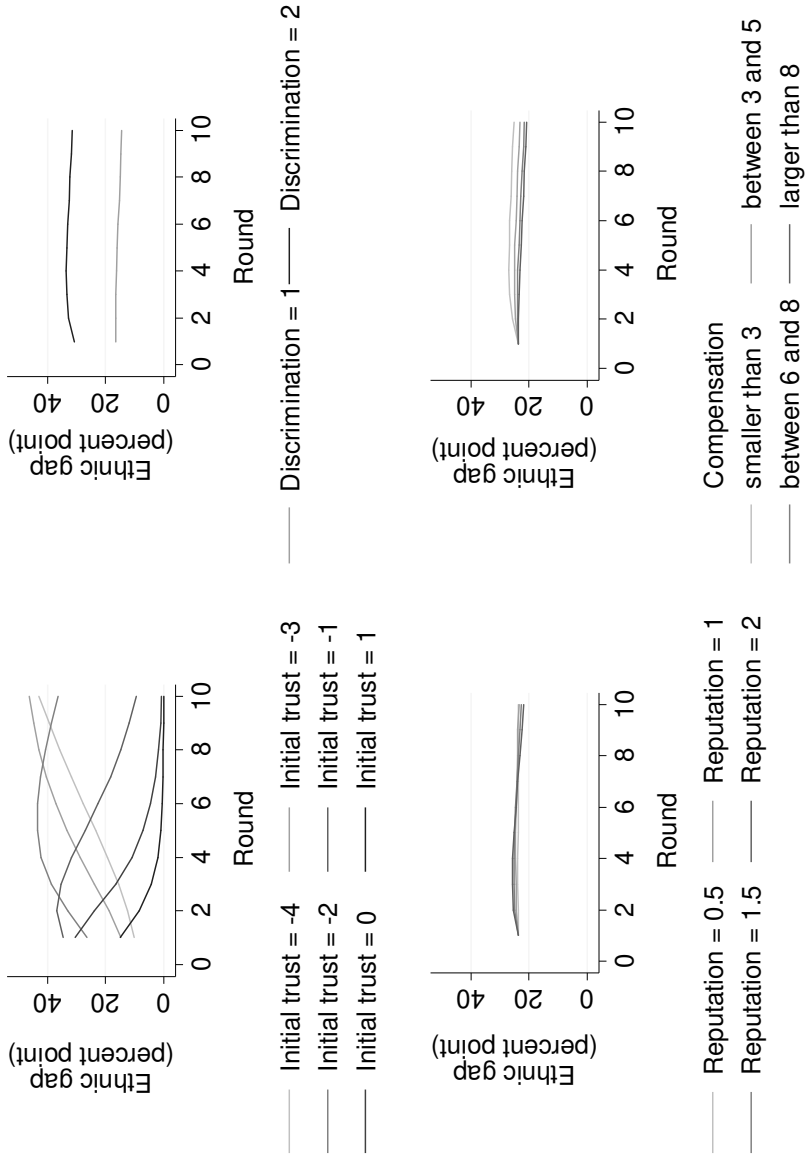

Figure A2.1: Change in the ethnic gap (difference in the probability that ethnic majority and minority renters are trusted) over the rounds, depending on trust in minority (top-left), discrimination (top-right), reputation effect (bottom-left) and compensation effect (bottom-right).

A

Table A2.2: Results of a linear regression with the ethnic gap (difference in probability that majority and minority renters are trusted) as the dependent variable. A positive number indicates a larger ethnic gap.

	Model 1	Model 2	Model 3	Model 4	Model 5
Main effects					
Round number	-0.002*	-0.002*	-0.002*	-0.002*	-0.002*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Initial level of trust in minority renters (G - D)	-0.074***	0.064***	-0.074***	-0.074***	-0.074***
	(0.002)	(0.004)	(0.002)	(0.002)	(0.002)
Discrimination (D)	0.060***	0.060***	0.054***	0.060***	0.060***
	(0.002)	(0.002)	(0.005)	(0.002)	(0.002)
Importance of reputation (R)	0.001	0.001	0.001	0.009	0.001
	(0.002)	(0.002)	(0.002)	(0.005)	(0.002)
Compensation (C)	-0.013***	-0.013***	-0.013***	-0.013***	-0.005
	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)
Interactions					
Round * Initial level of trust in minority renters	-	-0.025***	-	-	-
		(0.001)			
Round * Discrimination	-	-	0.001	-	-
			(0.001)		
Round * Importance reputation	-	-	-	-0.001	-
				(0.001)	
Round * Compensation	-	-	-	-	-0.001*
					(0.001)
Constant	0.250***	0.250***	0.250***	0.250***	0.250***
	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)
Observations	4,400	4,400	4,400	4,400	4,400
R-squared	0.370	0.528	0.371	0.371	0.371

Note: The following variables are standardized: initial level of trust, discrimination, reputation, compensation. Standard errors in parentheses.
*** $p<0.001$, ** $p<0.01$, * $p<0.05$

Additional regression results

Table A2.3: Results of the multilevel cross-classified linear probability model of outcome of the request: accepted (0/1) with request level control variables

Independent variables	Model 1
Threshold	0.133**
	(0.034, 0.302)
Main effects	
Renter # positive ratings (ref. cat. = no reviews)	0.037***
	(0.020, 0.055)
Renter # negative ratings	-0.031
	(-0.089, 0.026)
Renter request #	-0.009***
	(-0.014, -0.004)
Control variables	
Rental characteristics	
# Rental days	0.008*
	(0.001, 0.015)
Cumulative total # of interactions (/1000)	-0.007*
	(-0.015, 0.001)
Renter characteristics	
Renter age	0.003***
	(0.002, 0.004)
Renter member #years	0.008
	(-0.025, 0.041)
Owner characteristics	
Owner # of positive reviews	-0.002
	(-0.007, 0.004)
Owner # of negative reviews	-0.027
	(-0.060, 0.007)
Owner age	0.006***
	(0.004, 0.008)
Owner member #years	-0.069***
	(-0.103, -0.035)

(Continued on next page)

A

Table A2.3: continued

Independent variables	**Model 1**
Motorcycle characteristics	
Day price	0.000
	(-0.002, 0.001)
Motorcycle # of positive reviews	0.002
	(-0.001, 0.005)
Motorcycle # of negative reviews	0.016*
	(0.001, 0.031)
Motorcycle age	-0.003*
	(-0.006, 0.000)
Motorcycle weight	0.000
	(-0.001, 0.000)
Motorcycle CC	0.000
	(0.000, 0.000)
Motorcycle HP	0.000
	(-0.001, 0.001)
Variance intercept renter-level	0.002***
	(0.001, 0.005)
Variance intercept owner-level	0.042***
	(0.035, 0.050)
N	7181

Note: 95% Confidence Interval in parentheses. *** $p<0.001$, ** $p<0.01$, * $p<0.05$ (one-tailed tests).

Appendix chapter 3

Consent form

The goal of this experiment is to better understand how people make decisions in groups. This experiment is conducted by researchers from Utrecht University, the Netherlands.

If you agree to be in this study, you will play two games in which you will be paired with seven other individuals. The decisions that you and the other participants make in the games will determine how much you earn. If you complete the experiment, you will receive a payment of $2.00 plus a bonus that varies between $3.60 and $12.80. You will receive your payment as a bonus payment on to the HIT on MTurk that you used to sign up for the study. It may happen that you finish reading the instructions, but after that cannot be assigned to a group to play the game. In that case you will receive a payment of $2.

There are no known risks associated with your participation in this research beyond those of everyday life. Taking part in this study is voluntary. Confidentiality of your research records will be strictly maintained. We collect your Amazon Mechanical Turk ID to calculate your bonus payment. After the study has been completed, we will separate your MTurkID from the experimental data so answers cannot be traced back to you. The data from the study will be kept at least 10 years.

If there is anything about the study or taking part in it that is unclear or that you do not understand, you may contact the researcher (Judith Kas, j.kas@uu.nl).

Instructions for participants

Please carefully read all of the information that follows.

In this experiment, you will earn points by making decisions in "games". How much you earn depends on your decisions, the decisions of others and on chance. The experiment will last about 45 minutes. At the end of the experiment, you will be paid $2.00, plus a bonus of $1.00 for every 100 points you earned in the games. Please note that this is an interactive experiment. That means that you may have to wait while other participants are making decisions.

We will now explain the rules of the game. After you have read the instructions you will be asked to answer a few questions that help you evaluate your understanding of the game.

The rules of the game

This game is played in groups of 8 participants, four participants in the role of A (A1, A2, A3 and A4) and four participants in the role of B (B1, B2, B3 and B4). Every participant is randomly assigned to a role and number. Throughout the game, all participants keep their role and number.

We understand that sometimes studies claim that participants will actually interact in real time, but in fact use simulated other people ("bots"). As a policy, we do not do this. In the game you will really be interacting with 7 other people.

The game proceeds in rounds. There are 8 rounds. In each round only one of the As is active and interacts with the Bs. In round 1, it is A1's turn to interact with the Bs; in round 2, it is A2's turn, and so on such that in round 5, it is again A1's turn. How many points the active A and the Bs get depends on their choices.

The active A (simply called "A" in Figure A3.1) chooses either "RIGHT" or "DOWN."

- If A chooses RIGHT, A gets 30 points and all four Bs get 30 points as well.
- If A chooses DOWN, A must select one of the four Bs.
- If A chooses DOWN and selects one of the Bs, the selected B chooses "RIGHT" or "DOWN."

If the selected B chooses DOWN, A and the selected B get 50 points each. If the selected B chooses RIGHT, A gets nothing (0 points) and the selected B gets 80 points. In either case, the other Bs – the ones that were not selected by the A at play – get 30 points each. The As who are not active in a round get 30 points in that round, just as the Bs that were not selected.

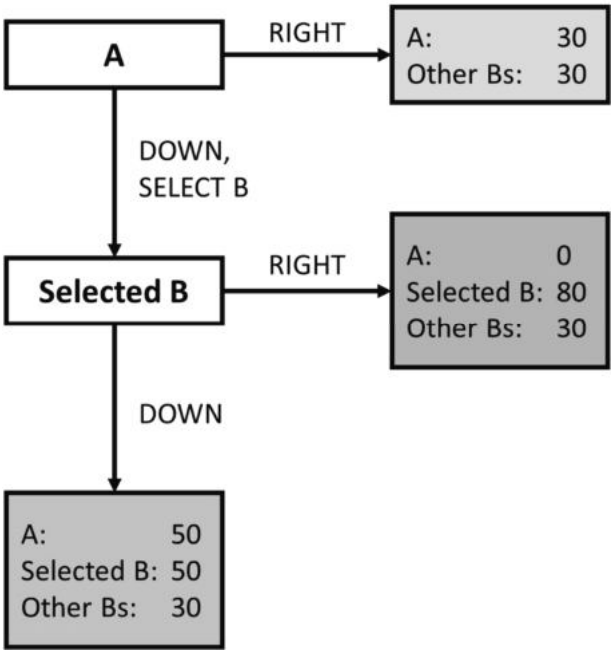

Figure A3.1: The game

The computer interface

We will now explain what your screen will look like during the game. What you will see on the right-hand side of your screen will be self-explanatory. On the left-hand side of your screen you see an example of a 'history window', which will always be visible during the game. Each of the four columns represents one B-participant and each row represents a round. At the top of the table you see the nationality and gender of each B-participant. The current round - round 5 in the example - is indicated by the bright green row color. In parentheses it is displayed which A is at play in which round. On the next pages we explain

the meaning of the colors. The background color of a cell in the table shows what choices were made.

Control condition only:

Each A gets informed about the results of his/her own interactions but not about the results of the interactions of any other A. Bs get informed about all choices.

Reputation condition only:

All As and Bs get informed about all choices of all participants.

In Figure A3.2 (control condition) / A3.3 (reputation condition) you see an example screen for an A (A1 in this case). The color of the background shows what choices were made. At the bottom of the table you can find a legend with an explanation of the colors.

Control condition only:

A1 sees question marks ("?") in the rounds in which one of the other As is at play. The question marks of past rounds are not highlighted in color because A1 never receives information about the choices in these rounds.

On the left you see an example screen for a B. Again, the colors indicate what decisions were made in previous rounds. In the example it's player A1's turn.

Control condition only:

The plus signs in past rounds indicate to a B which past rounds A has information about. Participant A1 only sees a plus sign ("+") in the rows in which he or she is at play (round 1 and 5). All participants in role B see plus signs ("+") and question marks ("?") in the same rounds as the A who is at play (A1 in the example). The plus signs in future rounds show when it's A's turn again. Note, the plus signs and question marks are placed in different rounds if it is a different A's turn.

Your Choice: RIGHT or DOWN

Time left to complete this page: **1:33**

	B1	B2	B3	B4
Nationality	US	India	US	India
1 (A1)	?	?	?	?
2 (A2)	?	?	?	?
3 (A3)	?	?	?	?
4 (A4)	+	+	+	+
5 (A1)	?	?	?	?
6 (A6)	?	?	?	?
7 (A1)	?	?	?	?
8 (A4)	+	+	+	+

Legend:

+ Result visible to you

? Result not visible to you

This B was not selected

This B was selected and chose RIGHT

This B was selected and chose DOWN

If in a round the active A chooses RIGHT, the colors of all Bs are grey.

You are Participant A4. It is your turn.
Make your choice - RIGHT or DOWN.

If you choose DOWN, also select a B participant.
Then click next.

- RIGHT
- DOWN - B1
- DOWN - B2
- DOWN - B3
- DOWN - B4

Next

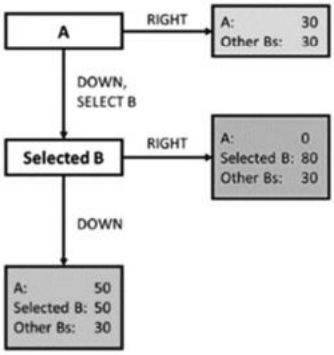

Figure A3.2: Decision screen for an A-player in the control condition.

A

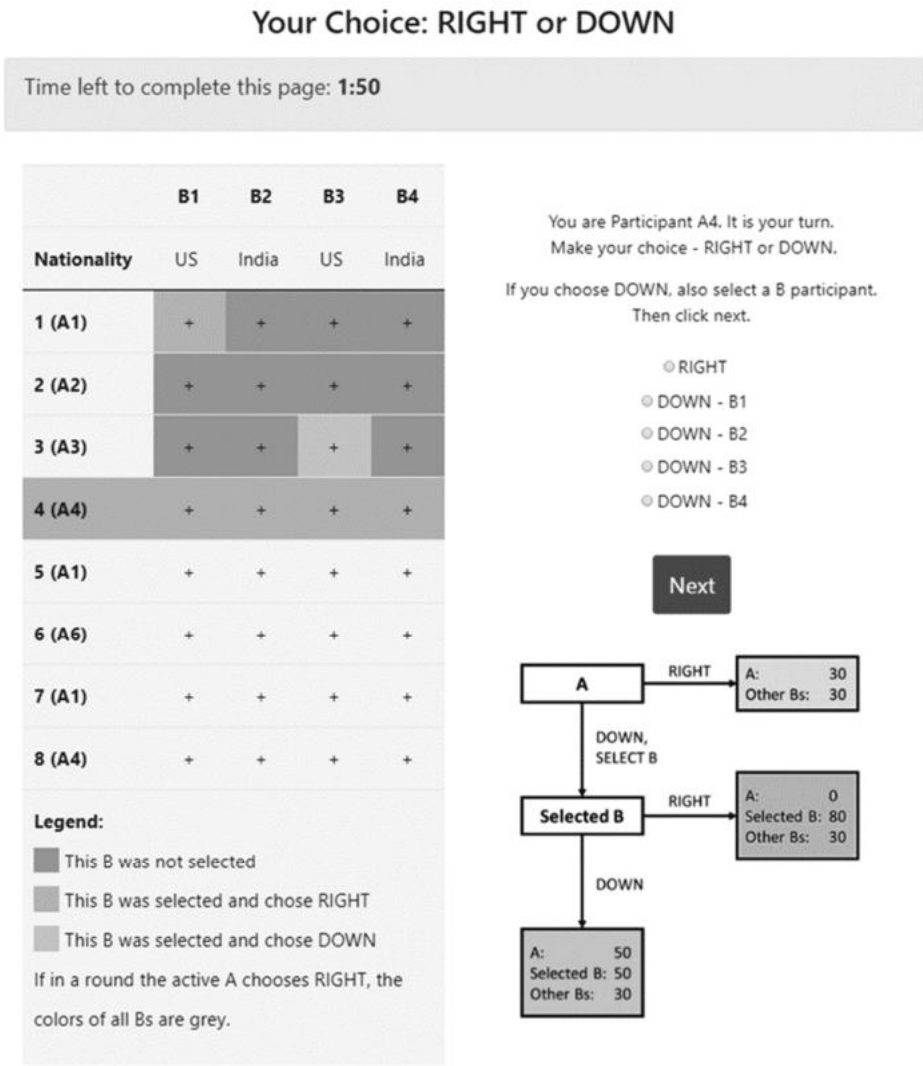

Figure A3.3: Decision screen for an A-player in the reputation condition**.**

The duration of the game

You will participate in 2 games, each lasting 8 rounds, one game played after the other. For each game you get randomly assigned to a new group of 8 participants and to your role. It is possible that you are matched with the same other participant in more than one game. However, should this happen, neither you nor the other participant will be able to notice this.

Interactive game

We would like to remind you that this is an interactive game: you are playing the game with 7 real other people. This means that you may have to wait sometimes while other people are making decisions. It also means that other people have to wait while you are making a decision. It is therefore important that you complete this study without interruptions. Please make your decision within the time limit shown on your screen. If you do not make your decision between RIGHT or DOWN within the time limit, the game will end and you will receive no payment and no bonus. The other participants will in that case receive the points that they earned so far, plus the points they would have earned in the remaining rounds (based on their average number of points earned per round). If this happens in the first round of the game (before you had the chance to earn any points) you will receive 30 points per round.

During the study, please do not close this window or leave the study's web pages in any other way. If you do close your browser or leave the task, you will not be able to re-enter and we will not be able to pay you!

To help you make a decision in time, we can send you a desktop notification when it is your turn to make a decision. If you wish to receive these notifications, please click the button below. After that click "Next".

Table A3.1: Results conditional logit models with the choice of the trustor as dependent variable. Logged odds

	American trustor				Indian trustor			
	Control	Reputation			Control	Reputation		
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Main effects								
Trustee India	-0.77***	-0.82***	-0.87**	-0.84***	0.15	0.26	0.25	0.11
	(0.22)	(0.22)	(0.30)	(0.23)	(0.21)	(0.24)	(0.29)	(0.25)
Number of times trustee honored trust	-	0.81***	0.76**	0.81***	-	0.30*	0.29	0.28
		(0.16)	(0.25)	(0.17)		(0.14)	(0.21)	(0.15)
Number of times trustee abused trust	-	-2.70***	-2.72***	-2.83**	-	0.42	0.42	-0.30
		(0.71)	(0.72)	(0.89)		(0.22)	(0.22)	(0.54)
Trustee India * Number of times trustee honored trust	-	-	0.10	-	-	-	0.02	-
			(0.32)				(0.30)	
Trustee India * Number of times trustee abused trust	-	-	-	0.47	-	-	-	0.97
				(1.32)				(0.54)
Second trustee								
Game number	0.09	1.28**	1.27**	1.28**	-0.76	-1.36*	-1.37*	-1.32*
	(0.51)	(0.49)	(0.49)	(0.49)	(0.65)	(0.54)	(0.54)	(0.55)
Round	0.24	0.07	0.06	0.07	-0.05	0.04	0.04	0.03
	(0.13)	(0.14)	(0.15)	(0.14)	(0.13)	(0.11)	(0.12)	(0.12)
Majority trustors India	0.84	-1.26*	-1.25*	-1.26*	-0.47	-0.40	-0.39	-0.32
	(0.65)	(0.61)	(0.62)	(0.61)	(0.63)	(0.61)	(0.63)	(0.61)
Constant	-1.50	-2.25*	-2.21*	-2.23*	1.92	2.35*	2.35*	2.19*
	(0.98)	(0.99)	(0.99)	(0.98)	(1.26)	(1.13)	(1.13)	(1.10)

(Continued on next page)

Table A3.1: continued

	American trustor				Indian trustor			
	Control	Reputation			Control	Reputation		
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Third trustee								
Game number	0.20	0.32	0.33	0.32	-0.84	-0.92	-0.92	-0.83
	(0.46)	(0.66)	(0.67)	(0.66)	(0.61)	(0.49)	(0.49)	(0.49)
Round	0.26*	0.11	0.11	0.11	-0.01	0.15	0.15	0.11
	(0.11)	(0.17)	(0.17)	(0.17)	(0.11)	(0.11)	(0.11)	(0.11)
Majority trustors India	0.46	-1.08	-1.08	-1.08	-0.67	-0.19	-0.19	-0.17
	(0.76)	(0.84)	(0.83)	(0.84)	(0.70)	(0.63)	(0.62)	(0.61)
Constant	-1.42	-1.79	-1.79	-1.81	2.21	1.08	1.08	0.95
	(1.11)	(1.34)	(1.33)	(1.34)	(1.34)	(1.11)	(1.11)	(1.06)
Fourth trustee								
Game number	0.43	1.01	0.99*	1.02*	-1.1	-0.01	-0.01	0.14
	(0.57)	(0.51)	(0.50)	(0.51)	(0.79)	(0.62)	(0.62)	(0.63)
Round	0.22	0.25	0.24	0.25	0.26*	0.23	0.23	0.22
	(0.12)	(0.14)	(0.15)	(0.14)	(0.12)	(0.12)	(0.12)	(0.12)
Majority trustors India	1.37	0.25	0.26	0.24	-0.67	-1.02	-1.02	-0.87
	(0.81)	(0.64)	(0.65)	(0.64)	(0.73)	(0.62)	(0.63)	(0.63)
Constant	-2.22*	-3.08**	-3.02**	-3.11**	0.66	-0.37	-0.37	-0.76
	(1.09)	(1.13)	(1.11)	(1.14)	(1.59)	(1.33)	(1.32)	(1.40)
Number of decisions	480	548	548	548	440	524	524	524
Number of trustors	120	137	137	137	110	131	131	131

Note: Standard error in parentheses. * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Appendix chapter 4

Table A4.1: Regression results study 1 including control variables

	Logged price		Sold		Number of bidders	
	(1)	(2)	(3)	(4)	(5)	(6)
Main effects						
Log positive ratings	-0.01***	-0.01***	0.33***	0.36***	0.03***	0.04***
	(-0.02, -0.01)	(-0.02, -0.01)	(0.31, 0.35)	(0.33, 0.38)	(0.02, 0.03)	(0.03, 0.04)
Log negative ratings	-0.02**	-0.04***	0.06*	0.05	-0.02	-0.02*
	(-0.04, -0.01)	(-0.06, -0.03)	(0.00, 0.12)	(-0.02, 0.13)	(-0.03, 0.00)	(-0.04, -0.00)
Denied	-0.00	-0.02**	0.19***	0.27***	0.02***	0.07***
	(-0.01, 0.00)	(-0.04, -0.01)	(0.15, 0.23)	(0.19, 0.35)	(0.02, 0.03)	(0.05, 0.09)
Acknowledged	-0.10***	-0.11***	0.26***	0.40***	0.02***	0.07***
	(-0.10, -0.09)	(-0.13, -0.10)	(0.23, 0.29)	(0.33, 0.47)	(0.01, 0.02)	(0.05, 0.08)
Interaction effects						
Positive * denied	-	0.00	-	-0.03*	-	-0.01***
		(-0.00, 0.01)		(-0.05, -0.00)		(-0.02, -0.01)
Negative * denied	-	0.01	-	0.01	-	0.00
		(-0.01, 0.02)		(-0.05, 0.08)		(-0.01, 0.02)
Positive * acknowledged	-	0.00	-	-0.04***	-	-0.02***
		(-0.00, 0.00)		(-0.06, -0.02)		(-0.02, -0.01)
Negative * acknowledged	-	0.03***	-	0.01	-	0.01*
		(0.02, 0.03)		(-0.05, 0.06)		(0.00, 0.02)

(Continued on next page)

Table A4.1: Continued

	Price		Sold		Number of bidders	
	(1)	(2)	(3)	(4)	(5)	(6)
Control variables						
Mileage (log)	-0.11***	-0.11***	0.16***	0.16***	0.01*	0.01*
	(-0.11, -0.11)	(-0.11, -0.11)	(0.14, 0.18)	(0.14, 0.18)	(0.00, 0.01)	(0.00, 0.01)
Tekst size (log)	0.14***	0.14***	0.08***	0.07***	0.06***	0.06***
	(0.13, 0.14)	(0.13, 0.14)	(0.04, 0.11)	(0.04, 0.10)	(0.06, 0.07)	(0.05, 0.07)
Featured listing	0.14***	0.14***	0.56***	0.57***	0.12***	0.12***
	(0.12, 0.17)	(0.12, 0.17)	(0.44, 0.69)	(0.44, 0.69)	(0.09, 0.15)	(0.09, 0.15)
Phone number	0.09***	0.09***	-0.06	-0.06	0.01	0.01
	(0.07, 0.10)	(0.07, 0.10)	(-0.13, 0.00)	(-0.13, 0.00)	(-0.01, 0.03)	(-0.01, 0.03)
Address	-0.06***	-0.06***	-0.13**	-0.13**	-0.01	-0.01
	(-0.07, -0.04)	(-0.07, -0.04)	(-0.22, -0.04)	(-0.21, -0.04)	(-0.03, 0.01)	(-0.03, 0.01)
Powerseller	-0.01	-0.01	-0.31***	-0.29***	-0.09***	-0.08***
	(-0.04, 0.02)	(-0.04, 0.02)	(-0.45, -0.18)	(-0.43, -0.16)	(-0.13, -0.06)	(-0.12, -0.05)
One or two options	0.10***	0.11***	-0.08*	-0.08*	0.00	0.00
	(0.09, 0.12)	(0.09, 0.12)	(-0.16, -0.00)	(-0.16, -0.01)	(-0.01, 0.02)	(-0.01, 0.02)
More options	0.18***	0.18***	-0.05	-0.05	0.02*	0.02*
	(0.17, 0.20)	(0.17, 0.20)	(-0.12, 0.02)	(-0.12, 0.02)	(0.00, 0.04)	(0.00, 0.04)
Duration (days)	0.04***	0.04***	-0.11***	-0.11***	0.01***	0.01***
	(0.03, 0.04)	(0.03, 0.04)	(-0.13, -0.09)	(-0.13, -0.09)	(0.01, 0.02)	(0.01, 0.02)
Constant	9.70***	9.70***	-1.24	-1.32	1.98***	1.95***
	(9.41, 9.98)	(9.41, 9.98)	(-2.69, 0.20)	(-2.77, 0.13)	(1.63, 2.33)	(1.60, 2.30)

(Continued on next page)

Table A4.1: Continued

	Price		Sold		Number of bidders	
	(1)	(2)	(3)	(4)	(5)	(6)
FE model/year/week	Yes	Yes	Yes	Yes	Yes	Yes
Random intercept seller	Yes	Yes	Yes	Yes	Yes	Yes
Observations	56,809	56,809	65,113	65,113	65,113	65,113
Sellers	37,285	37,285	41,903	41,903	41,903	41,903

Note: 95% Confidence Interval in parentheses. *** p<0.001, ** p<0.01, * p<0.05. Models 1 and 2: Multilevel linear regression. Models 3 and 4: Multilevel logistic regression. Model 5 and 6: Multilevel negative binomial regression.

Consent form (study 2)

The goal of this experiment is to better understand how people make decisions in groups. This experiment is conducted by researchers from Utrecht University, the Netherlands.

If you agree to be in this study, you will play 36 rounds of a game in which you will be assigned to a group with five to nine other individuals. The decisions that you and the other participants make in the games will determine how much you earn. If you complete the experiment, you will receive a payment of £2.50 plus a bonus that varies between £2.40 and £8.40. It may happen that you finish reading the instructions, but after that cannot be assigned to a group to play the game. In that case you will receive a payment of £2.50.

There are no known risks associated with your participation in this research beyond those of everyday life. Taking part in this study is voluntary. Confidentiality of your research records will be strictly maintained. We collect your Prolific Participant ID to calculate your bonus payment. After the study has been completed, we will separate your Prolific ID from the experimental data so answers cannot be traced back to you. The data from the study will be kept at least 10 years.

If there is anything about the study or taking part in it that is unclear or that you do not understand, you may contact the researcher (Judith Kas, j.kas@uu.nl, Utrecht University, the Netherlands, department of Sociology).

I have read and understand the explanation provided to me and I voluntarily agree to participate in this study: Yes / No

Instructions for participants (study 2)

Welcome and thank you for coming!

Please carefully read all of the information that follows. In this experiment, you will earn points by making decisions in "games". How much you earn depends on your decisions, the decisions of others and on chance. The experiment will last about 60 minutes. At the end of the experiment, you will be paid £2.50, plus a bonus of £1 for every 15 points you earned in the game.

Please note that this is an interactive experiment. That means that you may have to wait while other participants are making decisions. Please do not use your phone or any function of the computer that is not necessary for carrying out the experiment. We will now explain the rules of the game.

These instructions are the same for all participants. We understand that sometimes studies claim that participants will actually interact in real time, but in fact use simulated other people ("bots"). As a policy, we do not do this. No other participant will be able to link your decisions to your identity or get to know your name or earnings. After you have read the instructions, there will be a quiz in which you can make sure that you understand everything correctly. Then, we turn to the next stage during which you can score points by playing several rounds of a game. At the end of the session you will be asked to answer a short survey.

Description of the game

The game proceeds in rounds and is played in groups of 4, 6 or 8 participants (depending on the total number of participants in the session). You will play in the same group throughout the experiment. The game proceeds in rounds. At the beginning of each round, half of the participants in your group is assigned to role A and half of the participants is assigned to role B. Then, you are randomly matched with another participant in the other role (if you are assigned to role A, you are matched to a participant with role B and vice versa). In case you are by chance matched multiple times with the same other participant, neither you nor the other participant knows this.

There are two games: Game 1 and Game 2. Figure A4.1 explains schematically how the games proceed. In both games, player A starts and chooses between

RIGHT and DOWN. If player A chooses RIGHT, player A and player B both receive 2 points. If player A chooses DOWN in Game 1, player B has to make a decision. If player B also chooses DOWN, both player receive 3 points. If player B chooses RIGHT, player A receives 0 points and player B receives 4 points. If player A chooses down in Game 2, player B must choose RIGHT.

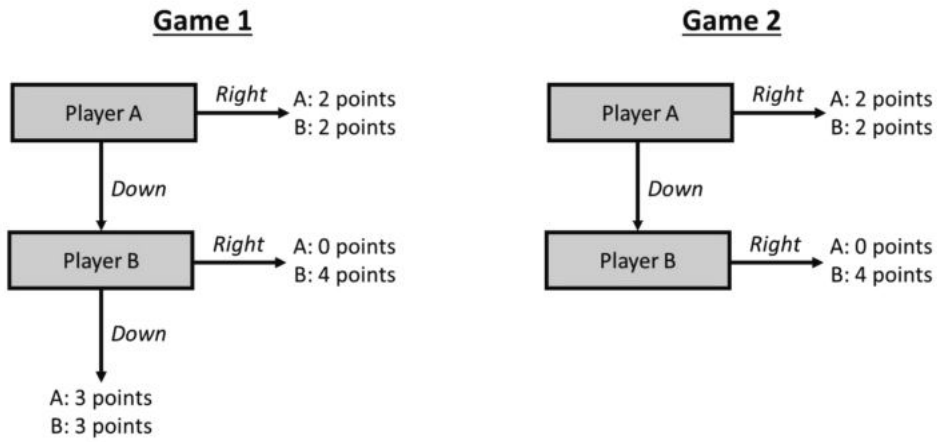

Figure A4.1: Game tree

At the beginning of each round the computer randomly determines which game you and the other participant are playing. You will be playing game 1 with probability 0.5 and game 2 also with probability 0.5. Only player B will know what game you are player. Player A will not be informed about this.

Communication conditions only: Before player A makes a decision, player B can choose to send a message to player A. In that message players B can tell player A whether they will choose DOWN or RIGHT if player A chooses DOWN. This information does not need to be truthful. Also when players B know that they are playing Game 2, and that they do not have the option to choose DOWN, they can promise to choose DOWN.

The game lasts for 36 rounds. Your role may change every round and the order in which you are playing as player A and as player B is determined randomly. At the beginning of each round you will see a screen with information about your role. Players B will also be informed about the game you have been assigned to.

Communication conditions only:

At this screen player B can also send a message to player A.

The computer interface

We will now explain what your screen will look like during the game. The right-hand side of the screen will be self-explanatory. On the top left of every screen you see an overview of your decisions and results of previous rounds. The information about the most recent round is in the top row of the table. You can use the light-blue scroll bar to scroll back to earlier rounds.

Reputation conditions only:

At the bottom left of every screen, player A will see information about the player B they are paired with in the current round. The table contains information about

Communication condition only:

the messages

and decisions of player B in role B. The table does not contain information about this player B's decisions in role A.

Below you find an explanation of the information in the columns:

- Round — Round number
- Role — Your role in that round
- Decision A — Decision of player A
- Decision B — Decision of player B
- Payoff — Total number of points you collected in that round.
  Communication condition only:
- Message — Message player B sent to player A

Quiz

You have finished reading the instructions. Feel free to have another look at the parts you found difficult to understand by using the 'back'-button. Before you begin the study we want to ensure that you understood the instructions.

On the following pages you find a number a questions that you need to answer correctly
before proceeding to the study. You cannot score (or lose) points by answering these questions. The goal is only to make sure that you understand everything correctly.

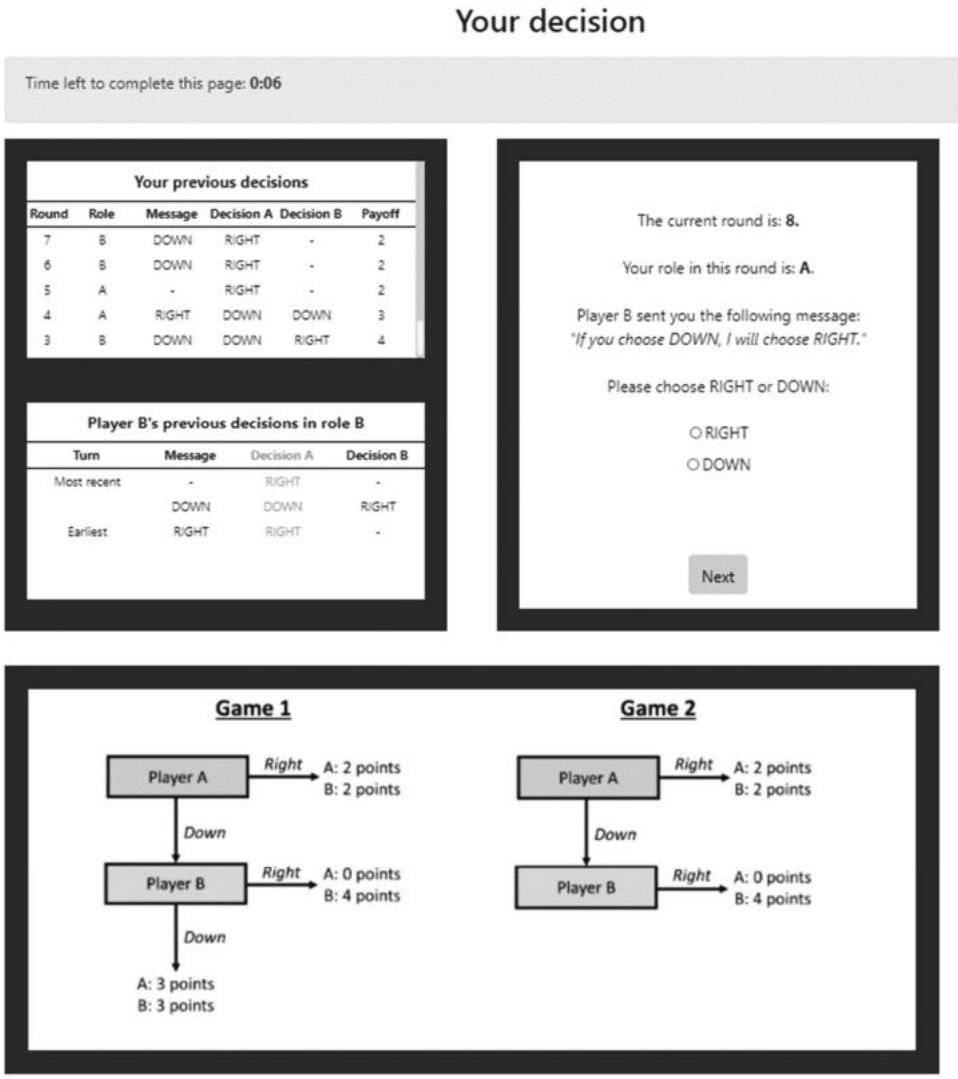

Figure A4.2: Decision screen of player A in the Combined condition (study 2)

Table A4.2: Logistic regression results per condition (study 2)

	Model 1 Control	Model 2 Comm.	Model 3 Comm.	Model 3 Reputation	Model 4 Combined	Model 4 Combined
Promises made by player B						
Promised to honor	-	16.77***	16.33***	-	6.96***	6.98***
		(8.72, 32.26)	(8.46, 31.52)		(4.50, 10.76)	(4.49, 10.85)
Promised to abuse	-	1.09	0.92	-	0.86	0.84
		(0.50, 2.36)	(0.42, 2.04)		(0.48, 1.55)	(0.46, 1.55)
Reputation player B						
# honored	0.97	1.05		1.19	1.15*	-
	(0.81, 1.17)	(0.91, 1.22)		(0.97, 1.45)	(1.01, 1.31)	
# abused	1.11	0.93		0.96	0.89	-
	(0.96, 1.28)	(0.76, 1.13)		(0.82, 1.13)	(0.78, 1.00)	
# honest	-	-	1.09	-	-	1.06
			(0.99, 1.21)			(0.97, 1.15)
# dishonest	-	-	0.88	-	-	0.88*
			(0.70, 1.10)			(0.78, 1.00)
Control variables						
Round	0.96***	0.96**	0.95**	0.96***	0.94***	0.95***
	(0.94, 0.98)	(0.93, 0.99)	(0.92, 0.98)	(0.93, 0.98)	(0.92, 0.97)	(0.92, 0.97)
Group size	1.13	1.29	1.30	1.03	0.92	0.92
	(0.82, 1.55)	(0.96, 1.74)	(0.97, 1.75)	(0.79, 1.34)	(0.76, 1.11)	(0.76, 1.11)

(Continued on next page)

Table A4.2: Continued

	Model 1 Control	Model 2 Comm.	Model 3 Comm.	Model 3 Reputation	Model 4 Combined	Model 4 Combined
Constants						
Constant	0.13	0.02**	0.02**	0.28	0.68	0.70
	(0.01, 1.44)	(0.00, 0.25)	(0.00, 0.24)	(0.04, 2.11)	(0.16, 2.97)	(0.16, 3.03)
Variance group	1.00	1.89	1.70	1.26	1.00	1.00
level	(1.00, 1.00)	(0.67, 5.32)	(0.68, 4.24)	(0.69, 2.29)	(1.00, 1.00)	(1.00, 1.00)
Variance subject	8.88***	4.72**	4.76**	4.80**	4.18***	4.15***
level	(2.64, 29.90)	(1.86, 12.01)	(1.87, 12.13)	(1.77, 12.98)	(1.95, 8.95)	(1.95, 8.83)
Number of decisions	1,047	849	849	953	1,000	1,000
Number of subjects	62	58	58	57	61	61
Number of groups	9	9	9	8	10	10

Note: Odds ratios. 95% Confidence Interval in parentheses. *** p<0.001, ** p<0.01, * p< 0.05

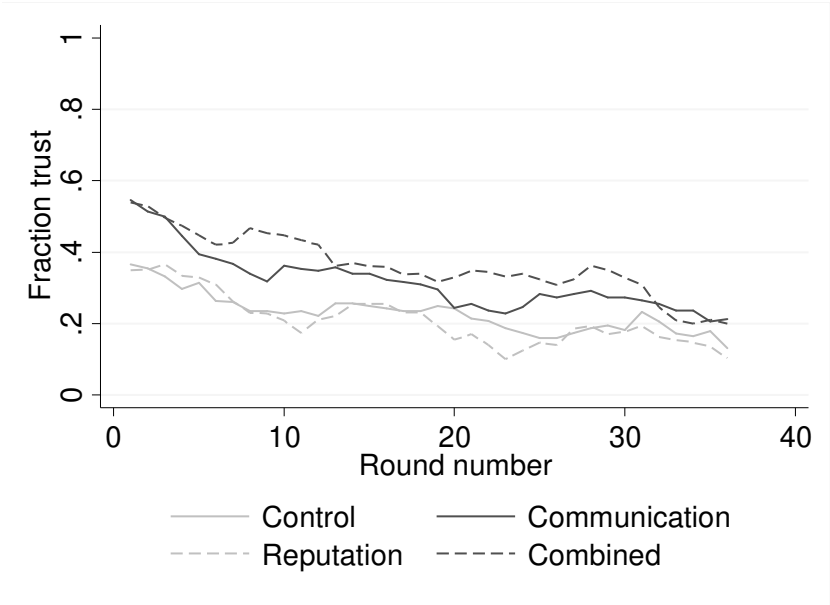

Figure A4.3: Change in trust over the rounds per condition (study 2)

A

Appendix chapter 5

Instructions for participants

Welcome and thank you for coming!

Please read the following instructions carefully. Please do not communicate with other participants. Turn off your phone and put it away. Also, you may not use any function of the computer that is not necessary for carrying out the experiment. Thank you very much for your participation. If at any point you have questions, raise your hand and we will assist you.

In this experiment you can earn money by scoring 'points' for certain decisions. These points will at the end of the experiment be exchanged for money. The number of points you score depends on your decisions and on the decisions of others. For every 100 points you score in the experiment, you will receive €0.60 (total amount rounded to €0.50).

These instructions are the same for all participants in the room. All of the payments are real and all participants are present in this room. All information that you read about in this study is truthful and real. No other participant will be able to link your decisions to your identity or get to know your name or earnings. After everybody has read the instructions, there will be a quiz in which you can make sure that you understand everything correctly. Then, we turn to the next stage during which you can score points by playing several rounds of a game. At the end of the session you will be asked to answer a short survey.

Description of the game

The game proceeds in rounds and is played in groups of 6, 8 or 10 participants (depending on the total number of participants in the session). At the beginning of each round, half of the participants in your group is assigned to role A and half of the participants is assigned to role B. Then, you are randomly matched with another participant in the other role (if you are assigned to role A, you are matched to a participant with role B and vice versa). In case you are by chance matched multiple times with the same other participant, neither you nor the other participant knows this.

Figure A5.1 / A5.2 shows the basic interaction situation. Player A starts and chooses between RIGHT and DOWN. If player A chooses RIGHT, player A and player B both receive 40 points. If player A chooses DOWN, player B has to make a decision. If player B also chooses DOWN,

Trust Game condition only:

Player A and player B both receive 50 points.

Lending Game condition only:

Player A receives 38a points and player B receives 50 points.

If player B chooses RIGHT, player A receives 0 points and player B receives 70 points.

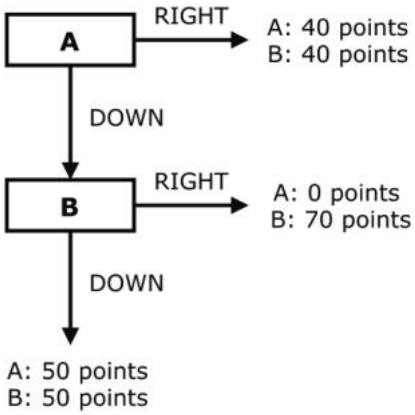

Figure A5.1: Trust Game

A

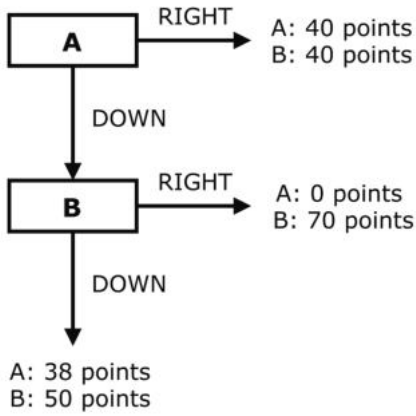

Figure A5.2: Lending Game

The duration of the game

The game lasts for 36 rounds. You will be assigned to role A 18 times and to role B also 18 times. However, your role may change every round and the order in which you are playing in role A and role B is determined randomly.

The computer interface

On top of the screen you can see the number of the current round, your role in this round and the number of times you have already played in each of the roles. At the bottom of the screen you can see the outcomes of previous rounds. Use the scroll bar to scroll back to earlier rounds.

> *Single-role reputation condition only:*
>
> At the beginning of every round, all players A are informed about the previous decisions in role B of the player B they are matched with.
>
> *Mixed-role reputation condition only:*
>
> At the beginning of every round, all players A are informed about the previous decisions in role A and B of the player B they are matched with.

This information is shown in the left panel.

In the right panel you can indicate your decision for RIGHT or DOWN. Make your decision by clicking with your mouse on the corresponding arrow. When you click on one of the arrows, the arrow will turn bold and blue and an 'OK'-button appears. Please click the 'OK'-button to proceed.

You will only see the decision screen of player B screen if the player in role A that you are matched with chooses DOWN. You can decide between RIGHT and DOWN by clicking with your mouse on the corresponding arrow. When you click on one of the arrows, the arrow will turn bold and blue and an 'OK'-button appears. Please click the 'OK'-button to proceed.

After every round, you can see the decisions and number of points scored by you and by the other participant in the current round. The other participant can also see your decision and the number of points scored by you in the current round. You also see the sum of your points in all rounds. You also see the sum of your points in all rounds. Please click the 'OK'-button to proceed to the next round.

Quiz

You have finished reading the instructions. Feel free to have another look at the parts you found difficult to understand. Please turn to the computer and answer a few questions that help you evaluate your understanding of the game. You cannot score (or lose) points by answering these questions. The goal is only to make sure that you understand everything correctly. If you have questions, please raise your hand. Otherwise, turn to the computer and click "OK.

Analyses decision trustee

Analytical strategy

We performed rank-sum tests and five multilevel logistic regressions to see whether there are differences in the level of trustworthiness between conditions. In the regressions we include the decision of the trustee (honor trust or not) as the dependent variable. In each analysis, we included random intercepts for subjects and groups. The results of these regressions are in Table A2. First, we did a regression pooling all conditions (model 1 in Table A2), with game type, reputation type, the interaction between the two and round number as independent variables. The next four models (models 2 through 5) regress the decision of the trustee on the round number and four variables indicating whether that trustee made a decision as a trustee and as a trustor ('Made trustee decision' and 'Made trustor decision) and the trustee's historical trustworthiness rate and trust rate ('Honor rate history' and 'Place rate history'.

Descriptive statistics and results rank-sum tests

In the Trust Game, significantly more trust is honored in the single-role reputation condition than in the mixed-role reputation condition (z = 2.722, p = .007), while in the Lending Game, the trustworthiness rate is higher in the mixed-role reputation condition than in the single-role reputation condition (z = -2.286, p = .022). There is no overall difference in trustworthiness rate between the Trust Game and the Lending Game (z = 1.360, p = .174).

A Wilcoxon rank-sum test shows that the success rate is significantly higher in the Trust Game than in the Lending Game (z=3.891, p<.001) and that there is no overall significant effect of reputation type (z = -0.327, p=0.744). In the Trust Game, there is no significant difference in the success rate between the single-role reputation condition and the mixed-role reputation condition (z = 1.363, p=0.173). In the Lending Game, the success rate is higher in the mixed-role reputation condition than in the single-role reputation condition (z =-2.575, p=0.010).

Individual trustee decision

In later rounds trustees become less and less likely to honor trust, although this effect is only significant in the Trust Game and not in the Lending Game. On

average, the probability that a trustee honors trust decreases with 0.07 per ten rounds. We find that a trustee's previous decisions in the trustee role ('Honor rate history') predict the trustee's honoring decision in both Trust Games and in the Lending Game in the single-role reputation condition. With every 0.1 increase in the trustee's historical trustworthiness rate, the probability that a trustee honors trust increases with 0.03 and 0.04 in the Trust Game with single-role reputation and mixed-role reputation respectively. In the Lending Game with single-role reputation, the probability that a trustee honors trust increases with 0.05 with every 0.1 increase in the trustee's historical trustworthiness rate. The trustee's previous decisions in the trustor role ('Place rate history') have a significant and positive effect on the probability that a trustee will honor trust in both games with mixed-role reputation. With every 0.1 increase in the trustee's historical trust rate, the probability that that trustee honors trust increases with 0.04 in both games with mixed-role reputation. In the games with single-role reputation the trustee's historical trust rate is not a significant predictor of the trustee's decision.

Overall trustworthiness rate

Model 1 in Table A5.1 shows that more trust is honored in the Trust Game than in the Lending Game ('LG') in the single-role reputation condition. On average, the probability that a trustee honors trust is 0.09 higher in the Trust Game than in the Lending Game with single-role reputation. By comparison, in the mixed-role reputation condition we find that the trustworthiness rate does not significantly differ between the Trust Game and the Lending Game ($b = .388$, $z = 0.95$, $p = .341$).

Confirming the results of the rank-sum test, we find that the interaction between game type and reputation system type ('LG * MRR') is significant: in the Trust Game the probability that trust is honored is 0.17 higher in the single-role reputation condition than in the mixed-role reputation condition, but in the Lending Game the probability that trust is honored is 0.14 higher in the mixed-role reputation condition ($b = 1.062$, $z = 1.97$, $p = 0.049$).

Interestingly, we found that not only trust levels were higher under mixed-role than single-role reputation conditions in the Lending Game, but levels at which trust was honored were also higher. One possible explanation for these additional beneficial effects of mixed-role reputation systems is that the

increased trust rate may have made it more worthwhile to invest in a reputation. The more frequently trust is placed, the larger the chance that a user with a good reputation receives trust, so the more it pays off to have a good reputation. These findings stand in stark contrast to what we found for paid service provision scenarios, where trust was honored *less* so under a mixed-role reputation regime than in a single-role reputation system.

Table A5.1: Determinants of trustworthiness

Independent variables	Model 1: all conditions	Model 2: TG, SRR	Model 3: TG, MRR	Model 4: LG, SRR	Model 5: LG, MRR
Treatment variables					
LG	-1.661**	-	-	-	-
	(0.481)				
MRR	-1.060*	-	-	-	-
	(0.429)				
LG * MRR	2.048**	-	-	-	-
	(0.666)				
Reputation variables					
Made trustee decision	-	0.049	-1.058*	-1.668	1.625*
		(0.426)	(0.459)	(0.854)	(0.783)
Honor rate history	-	1.766***	1.998***	3.243*	0.648
		(0.392)	(0.574)	(1.131)	(1.007)
Made trustor decision	-	-1.400*	-0.180***	-0.173	-3.890***
		(0.598)	(0.578)	(1.499)	(0.904)
Place rate history	-	0.299	1.924**	3.822	2.639*
		(0.430)	(0.547)	(3.269)	(1.213)

(Continued on next page)

Table A5.1: Continued

Independent variables	Model 1: all conditions	Model 2: TG, SRR	Model 3: TG, MRR	Model 4: LG, SRR	Model 5: LG, MRR
Round number	-0.048***	-0.031**	-0.050***	-0.044	-0.004
	(0.009)	(0.012)	(0.014)	(-0.029)	(0.026)
Constant	1.981	1.825	0.464	0.099	2.500
	(0.328)	(0.573)	(0.530)	(1.431)	(0.967)
Observations	1197	468	404	109	216

Note: Standard errors in parentheses. *** $p<0.001$, ** $p<0.01$, * $p<0.05$. TG refers to Trust Game, LG refers to Lending Game. SRR refers to single-role reputation system and MRR to mixed-role reputation system. Made trustee decision (Made trustor decision) indicates whether the trustee has made a decision in the trustee (trustor) role yet. Honor rate history (Place rate history) indicates the rate at which the trustee has honored (placed) trust in the past

NEDERLANDSE SAMENVATTING

(Dutch summary)

Vertrouwen en reputatie in de peer-to-peer platformeconomie

Introductie

In het voorjaar van 2019 was ik een paar maanden in de VS als onderdeel van mijn PhD. Ik maakte gebruik van het peer-to-peer platform Couchsurfing om anderen te ontmoeten die ook de nationale parken rondom San Francisco wilden bezoeken in de weekenden. Op het platform vond ik een advertentie van iemand die per motor naar de Grand Canyon wilde rijden om daar drie dagen te kamperen. Hij had plek voor een extra passagier en hij had ook nog een plaats vrij in zijn tent voor iemand die wel met hem mee wilde.

Dit voorbeeld illustreert het belang van vertrouwen in de peer-to-peer platformeconomie. 'Platformeconomie' is een verzamelnaam voor allerlei activiteiten. Platformen zijn websites waar aanbieders en vragers van diensten en goederen bij elkaar gebracht worden (Leong & Belzer, 2017). Het verschil met traditionele webbedrijven zit 'm in het 'peer-to-peer' aspect: niet (alleen) bedrijven, maar (ook) particulieren op platformen hun producten aan kunnen bieden. Via Airbnb kan iedereen die ruimte in huis over heeft een kamer verhuren aan vakantiegangers en via SnappCar kan iedereen zijn eigen auto verhuren. Vaak kennen de aanbieders en vragers op dit soort platformen elkaar niet en transacties zijn vaak eenmalig. Daardoor lopen zowel vragers als aanbieders een risico als zij deelnemen aan dit soort transacties (Kuwabara, 2015; Parigi, Santana, & Cook, 2017). Daarnaast is er op bedrijven strenger toezicht en gelden er strengere regels dan voor gebruikers op dit soort platformen (Katz, 2015; ter Huurne, Ronteltap, Corten, & Buskens, 2017)

De meeste gebruikers van dit soort platformen zullen goede intenties hebben. Toch is er altijd een kans dat je iemand ontmoet die je kwaad doet. In het voorbeeld van de reis naar de Grand Canyon zou ik me zorgen maken over mijn persoonlijke veiligheid. Ik wist namelijk niet hoe goed mijn onbekende reisgenoot is in het besturen van de motor. Daarnaast zou de overnachting in een kleine tent me behoorlijk kwetsbaar maken. Hetzelfde gold natuurlijk voor hem: aangezien we elkaar niet kenden, wist hij ook niets over mijn intenties, dus ik zou ook een bedreiging voor hem kunnen betekenen. Mensen die dure spullen (zoals auto's, huizen en motoren) verhuren via dit soort platformen lopen het risico dat de huurder de spullen niet terugbrengt, of dat de spullen beschadigd worden. Betrouwbare gebruikers doen wat ze beloven, terwijl onbetrouwbare gebruikers het product beschadigen of niet terugbrengen, of

zelfs een fysieke bedreiging vormen voor anderen. Om te voorkomen dat zij opgelicht worden, zullen mensen die iets willen (ver)kopen, (uit)lenen of (ver)huren onbetrouwbare gebruikers vermijden. Het probleem is dat gebruikers van dit soort platformen geen onderscheid kunnen maken tussen betrouwbare en onbetrouwbare gebruikers (Macy & Skvoretz, 1998).

Het vertrouwensprobleem in de platformeconomie is niet uniek. In veel alledaagse beslissingen speelt vertrouwen een rol en er is veel onderzoek gedaan naar dit onderwerp (Akerlof, 1970; Diekmann, Jann, Przepiorka, & Wehrli, 2014; Kreps, 1996). In dit proefschrift bestudeer ik het klassieke vertrouwensprobleem aan de hand van het voorbeeld van de platformeconomie. Dit onderzoek vergroot onze kennis van vertrouwensproblemen in het algemeen en biedt inzicht in specifieke oplossingen voor het vertrouwensprobleem in de platformeconomie.

Platformen helpen hun gebruikers onderscheid te maken tussen betrouwbare en onbetrouwbare gebruikers. De meeste online platformen maken gebruik van reputatiesystemen om dit onderscheid duidelijk te maken. Reputatiesystemen verzamelen en verspreiden feedback over gebruikers op het platform (Resnick, Kuwabara, Zeckhauser, & Friedman, 2000). Volgens klassieke theorieën over reputatiesystemen hebben mensen meer vertrouwen in personen die in het verleden betrouwbaarder zijn geweest (Buskens, 2003; Buskens & Raub, 2002; Cook, Hardin, & Levi, 2005; Weigelt & Camerer, 1988). Recensies op Couchsurfing bevatten bijvoorbeeld zinnen als 'Deze host is betrouwbaar'. Deze recensies zijn geschreven door andere Couchsurfers die in het verleden met de host gereisd hebben. Uit eerder onderzoek weten we dat platformgebruikers met een betere online reputatie inderdaad meer vertrouwd worden (Boero, Bravo, Castellani, & Squazzoni, 2009; Bolton, Katok, & Ockenfels, 2004; Charness, Du, & Yang, 2011; Duffy, Xie, & Lee, 2013; Fehrler & Przepiorka, 2013). Zij kunnen meer verkopen of verhuren voor hogere prijzen op dit soort platformen en zij krijgen meer kansen om spullen te lenen of te huren.

In dit proefschrift bestudeer ik drie nieuwe ideeën over de manier waarop reputatie vertrouwen beïnvloedt. Deze ideeën zijn gebaseerd op drie kenmerken die de platform economie typeren, maar die ook van toepassing zijn op andere contexten.

Ten eerste onderzoek ik in hoofdstuk 2 en 3 of reputatiesystemen kunnen helpen om discriminatie van (etnische) minderheden te verminderen. Aan de ene kant kan reputatie-informatie ervoor zorgen dat etniciteit een minder grote rol speelt bij de beslissing om iemand te vertrouwen. Aan de andere kant kan het zo zijn dat het niet voor iedereen even makkelijk is om een reputatie op te bouwen. Ten tweede onderzoek ik in hoofdstuk 4 of reputatiesystemen ervoor kunnen zorgen dat gebruikers van veilingplatforms vaker de waarheid spreken. Ik onderzoek of aan de hand van reputatie-informatie onderscheiden kan worden wie meestal de waarheid spreekt en wie er soms liegt. Ten derde onderzoek ik in hoofdstuk 5 of reputatiesystemen anders functioneren als de rollen van aanbieder en verkoper niet strikt van elkaar gescheiden zijn. In die situatie kan het reputatiesysteem meerdere soorten informatie bevatten. Ik heb onderzocht in hoeverre die extra informatie invloed heeft op vertrouwen.

Ik geef nu eerst een kort overzicht van de bestaande literatuur op het gebied van vertrouwen en reputatiesystemen. Daarna vat ik de vier hoofdstukken samen. Ten slotte trek ik een aantal algemene conclusies.

Overzicht van de bestaande literatuur

Vertrouwen wordt vaak gedefinieerd als 'de bereidheid om zich kwetsbaar op te stellen ten aanzien van anderen, gebaseerd op verwachtingen over de intenties en vaardigheden van die anderen' (Mayer, Davis, & Schoorman, 1995; Rousseau, Sitkin, Burt, & Camerer, 1998). Platformgebruikers lopen een risico, omdat de regels die van toepassing zijn op individuen minder duidelijk zijn dan de regels voor bedrijven (Katz, 2015; ter Huurne, Ronteltap, Corten, & Buskens, 2017) en omdat particulieren over het algemeen minder transacties aangaan dan bedrijven. Daarnaast zijn er veel verschillen in de kwaliteit van de aangeboden producten op dit soort platformen. Consumenten zijn beschermd tegen grote verliezen door verzekeringen en veilige betaalsystemen. Toch kan de kwaliteit van hun aankopen tegenvallen. Hetzelfde geldt voor aanbieders. In veel gevallen zijn zij beschermd tegen diefstal en zichtbare schade aan hun spullen. Slijtage en andere schade die niet meteen zichtbaar is worden echter niet gedekt door de verzekering. Daarnaast worden veel van dit soort online transacties gevolgd door fysieke ontmoetingen tussen aanbieders en vragers.

Die ontmoetingen kunnen een gevaar voor de veiligheid betekenen (Ranchordas, 2015). Het probleem is niet dat alle gebruikers van platformen egoïstisch en gevaarlijk zijn. De meeste gebruikers zijn waarschijnlijk te vertrouwen. Het probleem is dat een eenmalige ontmoeting met een onbetrouwbare gebruiker kosten met zich mee kan brengen die vele malen hoger zijn dan de voordelen van een succesvolle interactie.

We noemen gebruikers die een risico lopen en die dus vertrouwen in anderen moeten plaatsen 'trustors'. Gebruikers die vertrouwd worden door anderen noemen we 'trustees'. Figuur 1 geeft schematisch een interactie tussen een trustor en een trustee weer. Dit onderscheid is een versimpeling van de werkelijkheid: in veel platforminteracties lopen beide partijen een risico. Zowel verhuurders als huurders op Airbnb lopen het risico dat de andere partij niet doet wat zij beloofd heeft. We beschouwen dit soort interacties als een serie van vertrouwensproblemen waarin soms de verhuurder en soms de huurder een risico loopt.

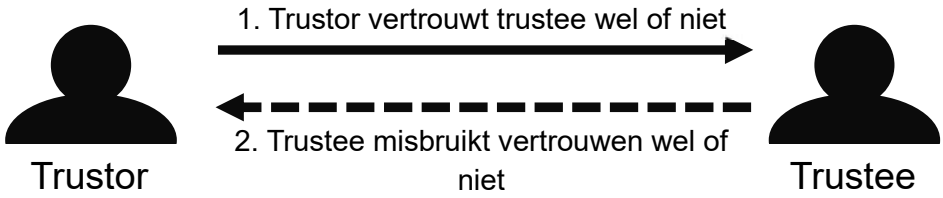

Figuur 1: Schematische weergave van het vertrouwensprobleem in de peer-to-peer platformeconomie

Het is lastig om te meten hoe groot de platformeconomie is en hoeveel impact deze nieuwe economie heeft. Volgens schattingen is het aantal platformen in het laatste decennium enorm gestegen en er gaan inmiddels miljoenen euro's om in dit soort platformen (McKinsey, 2016; Roof, 2019). Dit suggereert dat dit soort platformen een manier hebben gevonden om het vertrouwensprobleem op te lossen. Reputatiesystemen worden vaak genoemd als belangrijkste oplossing voor dit probleem.

Het idee is dat deze systemen informatie verschaffen over de mate waarin trustees zich in het verleden betrouwbaar hebben gedragen. Trustors kunnen

deze informatie gebruiken om een onderscheid te maken tussen betrouwbare en onbetrouwbare trustees. Daarnaast geven reputatiesystemen trustees een reden om zich netjes te gedragen, zelfs als ze dat zelf eigenlijk niet willen. Om te voorkomen dat zij een slechte beoordeling krijgen en in de toekomst niet meer deel kunnen nemen aan transacties, zullen zij zich betrouwbaar op moeten stellen (Buskens, 2003; Cook et al., 2005; Weigelt & Camerer, 1988).

Mijn bijdrage aan de theorie

Klassieke theorieën over reputatiesystemen richten zich op het effect van de mate waarin een persoon in het verleden betrouwbaar is geweest op vertrouwen. De opkomst van de platformeconomie heeft een aantal kenmerken van reputatiesystemen aan het licht gebracht waar tot nu toe geen aandacht aan geschonken is. Deze kenmerken hebben effect op de manier waarop reputatiesystemen vertrouwen beïnvloeden.

Hoofdstuk 2 en 3: Reputatiesystemen en discriminatie

Ten eerste, toen ik besloot of ik mee wilde gaan kamperen in de Grand Canyon wist ik niet alleen welke beoordelingen de organisator in het verleden had gekregen. Hij had ook een aantal foto's in zijn profiel staan en ik wist zijn voornaam. Deze nadruk op persoonlijke informatie is kenmerkend voor online platforms en is bedoeld om vertrouwen tussen verschillende gebruikers van het platform te kweken. Uit eerder onderzoek blijkt echter dat op basis van informatie die wordt afgeleid uit namen en foto's wordt gediscrimineerd. De namen en foto's en de profielen zeggen vaak iets over de etniciteit, leeftijd en geslacht van personen. 'Judith' is bijvoorbeeld een typische vrouwennaam, terwijl de meeste mensen die 'Rense' of 'Arnout' heten mannen zijn. Er is bijvoorbeeld enig bewijs dat Couchsurfers liever bij vrouwen of families overnachten dan bij mannen (Tan, 2010). In het voorbeeld van de reis naar de Grand Canyon zou ik me meer zorgen kunnen maken om mijn persoonlijke veiligheid omdat de organisator een man was. Andere voorbeelden van discriminatie in de platformeconomie zijn etnische discriminatie (bijv. Edelman & Luca, 2014) en discriminatie op basis van iemands seksuele voorkeur (Ahuja & Lyons, 2019). Mensen geven er in het algemeen de voorkeur aan om om te gaan mensen die op hen lijken (Tajfel, Billig, Bundy, & Flament, 1971). Dat is

niet alleen een probleem voor mensen die lijden onder negatieve discriminatie, maar ook voor de platforms en voor het potentieel van de platformeconomie om een groter publiek te bereiken.

Eerder onderzoek stelt dat reputatiesystemen discriminatie kunnen verminderen (Abrahao, Parigi, Gupta, & Cook, 2017; Cui, Li, & Zhang, 2017; Ert, Fleischer, & Magen, 2016; Mohammed, 2017; Tjaden, Schwemmer, & Khadjavi, 2018). De verklaring daarvoor is dat de aanwezigheid van reputatie-informatie de aandacht voor demografische informatie kan verkleinen. Toegepast op het voorbeeld van de reis naar de Grand Canyon zou dit betekenen dat ik me minder zorgen zou maken over het geslacht van de organisator van de reis als hij één of meer positieve reviews had gehad, omdat ik dan zou weten dat hij te vertrouwen was. Reputatiesystemen zouden er op die manier voor kunnen zorgen dat de verschillen in kansen om mee te doen tussen mensen met verschillende demografische kenmerken kleiner worden.

Dit is echter niet het hele verhaal. In hoofdstuk 2 en 3 beargumenteren we waarom reputatiesystemen mogelijk toch geen oplossing bieden voor het discriminatieprobleem. De reden daarvoor is dat boordelingen alleen geschreven kunnen worden na het voltooien van een transactie. We weten dat het niet voor iedereen even makkelijk is om deel te nemen aan een eerste transactie. Verschillen in de kans om deel te nemen (zoals de kleinere kans om deel te nemen voor gasten met een Afro-Amerikaanse naam) vertalen zich in verschillen in de kans om een eerste recensie te ontvangen. Alleen na het voltooien van een transactie kan namelijk een recensie geschreven worden. Het is dus aannemelijk dat geslacht en etniciteit mede bepalen hoe makkelijk het is voor mensen om een eerste recensie te ontvangen. Die recensies zijn vervolgens belangrijk voor de mogelijkheden om in de toekomst deel te nemen aan transacties. We weten uit eerder onderzoek dat de kans dat een transactieverzoek geaccepteerd wordt veel groter is als de afzender van het verzoek ten minste één positieve beoordeling heeft (Duffy et al., 2013; Frey & Van De Rijt, 2016).

Aan de ene kant kan de aanwezigheid van het reputatiesysteem er dus voor zorgen dat er minder aandacht is voor de etniciteit, het geslacht en de leeftijd van de gebruiker. Aan de andere kant kan het zo zijn dat de demografische kenmerken van de gebruiker bepalen hoeveel recensies die gebruiker verzamelt.

In hoofdstuk 2 en 3 onderzoek ik welk van de twee mechanismen sterker weegt. Ik beantwoord de vraag of reputatiesystemen bestaande ongelijkheden vergroten of verkleinen.

Hoofdstuk 2: Veldonderzoek

In hoofdstuk 2 gebruiken we data van een Nederlands motordeelplatform om te onderzoeken of de aanwezigheid van een reputatiesysteem invloed heeft op ongelijkheid. Via dit platform kunnen motorbezitters hun motor uitlenen. Huurders kunnen de eigenaar van de motor waar zij interesse in hebben een verhuurverzoek sturen. De eigenaar kan dat verzoek accepteren of afwijzen. De dataset bevat informatie over de uitkomst van de aanvraag (geaccepteerd of niet) en over de kenmerken van de huurders en verhuurders (o.a. etniciteit en reputatie).

We hebben onderzocht of er een verschil is in de kans dat verhuurders hun motor uitlenen aan huurders met en zonder een migratieachtergrond. We vinden dat verhuurders meer bereid zijn hun motor te verhuren aan Nederlanders zonder migratieachtergrond. Ongeveer 39% van de aanvragen verstuurd door huurders zonder migratieachtergrond werd geaccepteerd. Respectievelijk 24% en 29% van de aanvragen verstuurd door huurders met een migratieachtergrond werd geaccepteerd. We vinden niet dat etniciteit er minder toe doet als de huurder eenmaal een positieve review heeft. We vinden ook niet dat het reputatiesysteem ervoor zorgt dat de verschillen in de kans dat een aanvraag geaccepteerd wordt tussen huurders met een andere etniciteit kleiner of groter wordt. Al met al vinden we geen bewijs dat het reputatiesysteem helpt om discriminatie te verminderen.

Hoofdstuk 3: Online experiment

In hoofdstuk 3 heb ik op een andere manier onderzocht of reputatiesystemen voor minder ongelijkheid zorgen. Ruim 500 Indiase en Amerikaanse deelnemers hebben een spel gespeeld waarin zij vertrouwen konden plaatsen in één van vier andere personen. Voordat zij een keuze maakten, werden zij op de hoogte gesteld van de nationaliteit van de vier personen waaruit zij mochten kiezen. Zij wisten ook hoe vaak elk van de vier personen in het verleden betrouwbaar en onbetrouwbaar was geweest.

Ook in dit experiment heb ik gevonden dat de nationaliteit en de reputatie van de personen waaruit gekozen kon worden ertoe doen. Gemiddeld genomen werden Amerikanen vaker gekozen dan Indiërs. Het maakte niet uit of de deelnemers in het verleden betrouwbaar waren geweest of niet: de voorkeur voor Amerikanen bleef bestaan, zelfs als één van de Indiase personen een goede reputatie had. Ik vond opnieuw niet dat reputatiesystemen helpen om bestaande ongelijkheid te verkleinen. Ik vond wel dat beslissingen in eerdere ronden van het experiment de uitkomsten in latere ronden beïnvloedden. Als een van de vier personen vroeg in het experiment gekozen werd en betrouwbaar was, werd die persoon later ook vaker gekozen.

Uit deze twee onderzoeken kan geconcludeerd worden dat reputatiesystemen niet per se voor meer gelijkheid zorgen. In beide studies neemt het belang van demografische informatie niet af als reputatie-informatie aanwezig is. Daarnaast krijgt niet iedereen de kans om een reputatie op te bouwen.

Hoofdstuk 4: Reputatiesysteem als stimulans om de waarheid te spreken

De persoonlijke profielen die kenmerkend zijn voor de peer-to-peer platformeconomie bevatten niet alleen namen en foto's. Op de meeste platforms schrijven gebruikers ook iets over zichzelf en eventueel over het product dat zij aanbieden. De auteur van de vraag wie er mee wilde naar de Grand Canyon schreef bijvoorbeeld dat hij een vaardige motorrijder was. Op veel platformen kunnen gebruikers ook met elkaar chatten voor ze besluiten een transactie aan te gaan. In deze beschrijvingen kunnen gebruikers zeggen wat ze willen. Meestal wordt niet gecontroleerd of de feiten kloppen. Ik had bijvoorbeeld geen idee of de auteur van de advertentie voor het weekend weg daadwerkelijk goed kon motorrijden. Toch laat eerder onderzoek zien dat deze beschrijvingen informatie vertrouwen beïnvloedt (Anderson, Friedman, Milam, & Singh, 2007; Lewis, 2011; Rawlins & Johnson, 2007). Hoe kan dit verklaard worden?

In hoofdstuk 4 onderzochten we of reputatiesystemen ervoor kunnen zorgen dat gebruikers niet liegen in hun productomschrijving. We beargumenteren dat reputatiesystemen ervoor kunnen zorgen dat mensen de waarheid spreken als aan twee voorwaarden voldaan wordt. Ten eerste moeten kopers de prijzen

bepalen, zoals bij veilingen het geval is. Ten tweede moet de kwaliteit van de aangeboden producten niet stabiel zijn, zoals bijvoorbeeld het geval is bij tweedehandsproducten. Verkopers hebben op die markten soms betere producten dan op andere momenten.

Onder deze voorwaarden kunnen verkopers hun opbrengst vergroten door een goed product te leveren wanneer dat mogelijk is. Wanneer dat niet mogelijk is, moeten zij het kopers eerlijk laten weten als zij geen goed product hebben. Zij zullen daardoor een reputatie van een eerlijk persoon krijgen. Op de lange termijn kunnen zij daar voordeel van hebben. In hoofdstuk 4 testen we of de aanwezigheid van een reputatiesysteem er inderdaad voor zorgt dat kopers meer belang hechten aan beweringen die gedaan worden door verkopers met een betere reputatie.

In een eerste studie hebben we gebruik gemaakt van data van online veilingplatform eBay. We vinden slechts zwak bewijs dat bieders in veilingen sterker reageren op beweringen van verkopers met een betere reputatie. Kopers gaan in sterke mate af of beweringen die verkopers doen, zelfs als die verkoper geen goede reputatie heeft. Om de mogelijkheid uit te sluiten dat bepaalde verkopers zowel een betere reputatie hebben als meer geloofd worden, testen we deze hypothese nog een keer in een online experiment.

In het experiment speelden ruim 200 deelnemers een spel waarin zij moesten kiezen of ze een verkoper konden vertrouwden of niet. Voordat de keuze gemaakt werd, kon de verkoper ervoor kiezen om een belofte te doen aan de koper. De koper was ook op de hoogte van beslissingen die de andere persoon in het verleden had genomen. Uit het experiment blijkt dat de uitkomst voor kopers het beste zou zijn als zij alleen vertrouwen hadden in verkopers die beloofden vertrouwen niet te misbruiken en als zij dat in het verleden ook niet gedaan hadden. De kopers in het experiment volgden echter niet deze optimale strategie. Ze hechtten altijd waarde aan de beloften van de verkoper, zelfs als die verkoper geen goede reputatie had.

Op basis van deze twee studies concluderen we dat de aanwezigheid van een reputatiesysteem niet kan verklaren waarom kopers beloften van verkopers geloven. Dit blijft dus een puzzel.

Hoofdstuk 5: Vertrouwen in het verleden als signaal van betrouwbaarheid

In het laatste hoofdstuk ga ik tenslotte in op een ander kenmerk van de peer-to-peer platformeconomie. Tot nu toe zijn reputatiesystemen met name bestudeerd in situaties waar de rollen van kopers en verkopers strikt van elkaar gescheiden zijn. Dit onderscheid is echter niet altijd duidelijk op online platformen. Veel gebruikers handelen in de rol van consument in de ene interactie, terwijl zij de rol van aanbieder innemen bij een volgende interactie (Cova, Dalli, & Zwick, 2011; Ritzer & Jurgenson, 2010). De meeste gebruikers van het gastvrijheidsplatform 'Couchsurfing' overnachten bijvoorbeeld bij anderen als zij op reis zijn, maar stellen zelf ook een slaapplaats in hun huis ter beschikking aan andere reizigers (Lauterbach, Truong, Shah, & Adamic, 2009).

Wanneer de grenzen tussen de verschillende rollen niet duidelijk zijn afgebakend, kan een nieuw type reputatiesysteem ontstaan. Dit nieuwe reputatiesysteem bevat zowel informatie over het gedrag van een gebruiker in de ene rol als informatie over het gedrag van die gebruiker in de andere rol. De reputatie van de organisator van de reis naar de Grand Canyon zou zowel informatie over zijn gedrag als organisator van reizen kunnen bevatten, maar ook over zijn gedrag als deelnemer. Dit nieuwe reputatiesysteem voorziet gebruikers van meer en andere informatie. Daardoor verandert het de manier waarop vertrouwen gegenereerd. Hoofdstuk 5 gaat over de vraag op welke manier dit nieuwe reputatiesysteem vertrouwen kan vergroten.

Volgens bestaande theorieën is deze informatie over iemands gedrag in een andere rol van ondergeschikt belang als ook informatie over gedrag in de huidige rol beschikbaar is: Het zou veel relevanter zijn om te weten dat iemand eerder heeft laten zien te vertrouwen te zijn dan dat de persoon zelf anderen heeft vertrouwd. In hoofdstuk 2 beargumenteren we dat individuen die anderen sneller vertrouwen mogelijk ook betrouwbaarder zijn. Dat is vooral het geval als de persoon die vertrouwen plaatste niets kreeg in ruil voor zijn of haar vertrouwen. Zowel netjes omgaan met het vertrouwen van een ander als vertrouwen plaatsen in een ander, levert die ander wat op. Degene die het vertrouwen geplaatst heeft of netjes met andermans vertrouwen is omgegaan heeft er zelf niets aan. Mensen die via spullendeelplatform Peerby spullen uitlenen zonder daar geld voor te vragen, verdienen daar niks aan. Wanneer

mensen geleende spullen op tijd en in een goede staat terugbrengen is vooral de uitlener daarmee geholpen.

Mensen die meer vertrouwen hebben in anderen zijn dus naar verwachting ook betrouwbaarder. Dat betekent dat trustees die in het verleden meer vertrouwen in anderen hebben geplaatst waarschijnlijk ook netjes met het vertrouwen van anderen omgaan in de toekomst. Trustors plaatsen naar verwachting meer vertrouwen in trustees die in het verleden zelf ook vertrouwen in anderen hebben geplaatst. Het reputatiesysteem biedt op die manier ook een mogelijkheden om de eigen reputatie te verbeteren door meer vertrouwen in anderen te plaatsen. In hoofdstuk 5 is deze hypothesen onderzocht met een experiment.

In het experiment kiezen ruim 200 deelnemers of zij vertrouwen in een ander willen plaatsen of niet. Zij weten of die ander in het verleden betrouwbaar is geweest én of die ander in het verleden zelf ook vertrouwen in anderen heeft geplaatst. De resultaten leveren bewijs voor de theorie: wanneer trustees in het verleden meer vertrouwen geplaatst hebben, worden zij in het heden meer vertrouwd. Dit is echter alleen het geval als zij zelf geen profijt hadden van het plaatsen van vertrouwen.

Conclusie en discussie

Voordat ik besloot of ik wel of niet mee wilde naar de Grand Canyon, kon ik de recensies lezen die voormalige reispartners van de organisator hadden geschreven. Volgens klassieke theorieën over reputatiesystemen vergroten deze recensies het vertrouwen doordat zij mij informeren over de mate waarin de organisator in het verleden betrouwbaar is geweest (Buskens, 2003; Buskens & Raub, 2002; Cook et al., 2005; Weigelt & Camerer, 1988). De opkomst van de platformeconomie heeft drie kenmerken aan het licht gebracht die ervoor zorgen dat reputatiesystemen op een andere manier invloed hebben op vertrouwen. Hoewel de opkomst van de platformeconomie onze aandacht op deze kwesties gevestigd heeft, spelen deze kenmerken ook in andere situaties een rol. Het doel van dit proefschrift was bestuderen hoe deze kenmerken beïnvloeden op welke manier reputatiesystemen voor meer vertrouwen zorgen.

In hoofdstuk 2 en 3 vonden we dat reputatiesystemen niet helpen om de ongelijkheid tussen gebruikers met verschillende etnische achtergronden te verkleinen. Dit komt doordat gebruikers alleen een goede reputatie op kunnen bouwen als zij eerst door anderen vertrouwd worden. Niet iedereen wordt in eerste instantie in dezelfde mate vertrouwd. In het voorbeeld van de reis naar de Grand Canyon zou dit bijvoorbeeld betekenen dat mannen minder snel vertrouwd worden, daardoor minder recensies verzamelen, wat er vervolgens weer voor zorgt dat zij ook in de toekomst minder snel vertrouwd worden dan vrouwen. In hoofdstuk 4 vonden we dat de aanwezigheid van een reputatiesysteem niet kan verklaren waarom kopers beloften over de kwaliteit van een product van verkopers geloven. Kopers zijn geneigd beloften van verkopers te geloven, zelfs als zij niet weten of de verkoper de waarheid spreekt. In de context van de reis naar de Grand Canyon betekent dit dat ik eerder achterop de motor zou stappen bij iemand die beweert goed te zijn in motorrijden dan bij iemand die dat niet beweert, zelfs als er geen enkele manier is om te weten of de chauffeur over het algemeen de waarheid spreekt. In hoofdstuk 5 vonden we dat reputatiesystemen die informatie bevatten over verschillende rollen voor meer vertrouwen kunnen zorgen. Dit was alleen het geval als het plaatsen van vertrouwen de persoon die vertrouwen plaatst niet direct voordeel oplevert. Op Couchsurfing betekent dit bijvoorbeeld dat hosts eerder iemand bij hen thuis uitnodigen die zelf ook gasten heeft ontvangen.

Theoretische implicaties

Reputatiesystemen hebben invloed op bestaande ongelijkheden

De resultaten van hoofdstuk 2 en 3 laten zien dat reputatiesystemen niet onbevooroordeeld zijn. Of reputatiesystemen discriminatie verhogen of verlagen hangt naar verwachting af van of het vertrouwensprobleem groot of klein is. Het is niet noodzakelijk het geval dat de meest betrouwbare individuen de meeste positieve reviews verzamelen. In tegendeel: reputatie wordt in opeenvolgende stappen opgebouwd. Beslissingen op een eerder moment beïnvloeden latere beslissingen. Dit betekent dat reputatiesystemen aanvankelijke verschillen tussen mensen kunnen vergroten. In het beste geval zijn deze verschillen willekeurig. De resulterende ongelijkheid is dan ook willekeurig verdeeld over mensen en groepen. Een realistischer scenario is echter dat aanvankelijke verschillen gebaseerd zijn op structurele verschillen

tussen (groepen) mensen, zoals verschillen in demografische eigenschappen. Het effect van reputatiesystemen op individuele beslissingen om iemand te vertrouwen kan in isolatie bestudeerd worden. Het is echter niet mogelijk om de bredere effecten van reputatiesystemen te onderzoeken zonder naar de context te kijken. Om te begrijpen wat het totale effect van reputatiesystemen is, is het belangrijk om te weten wie de gebruikers van een platform zijn, welke vooroordelen een rol spelen en hoe groot het vertrouwensprobleem is.

De beslissing om vertrouwen te plaatsen is beter als die gebaseerd is op verschillende soorten reputatie-informatie

Eerder onderzoek ging ervan uit dat reputatiesystemen informatie verspreiden over de mate waarin een trustee in het verleden betrouwbaar is geweest. Het type informatie dat verspreid werd door een reputatiesysteem was meestal niet het onderwerp van onderzoek. Hoofdstuk 4 en 5 laten zien dat reputatiesystemen ook een andere rol kunnen vervullen. Ze kunnen ook informatie bevatten over de mate waarin een trustee eerlijk is geweest in het verleden. Ook kunnen ze informatie bevatten over de mate waarin een trustee vertrouwen in anderen heeft geplaatst. In beide hoofdstukken laten we zien dat trustors deze informatie zouden kunnen gebruiken om een betere inschatting van de betrouwbaarheid van de trustee te maken. In hoofdstuk 4 maken de trustors niet optimaal gebruik van de beschikbare informatie: ze zijn geneigd beloftes van trustees altijd te geloven, zelfs als ze niet weten of de trustee over het algemeen de waarheid spreekt. In hoofdstuk 5 gebruiken de trustors de informatie wel zoals verwacht: ze hebben meer vertrouwen in trustees die in het verleden ook vertrouwen in anderen hebben gehad, maar alleen als die trustees daar zelf niks voor terugkregen.

De vraag is waarom trustors soms wel en soms niet gebruik maken van de beschikbare informatie, aangezien beide theorieën gebaseerd waren op dezelfde aannamen over het gedrag. We hebben in beide hoofdstukken aangenomen dat trustors en trustees rationeel en egoïstisch zijn. Een eerste verklaring voor het verschil in resultaten is dat trustors minder rationeel zijn dan wij aannamen. Daardoor konden zij mogelijk niet alle gevolgen van hun keuzes overzien. Het is minder ingewikkeld om te begrijpen waarom trustees die in het verleden zelf ook mensen vertrouwd hebben betrouwbaarder zijn, dan om te begrijpen hoe

informatie over een trustee's gedrag in het verleden het best gecombineerd kan worden met nieuwe beloften.

Een andere verklaring is dat trustors niet egoïstisch zijn, of dat zij geloven dat trustees niet egoïstisch zijn. Als de trustors in hoofdstuk 4 geloofden dat de trustees over het algemeen de waarheid spraken, was er geen reden om hun beloften niet te geloven. Rationele trustors zouden met de tijd moeten leren dat niet alle trustees de waarheid spreken. Trustors zouden hun gedrag daaraan aan moeten passen. Klassieke theorieën over reputatiesystemen gaan er meestal vanuit dat mensen rationeel en egoïstisch zijn. Als deze aannamen losgelaten worden, is er ruimte voor nieuwe theorieën die uitgaan van andere mechanismen, zoals wederkerigheid en het gebruik van vuistregels. Toekomstig onderzoek kan ingaan op de vraag of beperkte rationaliteit en egoïsme de verschillen tussen de studies kunnen verklaren.

Informatie in reputatiesystemen beïnvloedt de strategische reputatieopbouw

De aanwezigheid van verschillende soorten reputatie-informatie betekent dat mensen op verschillende manier hun reputatie kunnen opbouwen. Trustees kunnen hun reputatie niet alleen verbeteren door betrouwbaar te zijn. Zij kunnen hun reputatie ook verbeteren door vertrouwen te plaatsen in anderen en door eerlijk te zijn. Deze toegenomen kansen voor het opbouwen van reputatie hebben niet alleen gevolgen voor vertrouwen. Ten eerste kan dit ervoor zorgen dat de correlatie tussen de intrinsieke motivatie van mensen en hun gedrag verzwakt wordt. Zonder reputatiesysteem zullen alleen mensen die dat zelf willen netjes met het vertrouwen van anderen omgaan. Door de aanwezigheid van het reputatiesysteem zullen ook mensen die belang hechten aan hun reputatie zich goed gedragen. Dit betekent dat het moeilijker is om aan de hand van gedrag een onderscheid te maken tussen mensen die het zelf belangrijk vinden om je goed te gedragen en mensen die dat alleen uit eigenbelang doen. Dit kan er op zijn beurt voor zorgen dat de reputatie-informatie zelf minder bruikbaar is.

Een gerelateerd probleem is dat de aanwezigheid van een reputatiesysteem ervoor kan zorgen dat de intrinsieke motivatie van mensen om zich goed te gedragen verzwakt kan worden. Er zijn talloze voorbeelden van onderzoeken waaruit blijkt dat de aanwezigheid van externe stimulansen voor goed gedrag

(zoals reputatie) de intrinsieke motivatie van mensen voor dat gedrag kunnen verminderen (Fehr & Gächter, 2005; Titmuss, 1970). Toekomstig onderzoek kan zich richten op de vraag in hoeverre reputatiesystemen ervoor zorgen dat mensen minder intrinsiek gemotiveerd zijn om betrouwbaar en eerlijk te zijn en om anderen te vertrouwen.

Praktische implicaties

Reputatiesystemen hebben invloed op ongelijkheid

Hoofdstuk 4 en 5 laten zien dat reputatiesystemen aanvankelijke verschillen tussen groepen individuen kunnen bestendigen of zelfs vergroten. Voor platformen die deze verschillen willen verkleinen is het dus niet voldoende om een effectief reputatiesysteem te gebruiken. Zij zouden ook moeten overwegen om foto's en namen van gebruikers een minder prominente plek te geven, zoals Airbnb sinds kort doet (Mohammed, 2017). Daarnaast zouden zij nieuwe gebruikers kunnen helpen bij het krijgen van een eerste review, vooral als die gebruiker tot een etnische minderheid behoort. Een mogelijke manier om dat te doen is het gebruiken van een reputatiesysteem met informatie over gedrag in verschillende rollen (zoals onderzocht in hoofdstuk 5). Op die manier kunnen gebruikers een reputatie opbouwen zonder dat zij eerst vertrouwd hoeven te worden door anderen. Daarna kunnen zij die reputatie gebruiken om het vertrouwen van anderen te winnen.

Platformen kunnen alternatieve designs voor reputatiesystemen overwegen

Bij het ontwerpen en implementeren van het reputatiesysteem zouden platformen in overweging moeten nemen dat reputatiesystemen op verschillende manieren ingericht kunnen worden. Uit dit proefschrift volgt dat manier waarop reputatiesystemen werken afhangt van andere kenmerken van het platform. De bevindingen van hoofdstuk 4 en 5 laten bijvoorbeeld zien dat reputatiesystemen niet alleen het vertrouwen kunnen vergroten omdat ze informatie bevatten over betrouwbaarheid in het verleden. Ze kunnen ook voor meer vertrouwen zorgen omdat ze informatie over de mate van vertrouwen en eerlijkheid bevatten. Dit betekent dat er geen one-size-fits-all oplossing is. Platformen moeten nagaan hoe hun platform er verder uit ziet en dus op welke manier het reputatiesysteem effect heeft op het gedrag van haar gebruikers.

Reflectie op gebruikte methoden

In dit proefschrift is gebruik gemaakt van verschillende methoden: analyse van data van echte platformen, laboratoriumexperimenten en online experimenten. Het voordeel van experimenten in het lab en online is dat het de mogelijkheid biedt om het effect van de aanwezigheid van een reputatiesysteem te bestuderen. In het veld is dat lastig, omdat echte platformen voor alle gebruikers hetzelfde ingericht zijn: ze hebben ofwel voor alle gebruikers een reputatiesysteem, of ze hebben helemaal geen reputatiesysteem. Een ander probleem met het gebruik van data van echte platformen is dat zij vaak hun data niet willen, kunnen of mogen delen met onderzoekers (Cox & Slee, 2016).

Het nadeel van experimenten is dat de bestudeerde situatie tamelijk abstract is. In deze experimenten hebben we geprobeerd om de belangrijkste kenmerken van de situatie op echte platformen samen te vatten in spellen. Door het wegnemen van de context zijn de resultaten mogelijk op een groot aantal verschillende platformen van toepassing waar in de basis hetzelfde vertrouwensprobleem speelt. Het is echter ook mogelijk dat door het wegnemen van de context belangrijke kenmerken wegvallen die de resultaten veranderen. Hoewel de mate waarin de resultaten van toepassing zijn op echte platformen dus mogelijk beperkt is, komen de resultaten van de experimenten en de veldstudies in dit proefschrift in de meeste gevallen overeen. Dit suggereert dat de abstracte experimenten inderdaad bruikbaar kunnen zijn voor het onderzoeken van specifieke platformen.

Beperkingen en toekomstig onderzoek

De onderzoeken in dit proefschrift hebben ook een aantal beperkingen. In dit proefschrift richt ik mij op de informatie die over gebruikers van platformen verschaft wordt. Ik heb onderzoek hoe dit informatie helpt het vertrouwensprobleem op te lossen. Ik focus niet op andere oplossingen voor het vertrouwensprobleem, zoals verzekeringen. Toekomstig onderzoek kan zich richten op de vraag op welke manier het functioneren van het reputatiesysteem afhangt van andere eigenschappen van platformen.

Een tweede beperking is dat het in de experimenten mogelijk was om te zoeken of vertrouwen misplaatst was of niet, terwijl dat in de empirische studies niet mogelijk was. Ik weet bijvoorbeeld niet of de motoren die verhuurd waren via

het motordeelplatform beschadigd werden tijdens de verhuurperiode (hoofdstuk 2). Toekomstig onderzoek zou bijvoorbeeld aan de hand van klachten die de platformen ontvangen kunnen onderzoeken in hoeverre het plaatsen van vertrouwen leidt tot problemen.

Een derde beperking is dat in de experimenten is aangenomen dat geen nieuwe gebruikers tot de markt toetreden en dat zij de markt ook niet verlaten. Het kan interessant zijn om te onderzoeken of nieuwe toetreders benadeeld zijn ten opzichte van gevestigde gebruikers.

Een andere vraag voor toekomstig onderzoek kan gaan over de mate waarin positieve recensies ervoor zorgen dat minder aandacht naar de etniciteit van de trustee gaat. Hoewel dit effect in eerdere onderzoeken werd gevonden, vinden wij hier in hoofdstuk 2 en 3 geen bewijs voor. Een reden hiervoor kan zijn dat eerdere onderzoeken niet naar etniciteit hebben gekeken (Abrahao et al., 2017), of niet naar daadwerkelijke transacties (Tjaden et al., 2018). Andere onderzoekers hebben geen echte gebruikersprofielen gebruikt, maar 'lege' profielen die zij zelf gemaakt hebben (Cui et al., 2016). Een laatste studie vond geen bewijs voor het effect, tenzij de verschillen in reputatiescores tussen gebruikers vergroot werden (Ert et al., 2016). Toekomstig onderzoek kan ingaan op de voorwaarden van het bestaan van dit 'compensatie-effect'.

Andere onderwerpen voor toekomstig onderzoek zijn de rol van algoritmes die gebruikers rangschikken op basis van bijvoorbeeld reputatie en reputatiesystemen die op meerdere platformen gebruikt kunnen worden. Met name de vraag op welke manier dit soort ontwikkelingen ongelijkheid tussen gebruikers beïnvloeden is belangrijk. Toekomstig onderzoek kan ook gaan over de overdraagbaarheid van reputatie van het ene naar het andere platform. Een probleem van het hebben van één reputatie per platform, is dat het lastig is voor gebruikers om over te stappen naar een ander platform. Als ze overstappen, verliezen ze namelijk hun beoordelingen. Reputatiesystemen die informatie van meer dan één platform bevatten zouden hier een antwoord op kunnen zijn. We weten echter nog weinig over de bruikbaarheid van informatie op een ander platform. We weten ook niet of dit voor meer of juist minder ongelijkheid kan zorgen.

REFERENCES

Abrahao, B., Parigi, P., Gupta, A., & Cook, K. S. (2017). Reputation offsets trust judgments based on social biases among Airbnb users. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(37), 9848–9853.

Ahuja, R., & Lyons, R. C. (2019). The Silent Treatment: LGBT Discrimination in the Sharing Economy. *Oxford Economic Papers*, *71*(3), 564–576.

Akerlof, G. (1970). The market for lemons: Qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics*, *84*(3), 235–251.

Anderson, S. T., Friedman, D., Milam, G. H., & Singh, N. (2007). Seller Strategies on eBay: Does size matter? *International Journal of Electronic Business*, *5*(6), 643–669.

Arrow, K. (1973). The Theory of Discrimination. *Discrimination in Labor Markets*, *3*(10), 3–33.

Axelrod, R., & Dion, D. (1988). The Further Evolution of Cooperation. *Science*, *242*(4884), 1385–1390.

Ba, S., & Pavlou, P. A. (2002). Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior. *MIS Quarterly*, *26*(3), 243–268.

Balliet, D. (2010). Communication and cooperation in social dilemmas: A meta-analytic review. *Journal of Conflict Resolution*, *54*(1), 39–57.

Balliet, D., Mulder, L. B., & Van Lange, P. A. M. (2011). Reward, punishment, and cooperation: a meta-analysis. *Psychological Bulletin*, *137*(4), 594–615.

Balliet, D., Wu, J., & De Dreu, C. K. W. (2014). Ingroup Favoritism in Cooperation: A Meta-Analysis. *Psychological Bulletin*, *141*(6), 1556–1581.

Becker, G. S. (1957). *The Economics of Discrimination.* Chicago: University of Chicago Press.

Benard, S. (2013). Reputation systems, aggression, and deterrence in social interaction. *Social Science Research*, *42*(1), 230–245.

Bente, G., Baptist, O., & Leuschner, H. (2012). To buy or not to buy: Influence of seller photos and reputation on buyer trust and purchase behavior. *International Journal of Human-Computer Studies*, *70*(1), 1–13.

Boero, R., Bravo, G., Castellani, M., & Squazzoni, F. (2009). Reputational cues in repeated trust games. *The Journal of Socio-Economics*, *38*(6), 871–877.

Bolton, G. E., Katok, E., & Ockenfels, A. (2004a). How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation. *Management Science*, *50*(11), 1587–1602.

Bolton, G. E., Katok, E., & Ockenfels, A. (2004b). Trust among Internet Traders. *Analyse & Kritik*, *26*(1), 185–202.

Bolton, G. E., Katok, E., & Ockenfels, A. (2005). Cooperation among strangers with limited information about reputation. *Journal of Public Economics*, *89*(8), 1457–1468.

Bridges, J., & Vásquez, C. (2018). If nearly all Airbnb reviews are positive, does that make them meaningless? *Current Issues in Tourism*, *21*(18), 2065–2083.

Brosig-Koch, J., & Heinrich, T. (2018). The role of communication content and reputation in the choice of transaction partners: A study based on field and laboratory data. *Games and Economic Behavior*, *112*, 49–66.

Buskens, V. (2003). Trust in triads: Effects of exit, control, and learning. *Games and Economic Behavior*, *42*(2), 235–252.

Buskens, V., & Raub, W. (2002). Embedded trust: Control and learning. *Group Cohesion, Trust and Solidarity*, *19*, 167–202.

Buskens, V., Raub, W., & van der Veer, J. (2010). Trust in triads: An experimental study. *Social Networks*, *32*(4), 301–312.

Buskens, V., & Weesie, J. (2000). An experiment on the effects of embeddedness in trust situations: Buying a used car. *Rationality and Society*, *12*(2), 227–253.

Cabral, L., & Hortaçsu, A. (2010). The dynamics of seller reputation: Evidence from eBay. *The Journal of Industrial Economics*, *58*(1), 54–78.

Cason, T. N., & Mui, V. L. (2014). Coordinating resistance through communication and repeated interaction. *Economic Journal*, *124*(574), 226–256.

Charness, G., Du, N., & Yang, C. L. (2011). Trust and trustworthiness reputations in an investment game. *Games and Economic Behavior*, *72*(2), 361–375.

Chaudhuri, A., & Gangadharan, L. (2007). An Experimental Analysis of Trust and Trustworthiness. *Southern Economic Journal*, *73*(4), 959–985.

Chen, D. L., Schonger, M., & Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88–97.

Cheshire, C. (2007). Selective Incentives and Generalized Information Exchange. *Social Psychology Quarterly*, *70*(1), 82–100.

Coleman, J. S. (1994). *Foundations of Social Theory*. Harvard University Press.

Collins, L. M., Schafer, J. L., & Kam, C. M. (2001). A comparison of inclusive and restrictive strategies in modern missing data procedures. *Psychological Methods*, *6*(4), 330–351.

Cook, K. S. (2005). Networks, Norms, and Trust: The Social Psychology of Social Capital. *Social Psychology Quarterly*, *68*(1), 4–14.

Cook, K. S., Hardin, R., & Levi, M. (2005). *Cooperation Without Trust?* New York: Russell Sage Foundation.

Corten, R. (2019). Social dilemmas in the sharing economy. In R. W. Belk, G. M. Eckhardt, & F. Bardhi (Eds.), *Handbook of the Sharing Economy* (pp. 278–289). Edward Elgar Publishing.

Cova, B., Dalli, D., & Zwick, D. (2011). Critical perspectives on consumers' role as "producers": Broadening the debate on value co-creation in marketing processes. *Marketing Theory*, *11*(3), 231–241.

Cox, M., & Slee, T. (2016). *How Airbnb's data hid the facts in New York City*.

Cui, R., Li, J., & Zhang, D. J. (2017). *Discrimination with Incomplete Information in the Sharing Economy: Evidence from Field Experiments on Airbnb*.

Currarine, S., & Mengel, F. (2016). Identity, homophily and in-group bias. *European Economic Review*, *60*, 40–55.

Dasgupta, P. (1988). Trust as a Commodity. In D. Gambetta (Ed.), *Trust: Making an Breaking Cooperative Relations* (pp. 49–72). Cornwall: Basil Blackwell.

Dellarocas, C. (2010). Online reputation systems: How to design one that does what you need. *MIT Sloan Management Review*, *51*(3), 33–38.

Demary, V. (2015). *Competition in the sharing economy* (IW Policy Paper No. 19). Cologne.

Denant-Boemont, L., Masclet, D., & Noussair, C. (2011). Announcement, observation and honesty in the voluntary contributions game. *Pacific Economic Review*, *16*(2), 207–228.

Dewally, M., & Ederington, L. (2006). Reputation, Certification, Warranties, and Information as Remedies for Seller-Buyer Information Asymmetries: Lessons from the Online Comic Book Market. *The Journal of Business*, *79*(2), 693–729.

Dewan, S., Hsu, V., Dewanw, S., Hsuz, V., & Rogers, M. (2004). Adverse Selection in Electronic Markets: Evidence from online stamp auctions. *The Journal of Industrial Economics*, *52*(4), 497–516.

Diekmann, A., Jann, B., Przepiorka, W., & Wehrli, S. (2014). Reputation Formation and the Evolution of Cooperation in Anonymous Online Markets. *American Sociological Review*, *79*(1), 65–85.

Difallah, D., Filatova, E., & Ipeirotis, P. (2018). Demographics and dynamics of Mechanical Turk workers. In *WSDM 2018 - Proceedings of the 11th ACM International Conference on Web Search and Data Mining* (pp. 135–143).

Diprete, T. A., & Eirich, G. M. (2006). Cumulative Advantage as a Mechanims of Inequality: A Review of Theoretical and Empirical Developments. *Annu. Rev. Sociol*, *32*, 271–97.

Dubois, D., Willinger, M., & Blayac, T. (2012). Does players' identification affect trust and reciprocity in the lab? *Journal of Economic Psychology*, *33*(1), 303–317.

Duffy, J., & Feltovich, N. (2006). Words, deeds, and lies: Strategic behaviour in games with multiple signals. *Review of Economic Studies*, *73*(3), 669–688.

Duffy, J., Xie, H., & Lee, Y. J. (2013). Social norms, information, and trust among strangers: theory and evidence. *Economic Theory*, *52*(2), 669–708.

eBay. (2020a). Selling through auctions. Retrieved February 19, 2020, from https://www.ebay.com/help/selling/listings/selling-auctions?id=4110

eBay. (2020b). Setting a reserve price . Retrieved February 19, 2020, from https://www.ebay.com/help/selling/listings/selling-auctions/reserve-prices?id=4143

Edelman, B., & Luca, M. (2014). *Digital Discrimination: The Case of Airbnb.com* (Harvard Business School NOM Unit Working paper No. 14–054).

Edelman, B., Luca, M., & Svirsky, D. (2017). Racial Discrimination in the

Sharing Economy: Evidence from a Field Experiment. *American Economic Journal: Applied Economics*, *9*(2), 1–22.

Elfenbein, D. W., Fisman, R., & Mcmanus, B. (2012). Charity as a substitute for reputation: Evidence from an online marketplace. *Review of Economic Studies*, *79*(4), 1441–1468.

Engelmann, D., & Fischbacher, U. (2009). Indirect reciprocity and strategic reputation building in an experimental helping game. *Games and Economic Behavior*, *67*(2), 399–407.

Ert, E., & Fleischer, A. (2019). The evolution of trust in Airbnb: A case of home rental. *Annals of Tourism Research*, *75*, 279–287.

Ert, E., Fleischer, A., & Magen, N. (2016). Trust and Reputation in the Sharing Economy: The Role of Personal Photos on Airbnb. *Tourism Management*, *55*, 62–73.

Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, *90*(4), 980–994.

Fehr, E., & Gächter, S. (2005). *Do incentive contracts undermine voluntary cooperation?* (Zurich IEER Working Paper No. 34).

Fehrler, S., & Przepiorka, W. (2013). Charitable giving as a signal of trustworthiness: Disentangling the signaling benefits of altruistic acts. *Evolution and Human Behavior*, *34*(2), 139–145.

Fehrler, S., & Przepiorka, W. (2016). Choosing a partner for social exchange: Charitable giving as a signal of trustworthiness. *Journal of Economic Behavior and Organization*, *129*, 157–171.

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.

Fradkin, A., Grewal, E., & Holtz, D. (2018). *The Determinants of Online Review Informativeness: Evidence from Field Experiments on Airbnb* (No. 2939064). *SSRN Electronic Journal*.

Frenken, K., & Schor, J. (2017). Putting the sharing economy into perspective. *Environmental Innovation and Societal Transitions*, *23*, 3–10.

Frey, B., & Oberholzer-Gee, F. (1997). The cost of price incentives: An empirical analysis of motivation crowding-out. *The American Economic Review*, *87*(4), 746–755.

Frey, V., & Van De Rijt, A. (2016). Arbitrary Inequality in Reputation Systems. *Scientific Reports*, *6*(1), 1–5.

Gambetta, D. (2009). Signaling. In *The Oxford Handbook of Analytical Sociology* (pp. 168–194). Oxford: Oxford University Press.

Ganzach, Y., Saporta, I., & Weber, Y. (2000). Interaction in Linear Versus Logistic Models: A Substantive Illustration Using the Relationship Between Motivation, Ability, and Performance. *Organizational Research Methods*, *3*(3), 237–253.

Ge, Y., Knittel, C. R., MacKenzie, D., & Zoepf, S. (2016). *Racial and gender discrimination in transportation network companies* (No. w22776). *NBER Working Paper Series*.

Geiger, A. (2015). Thank Me For Hosting: the Role of Reciprocity in Sharing. *Advances in Consumer Research*, *43*, 534–535.

Geiger, A., & Germelmann, C. C. (2015). Reciprocal Couchsurfing Versus Sharing's Non-Reciprocity Principle. In *Proceedings of the 44th EMAC Conference*.

Gong, B., & Yang, C. L. (2010). *Reputation and Cooperation: An Experiment on Prisoner's Dilemma with Second-Order Information* (No. 1549605). *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.1549605

Gracia, P., Vázquez-Quesada, L., & Van de Werfhorst, H. G. (2016). Ethnic penalties? The role of human capital and social origins in labour market outcomes of second-generation Moroccans and Turks in the Netherlands. *Journal of Ethnic and Migration Studies*, *42*(1), 69–87.

Greiner, B. (2015). Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE. *Journal of the Economic Science Association*, *1*(1), 114–125.

Grybaitė, V., & Stankevičienė, J. (2016). Motives for participation in the sharing economy – evidence from Lithuania. *Ekonomia i Zarzadzanie*, *8*(4), 7–17.

Guttentag, D. (2015). Airbnb: disruptive innovation and the rise of an informal tourism accommodation sector. *Current Issues in Tourism*, *18*(12), 1192–1217.

Hamari, J., Sjöklint, M., & Ukkonen, A. (2015). The Sharing Economy: Why People Participate in Collaborative Consumption. *Journal of the Association*

for Information Science and Technology, *67*(9), 2047–2059.

Hars, A., & Ou, S. (2002). Working for Free? Motivations for Participating in Open-Source Projects. *International Journal of Electronic Commerce*, *6*(3), 25–39.

Hoffman, E., McCabe, K., & Smith, V. L. (1996). Social Distance and Other-Regarding Behavior in Dictator Games. *American Economic Review*, *86*(3), 653–660.

Hofstra, B., & de Schipper, N. C. (2018). Predicting ethnicity with first names in online social media networks. *Big Data & Society*, *5*(1), 1–14.

Hooshmand, D. (2018). Everything you wanted to know about Riders Share: How to Make Passive Income with a Motorcycle…. Retrieved November 15, 2018, from https://medium.com/@hoosh/everything-you-wanted-to-know-about-riders-share-how-to-make-passive-income-with-a-motorcycle-c6f59013e978

Houser, D., & Wooders, J. (2006). Reputation in auctions: Theory, and evidence from eBay. *Journal of Economics and Management Strategy*, *15*(2), 353–369.

Jaeger, B., & Sleegers, W. (2020). Racial discrimination in the sharing economy: Evidence from Airbnb markets across the world.

Jaeger, B., & van Beest, I. (2019). The effects of facial attractiveness and trustworthiness in the sharing economy. *Journal of Economic Psychology*, *75*, 102125.

Jin, G. Z., & Kato, A. (2006). Price, quality, and reputation: evidence from an online field experiment. *The RAND Journal of Economics*, *37*(4), 983–1005.

Kahneman, D. (2002). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, *93*(5), 1449–1475.

Katz, V. (2015). Regulating the Sharing Economy. *Berkeley Technology Law Journal*, *30*(4), 11–29.

Kerr, N. L., Garst, J., Lewandowski, D. A., & Harris, S. E. (1997). That Still, Small Voice: Commitment to Cooperate as an Internalized Versus a Social Norm. *Personality and Social Psychology Bulletin*, *23*(12), 1300–1311.

Kiyonari, T., Tanida, S., & Yamagishi, T. (2000). Social exchange and reciprocity: Confusion or a heuristic? *Evolution and Human Behavior*, *21*(6), 411–427.

Klinger, S. (2018). Riders Share: New Motorcycle Sharing Service Launched - ADV Pulse. Retrieved November 15, 2018, from https://www.advpulse.com/adv-news/motorcycle-sharing-service-riders-share/

Kollock, P. (1998). Social Dilemmas: The Anatomy of Cooperation. *Annual Review of Sociology*, *24*(1), 183–214.

Kreps, D. M. (1996). Corporate culture and economic theory. In J. Alt & K. Shepsle (Eds.), *Firms, Organizations and Contracts* (pp. 221–175). Oxford: Oxford University Press.

Kroher, M., & Wolbring, T. (2015). Social control, social learning, and cheating: Evidence from lab and online experiments on dishonesty. *Social Science Research*, *53*, 311–324.

Kuwabara, K. (2015). Do Reputation Systems Undermine Trust? Divergent Effects of Enforcement Type on Generalized Trust and Trustworthiness. *American Journal of Sociology*, *120*(5), 1390–1428.

Laouénan, M., Rathelot, R., Autor, D., Bagues, M., Boustan, L., De Chaisemartin, C., … Zinovyeva, N. (2017). *Ethnic Discrimination on an Online Marketplace of Vacation Rental Ethnic Discrimination on an Online Marketplace of Vacation Rentals*.

Lauterbach, D., Truong, H., Shah, T., & Adamic, L. (2009). Surfing a web of trust: Reputation and reciprocity on Couchsurfing.com. In *2009 International Conference on Computational Science and Engineering* (Vol. 4, pp. 346–353).

Leong, N., & Belzer, A. (2017). The new public accommodations: race discrimination in the platform economy. *Georgetown Law Journal*, *105*(5), 1271–1322.

Lewis, G. (2011). Asymmetric information, adverse selection and online disclosure: The case of eBay motors. *American Economic Review*, *101*(4), 1535–1546.

Lewis, G. (2019). Replication data for: Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors. Nashville: TN: American Economic Association [publisher], 2011. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor].

Macy, M. W., & Skvoretz, J. (1998). The Evolution of Trust and Cooperation

between Strangers: A Computational Model. *American Sociological Review*, *63*(5), 638–660.

Malinowski, B. (1922). *Argonauts of the Western Pacific: An account of native enterprise and adventure in the archipelagoes of Melanesian New Guinea*. London: Routlegde & Sons.

Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, *20*(3), 709–734.

McKinsey. (2016). *Independent work: Choice, necessity, and the gig economy*.

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, *27*(1), 415–444.

Mohammed, A. R. S. (2017). *Designing for Racial Impartiality: The Impact of Relocating Host Photos on the Airbnb Website*.

Molm, L. D. (2010). The structure of reciprocity. *Social Psychology Quarterly*, *73*(2), 119–131.

Molm, L. D., Collett, J. L., & Schaefer, D. R. (2007). Building Solidarity through Generalized Exchange: A Theory of Reciprocity. *American Journal of Sociology*, *113*(1), 205–242.

Molm, L. D., Schaefer, D. R., & Collett, J. L. (2009). Fragile and Resilient Trust: Risk and Uncertainty in Negotiated and Reciprocal Exchange. *Sociological Theory*, *27*(1), 1–32.

MTurk. (2019). Retrieved August 26, 2020, from https://www.mturk.com/

Murphy, L. (2016). *Airbnb's Work to Fight Discrimination and Build Inclusion: A Report Submitted to Airbnb*.

Muthén, L. K., & Muthén, B. O. (2017). Mplus User's Guide (Eigth edition). Los Angeles, CA: Muthén & Muthén.

Neto, J. Q. F., Bloemhof, J., & Corbett, C. (2016). Market prices of remanufactured, used and new items: Evidence from eBay. *International Journal of Production Economics*, *171*(3), 371–380.

Norbutas, L., Ruiter, S., & Corten, R. (2020a). Believe it when you see it: dyadic embeddedness and reputation effects on trust in cryptomarkets for illegal drugs. *Social Networks*, *63*, 150–161.

Norbutas, L., Ruiter, S., & Corten, R. (2020b). Reputation transferability across contexts: Maintaining cooperation among anonymous

cryptomarket actors when moving between markets. *International Journal of Drug Policy*, *76*, 102635.

Nunley, J. M., Owens, M. F., & Howard, R. S. (2011). The effects of information and competition on racial discrimination: Evidence from a field experiment. *Journal of Economic Behavior & Organization*, *80*(3), 670–679.

Pager, D., & Shepherd, H. (2008). The Sociology of Discrimination: Racial Discrimination in Employment, Housing, Credit, and Consumer Markets. *Annual Review of Sociology*, *34*(1), 181–209.

Parigi, P., Santana, J. J., & Cook, K. S. (2017). Online Field Experiments: Studying Social Interactions in Context. *Social Psychology Quarterly*, *80*(1), 1–19.

Pope, D. G., & Sydnor, J. R. (2011). What's in a Picture? Evidence of Discrimination from Prosper.com. *Journal of Human Resources*, *46*(1), 53–92.

Prolific. (2019). Retrieved from www.prolific.co

Przepiorka, W., & Berger, J. (2017). Signalling theory evolving: signals and signs of trustworthiness in social exchange. In *Social Dilemmas, Institutions and the Evolution of Cooperation.* (pp. 373–392). Berlin: De Gruyter Oldenbourg.

Przepiorka, W., & Diekmann, A. (2013). Temporal embeddedness and signals of trustworthiness: Experimental tests of a game theoretic model in the United Kingdom, Russia, and Switzerland. *European Sociological Review*, *29*(5), 1010–1023.

Przepiorka, W., Norbutas, L., & Corten, R. (2017). Order without Law: Reputation Promotes Cooperation in a Cryptomarket for Illegal Drugs. *European Sociological Review*, *33*(6), 752–764.

Quora. (2016). CouchSurfing. Retrieved December 15, 2016, from https://www.quora.com/As-a-couchsurfing-host-I-always-end-up-out-of-pocket-because-surfers-either-have-to-be-reminded-to-pay-for-their-own-food-or-at-least-chip-in-for-it-Is-anyone-else-finding-the-same

Ranchordas, S. (2015). Does Sharing Mean Caring: Regulating Innovation in the Sharing Economy. *Minnesota Journal of Law, Science and Technology*, *16*(1), 413–476.

Raub, W. (2004). Hostage Posting as a Mechanism of Trust: Binding, Compensation, and Signaling. *Rationality and Society*, *16*(3), 319–365.

Rawlins, C., & Johnson, P. (2007). Selling On eBay: Persuasive Communication Advice Based On Analysis Of Auction Item Descriptions. *Journal of Strategic E-Commerce*, *5*(2), 75–82.

Reddit. (2015). Couchsurfing. Retrieved December 15, 2016, from https://www.reddit.com/r/couchsurfing/comments/3dlfqi/why_is_it_so_hard_for_guys_to_find_a_host/

Resnick, P., Kuwabara, K., Zeckhauser, R., & Friedman, E. (2000). Reputation systems. *Communications of the ACM*, *43*(12), 45–48.

Resnick, P., & Zeckhauser, R. (2002). Trust among strangers in internet transactions: Empirical analysis of eBay' s reputation system. *Advances in Applied Microeconomics*, *11*(2), 127–157.

Resnick, P., Zeckhauser, R., Swanson, J., & Lockwood, K. (2006). The value of reputation on eBay: A controlled experiment. *Experimental Economics*, *9*(2), 79–101.

Rich, J. (2014). *What Do Field Experiments of Discrimination in Markets Tell Us? A Meta Analysis of Studies Conducted since 2000* (IZA Discussion Papers No. 8584).

Ritzer, G., & Jurgenson, N. (2010). Production, Consumption, Prosumption: The nature of capitalism in the age of the digital "prosumer." *Journal of Consumer Culture*, *10*(1), 13–36.

Robbins, B. G. (2016). Probing the Links Between Trustworthiness, Trust, and Emotion: Evidence From Four Survey Experiments. *Social Psychology Quarterly*, *79*(3), 284–308.

Robbins, B. G. (2017). Status, identity, and ability in the formation of trust. *Rationality and Society*, *29*(4), 408–448.

Roof, K. (2019, August 19). Airbnb's First-Quarter Financials Show Strong Growth Rate, Cash Reserves. *The Wall Street Journal.*

Rooks, G., Raub, W., & Tazelaar, F. (2006). Ex post problems in buyer-supplier transactions: Effects of transaction characteristics, social embeddedness, and contractual governance. *Journal of Management and Governance*, *10*(3), 239–276.

Rousseau, D. M., Sitkin, S. B., Burt, R. S., & Camerer, C. (1998). Not so

different after all: A cross-discipline view of trust. *Academy of Management Review*, *23*(3), 393–404.

Schniter, E., Sheremeta, R. M., & Sznycer, D. (2013). Building and rebuilding trust with promises and apologies. *Journal of Economic Behavior and Organization*, *94*, 242–256.

Seinen, I., & Schram, A. (2006). Social status and group norms: Indirect reciprocity in a repeated helping experiment. *European Economic Review*, *50*(3), 581–602.

Simon, H. A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, *69*(1), 99–118.

Snijders, C., & Zijdeman, R. (2004). Reputation and internet auctions: eBay and beyond. *Analyse & Kritik*, *26*(1), 158–184.

Stanca, L. (2009). Measuring indirect reciprocity: Whose back do we scratch? *Journal of Economic Psychology*, *30*(2), 190–202.

Statistics Netherlands. (2018). Definitions: Native Dutch. Retrieved December 12, 2018, from https://jeugdmonitor.cbs.nl/en/definitions/Native Dutch

Statistics Netherlands. (2019). Hoeveel mensen met een migratieachtergrond wonen in Nederland? Retrieved October 29, 2019, from https://www.cbs.nl/nl-nl/dossier/dossier-asiel-migratie-en-integratie/hoeveel-mensen-met-een-migratieachtergrond-wonen-in-nederland-

Tadelis, S. (2016). Reputation and Feedback Systems in Online Platform Markets. *Annual Review of Economics*, *8*(1), 321–340.

Tajfel, H., Billig, M. G., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behaviour. *European Journal of Social Psychology*, *1*, 149–178.

Takahashi, N. (2000). The Emergence of Generalized Exchange. *American Journal of Sociology*, *105*(4), 1105–1134.

Takahashi, N., & Mashima, R. (2003). The Emergence of Indirect Reciprocity: Is the standing strategy the answer? *Working Paper Series. Center for the Study of Cultural and Ecological Foundations of the Mind, Hokkaido University, Japan.*, *29*.

Tan, J. E. (2010). The Leap of Faith from Online to Offline: An Exploratory

Study of Couchsurfing.org. In *Proceedings of International Conference on Trust and Trustworthy Computing* (pp. 367–380).

ter Huurne, M., Ronteltap, A., Corten, R., & Buskens, V. (2017). Antecedents of trust in the sharing economy: A systematic review. *Journal of Consumer Behaviour*, *16*(6), 485–498.

ter Huurne, M., Ronteltap, A., Guo, C., Corten, R., & Buskens, V. (2018). Reputation Effects in Socially Driven Sharing Economy Transactions. *Sustainability*, *10*(8), 2674.

Teubner, T. (2017). The web of host-guest connections on Airbnb - A social network perspective. *Journal of Systems and Information Technology*, *20*(3), 262–277.

Teubner, T., Adam, M. T. P., & Hawlitschek, F. (2019). Unlocking Online Reputation: On the Effectiveness of Cross-Platform Signaling in the Sharing Economy. *Business and Information Systems Engineering*, *61*(5), 1–13.

Teubner, T., & Glaser, F. (2018). Up or out: The dynamics of star rating scores on Airbnb. In *ECIS 2018 Proceedings*.

Teubner, T., Hawlitschek, F., & Dann, D. (2017). Price determinants on Airbnb: How repuation pays off in the sharing economy. *Journal of Self-Governance and Management Economics*, *5*(4), 53–80.

Titmuss, R. M. (1970). *The gift relationship: from human blood to social policy*. London: Allen and Unwin.

Tjaden, J. D., Schwemmer, C., & Khadjavi, M. (2018). Ride with me - Ethnic discrimination in Social Markets. *European Sociological Review*, *34*(4), 418–432.

Traity. (2020). Retrieved August 19, 2020, from https://traity.com/

Weigelt, K., & Camerer, C. (1988). Reputation and corporate strategy: A review of recent theory and applications. *Strategic Management Journal*, *9*(5), 443–454.

Wilson, R. K., & Sell, J. (1997). "Liar, Liar... ": Cheap Talk and Reputation in Repeated Public Goods Settings. *Journal of Conflict Resolution*, *41*(5), 695–717.

Wu, J., & Jin, M. (2018). Signaling Peer Trust in Accommodation-sharing Services: Effects of Similarity and Reviews on Listing Sales. In *Wuhan International Conference on e-Business*.

Wu, J., Ma, P., & Xie, K. L. (2017). In sharing economy we trust: the effects of host attributes on short-term rental purchases. *International Journal of Contemporary Hospitality Management*, *29*(11), 2962–2976.

Yamagishi, T., & Yamagishi, M. (1994). Trust and commitment in the United States and Japan. *Motivation and Emotion*, *18*(2), 129–166.

Zervas, G., Proserpio, D., & Byers, J. (2015). *A First Look at Online Reputation on Airbnb, Where Every Stay is Above Average* (No. 2554500). *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2554500

Zschirnt, E., & Ruedin, D. (2016). Ethnic discrimination in hiring decisions: a meta-analysis of correspondence tests 1990–2015. *Journal of Ethnic and Migration Studies*, *42*(7), 1115–1134.

DANKWOORD

(Acknowledgements)

Hoewel alleen mijn naam op de voorkant van dit boekje staat, zou het schrijven van dit proefschrift zeker niet gelukt zijn zonder een aantal andere personen, die ik hier graag zou willen bedanken. Ik wil ook graag mijn dank uitspreken voor iedereen die ervoor gezorgd heeft dat de afgelopen vier jaar in Utrecht voorbij gevlogen zijn!

Ten eerste wil ik graag Rense en Arnout bedanken voor hun begeleiding de afgelopen jaren. Ik kwam naar Utrecht met een nogal brede achtergrond die niet echt gericht was op het doen van onderzoek vanuit een theoretisch perspectief. Ik wil jullie bedanken omdat jullie me de kans hebben gegeven om dit project te doen en voor al jullie begeleiding. Jullie hebben echt de tijd genomen om mij te leren hoe ik een goed onderzoek kan opzetten en ik denk met veel plezier terug aan onze afspraken. Ik vond het leuk dat jullie soms op heel verschillende manieren naar een probleem keken. Ik heb genoten van de discussies en de nieuwe inzichten die dat opleverde. Ook buiten de afspraken kon ik bij jullie terecht. Het gebeurde regelmatig dat we tijdens de afspraak niet direct een oplossing voor aan probleem vonden, maar als ik dan op dezelfde dag of een paar dagen later een nieuw idee wilde bespreken, hadden jullie daar bijna altijd wel tijd voor.

Rense, ik wil je graag bedanken voor het geduld waarmee je mij geleerd hebt wat een goede puzzel is en voor alle kennis over de deeleconomie en vertrouwen die je gedeeld hebt. Naast de inhoudelijke begeleiding wil ik je ook bedanken voor de begeleiding op andere vlakken. Je hebt me een hoop geleerd over hoe dingen werken in de wetenschap en over wat er allemaal komt kijken bij het doen van goed onderzoek. Ik vond het fijn dat ik altijd bij je terecht kon met al mijn vragen!

Arnout, het kwam regelmatig voor dat het niet lukte om tijdens onze reguliere afspraak tot een oplossing van een probleem te komen. Als ik er dan nog een tijdje op ging zitten broeden en met een nieuw idee kwam, had jij meestal wel tijd om dat meteen even te bespreken. Ik vond het leuk om samen te puzzelen en om te horen wat jouw ideeën zijn!

Ten tweede wil ik graag Joyce bedanken. Ik vond het leuk om samen aan het Airbnbproject te werken (ik hoop dat daar nog een vervolg op komt!) en om onze eigen deelprojecten binnen OCTSE te bespreken. Je hebt regelmatig meegedacht over experimentele designs en je hebt me ook erg geholpen door

af en toe feedback te geven op mijn artikelen. Bovendien vond ik het gezellig om gewoon even in je kantoor te komen kletsen, een rondje te gaan lopen en om naar conferenties te gaan.

Natuurlijk wil ik ook graag de andere PhD studenten en collega's van de afdeling Sociologie bedanken, in het bijzonder Eva, Siyang, Mathijs, Jannes, Joris, Marcus, Nikki, Leonie, Lex, Müge en Jelle. Toen ik begon aan mijn PhD kende ik nog niet zoveel mensen in Utrecht, maar al in de eerste maand zat de sfeer er goed in. Al die koffiepauzes, wandelrondjes, feestjes, weekenden weg, pinpassenbingo, zwempartijen en borrels hebben van die vier jaar een gezellige boel gemaakt!

Eva, naast al die gezelligheid hebben we ook veel gepraat over onze onderzoeken. Ik heb veel gehad aan je kennis over o.a. statistiek en aan je inzichten in de onderwerpen van ons onderzoek. Ik wist al snel dat jij mijn paranimf zou zijn! Merle, als mijn oudste vriendin ben jij natuurlijk mijn andere paranimf. Ik wil je bedanken voor je hulp bij het testen van experimenten en het geven van feedback op stukken tekst, maar vooral voor de vakanties, atletiekwedstrijden, feestjes, wandelingen en etentjes.

Een andere persoon die ik graag wil bedanken is Martijn Arets. Jouw hulp bij het leggen van contact met verschillende platforms heeft ervoor gezorgd dat we toegang hebben gekregen tot de data van twee platforms. Zonder die data had ik het tweede hoofdstuk van mijn proefschrift, over discriminatie op een motordeelplatform, nooit kunnen schrijven.

Ik wil graag student-assistenten Jaap van Slageren en Marissa Bultman graag bedanken voor hun hulp bij de verschillende projecten, zoals het zoeken naar literatuur en het uitvoeren van het labexperiment (hoofdstuk 5). Daarnaast wil ik graag Sandor Spruit bedanken voor zijn hulp bij het opzetten van de server voor het draaien van het experiment in hoofdstuk 4. Jacqueline Tenkink-de Jong en Dennis Hofman wil ik graag bedanken bij hun hulp van het vastleggen van de samenwerking met verschillende platforms. Ten slotte wil ik Jeroen Weesie en Eva Vriens graag bedanken omdat ze met mee hebben laten doen met de statistiekcursus die zij geven, dat heeft erg geholpen!

Roel, bedankt voor het kaften van mijn proefschrift!

I would also like to say thanks to all the other employees of ICS and the Sociology department, especially to the members of the Cooperative Relations group for their feedback on several research designs and papers. Your feedback greatly helped to improve this dissertation. Another group of people I would like to thank are the members of the Sharing Economy Reading Group. I have enjoyed reading and discussing research papers from different fields. Special thanks to Bilgehan Uzunca and Ivanka Visnjic for inviting me to work on your project about Uber. It was very interesting to work with you!

Thanks to Karen Cook, Robb Willer and Paolo Parigi for giving me the opportunity to come to Stanford during the third year of my PhD. I have enjoyed the conversations we have had about my PhD project in general and about the experiment about inequality in particular. Karen, I appreciate that you took the time to meet me so regularly. These meetings gave me the opportunity to pitch issues I encountered while setting up the experiment and to get immediate feedback. Robb, I would like to thank you for convincing me that I could answer the research question by using an interactive experiment rather than a vignette study. I think that advice greatly improved the research design. I would also like to thank Karen and Robb for introducing me to a lot of other people at the department and for suggesting seminars and workshops that I could visit. I enjoyed talking to all these interesting people, and I have learned a lot! Finally, I would like to thank Paolo for the cooperation on the Airbnb project. It was nice to meet you in person after all the Skype talks!

I would also like to say thanks to the members of the Circular Economy Lab of SMO Promovendi. I enjoyed our monthly meetings and it was great to see how the plans we made in the beginning of the year resulted in a report and a very nice Circular Minds Conference a year later.

I would like to thank Michel Handgraaf and David Hardisty for motivating me to apply for a PhD position. You gave me the opportunity to visit UBC for some months during my masters and that really sparked my interest in doing research.

Met het einde van mijn PhD komt er (voorlopig) ook een eind aan mijn tijd in Utrecht. Ik wil oude en nieuwe vrienden hier bedanken voor de leuke tijd. Ook Mustapha, Miriam, Egbert en al mijn trainingsmaatjes wil ik graag bedanken voor alle uren op de atletiekbaan. Mijn ouders en Marleen wil ik graag bedanken

omdat jullie het altijd leuk vonden om te horen waar ik mee bezig was de afgelopen jaren. Ten slotte wil ik Douwe natuurlijk bedanken, omdat je zin had om mee te verhuizen naar Utrecht en omdat je er nu ook samen een leuke tijd van wilt maken in Berlijn!

ABOUT THE AUTHOR

Judith Kas was born in Wageningen, the Netherlands, on September 22, 1992. In 2013, she obtained her bachelor's degree in Management, Economics and Consumer Studies at Wageningen University (*with distinction)*. She obtained her master's degree in Management, Economics and Consumer Studies at Wageningen University in 2016. In 2015, she visited Sauder School of Business at the University of British Colombia, hosted by Dr. David Hardisty. In 2016, she started working as a Ph.D. candidate at the Interuniversity Centre for Social Science Theory and Methodology (ICS) at the department of Sociology at Utrecht University. She wrote her dissertation under supervision of Rense Corten and Arnout van de Rijt. As part of her training, she visited Prof. Karen Cook and Prof. Robb Willer at Stanford University. She currently works at Berlin Social Science Center (WZB). Her research interests include experiments, trust and discrimination.

ICS
DISSERTATION
SERIES

The ICS series presents dissertations of the Interuniversity Center for Social Science Theory and Methodology. Each of these studies aims at integrating explicit theory formation with state of the art empirical research or at the development of advanced methods for empirical research. The ICS was founded in 1986 as a cooperative effort of the universities of Groningen and Utrecht. Since 1992, the ICS expanded to the University of Nijmegen and since 2017 to the University of Amsterdam (UvA). Most of the projects are financed by the participating universities or by the Netherlands Organization for Scientific Research (NWO). The international composition of the ICS graduate students is mirrored in the increasing international orientation of the projects and thus of the ICS series itself.

1. Kees van Liere (1990), *"Lastige Leerlingen. Een empirisch onderzoek naar sociale oorzaken van probleemgedrag op basisscholen."* Amsterdam: Thesis Publishers
2. Marco van Leeuwen (1990), *"Bijstand in Amsterdam, ca. 1800 - 1850. Armenzorg als beheersings en overlevingsstrategie."* ICS-dissertation, Utrecht
3. Ineke Maas (1990), *"Deelname aan podiumkunsten via de podia, de media en actieve beoefening. Substitutie of leereffecten?"* Amsterdam: Thesis Publishers
4. Marjolein Broese van Groenou (1991), *"Gescheiden Netwerken. De relaties met vrienden en verwanten na echtscheiding"* Amsterdam: Thesis Publishers
5. Jan van den Bos (1991), *"Dutch EC Policy Making. A Model Guided Approach to Coordination and Negotiation."* Amsterdam: Thesis Publishers
6. Karin Sanders (1991), *"Vrouwelijke Pioniers. Vrouwen en mannen met een 'mannelijke' hogere beroepsopleiding aan het begin van hun loopbaan."* Amsterdam: Thesis Publishers
7. Sjerp de Vries (1991), *"Egoism, Altruism, and Social Justice. Theory and Experiments on Cooperation in Social Dilemmas."* Amsterdam: Thesis Publishers
8. Ronald Batenburg (1991), *"Automatisering in bedrijf."* Amsterdam: Thesis Publishers
9. Rudi Wielers (1991), *"Selectie en allocatie op de arbeidsmarkt. Een uitwerking voor de informele en geïnstitutionaliseerde kinderopvang."* Amsterdam: Thesis Publishers
10. Gert Westert (1991), *"Verschillen in ziekenhuisgebruik."* ICS-dissertation, Groningen

11. Hanneke Hermsen (1992), *"Votes and Policy Preferences. Equilibria in Party Systems."* Amsterdam: Thesis Publishers

12. Cora Maas (1992), *"Probleemleerlingen in het basisonderwijs"* Amsterdam: Thesis Publis¬hers

13. Ed Boxman (1992), *"Contacten en carrière. Een empirisch theoretisch onderzoek naar de relatie tussen sociale netwerken en arbeidsmarktposities"* Amsterdam: Thesis Publishers

14. Conny Taes (1992), *"Kijken naar banen. Een onderzoek naar de inschatting van arbeidsmarktkansen bij schoolverlaters uit het middelbaar beroepsonderwijs."* Amsterdam: Thesis Publishers

15. Peter van Roozendaal (1992), *"Cabinets in Multi Party Democracies. The Effect of Dominant and Central Parties on Cabinet Composition and Durability."* Amsterdam: Thesis Publishers

16. Marcel van Dam (1992), *"Regio zonder regie. Verschillen in en effectiviteit van gemeentelijk arbeidsmarktbeleid."* Amsterdam: Thesis Publishers

17. Tanja van der Lippe (1993), *"Arbeidsverdeling tussen mannen en vrouwen."* Amsterdam: Thesis Publishers

18. Marc Jacobs (1993), *"Software: Kopen of Kopiëren? Een sociaal wetenschappelijk onderzoek onder PC gebruikers."* Amsterdam: Thesis Publishers

19. Peter van der Meer (1993), *"Verdringing op de Nederlandse arbeidsmarkt. Sector- en sekseverschillen."* Amsterdam: Thesis Publishers

20. Gerbert Kraaykamp (1993), *"Over lezen gesproken. Een studie naar sociale differentiatie in leesgedrag."* Amsterdam: Thesis Publishers

21. Evelien Zeggelink (1993), *"Strangers into Friends. The Evolution of Friendship Networks Using an Individual Oriented Modeling Approach."* Amsterdam: Thesis Publishers

22. Jaco Berveling (1994), *"Het stempel op de besluitvorming. Macht, invloed en besluitvorming op twee Amsterdamse beleidsterreinen."* Amsterdam: Thesis Publishers

23. Wim Bernasco (1994), *"Coupled Careers. The Effects of Spouse's Resources on Success at Work."* Amsterdam: Thesis Publishers

24. Liset van Dijk (1994), *"Choices in Child Care. The Distribution of Child Care Among Mothers, Fathers and Non Parental Care Providers."* Amsterdam: Thesis Publishers

25. Jos de Haan (1994), *"Research Groups in Dutch Sociology."* Amsterdam: Thesis Publishers

26. Kwasi Boahene (1995), *"Innovation Adoption as a Socio Economic Process. The Case of the Ghanaian Cocoa Industry."* Amsterdam: Thesis Publishers

27. Paul Ligthart (1995), *"Solidarity in Economic Transactions. An Experimental Study of Framing Effects in Bargaining and Contracting."* Amsterdam: Thesis Publishers

28. Roger Leenders (1995), *"Structure and Influence. Statistical Models for the Dynamics of Actor Attributes, Network Structure, and their Interdependence."* Amsterdam: Thesis Publishers

29. Beate Volker (1995), *"Should Auld Acquaintance Be Forgot...? Institutions of Communism, the Transition to Capitalism and Personal Networks: the Case of East Germany."* Amsterdam: Thesis Publishers

30. Anneke Cancrinus-Matthijsse (1995), *"Tussen hulpverlening en ondernemerschap. Beroepsuitoefening en taakopvattingen van openbare apothekers in een aantal West Europese landen."* Amsterdam: Thesis Publishers

31. Nardi Steverink (1996), *"Zo lang mogelijk zelfstandig. Naar een verklaring van verschillen in oriëntatie ten aanzien van opname in een verzorgingstehuis onder fysiek kwetsbare ouderen."* Amsterdam: Thesis Publishers

32. Ellen Lindeman (1996), *"Participatie in vrijwilligerswerk."* Amsterdam: Thesis Publishers

33. Chris Snijders (1996), *"Trust and Commitments."* Amsterdam: Thesis Publishers

34. Koos Postma (1996), *"Changing Prejudice in Hungary. A Study on the Collapse of State Socialism and Its Impact on Prejudice Against Gypsies and Jews."* Amsterdam: Thesis Publishers

35. Jooske van Busschbach (1996), *"Uit het oog, uit het hart? Stabiliteit en verandering in persoonlijke relaties."* Amsterdam: Thesis Publishers

36. René Torenvlied (1996), *"Besluiten in uitvoering. Theorieën over beleidsuitvoering modelmatig getoetst op sociale vernieuwing in drie gemeenten."* Amsterdam: Thesis Publishers

37. Andreas Flache (1996), *"The Double Edge of Networks. An Analysis of the Effect of Informal Networks on Cooperation in Social Dilemmas."* Amsterdam: Thesis Publishers

38. Kees van Veen (1997), *"Inside an Internal Labor Market: Formal Rules, Flexibility and Career Lines in a Dutch Manufacturing Company."* Amsterdam: Thesis Publishers

39. Lucienne van Eijk (1997), *"Activity and Well being in the Elderly."* Amsterdam: Thesis Publishers

40. Róbert Gál (1997), *"Unreliability. Contract Discipline and Contract Governance under Economic Transition."* Amsterdam: Thesis Publishers

41. Anne-Geerte van de Goor (1997), *"Effects of Regulation on Disability Duration."* ICS-dissertation, Utrecht

42. Boris Blumberg (1997), *"Das Management von Technologiekooperationen. Partnersuche und Verhandlungen mit dem Partner aus Empirisch Theoretischer Perspektive."* ICS-dissertation, Utrecht

43. Marijke von Bergh (1997), *"Loopbanen van oudere werknemers."* Amsterdam: Thesis Publishers

44. Anna Petra Nieboer (1997), *"Life Events and Well Being: A Prospective Study on Changes in Well Being of Elderly People Due to a Serious Illness Event or Death of the Spouse."* Amsterdam: Thesis Publishers

45. Jacques Niehof (1997), *"Resources and Social Reproduction: The Effects of Cultural and Material Resources on Educational and Occupational Careers in Industrial Nations at the End of the Twentieth Century."* ICS-dissertation, Nijmegen

46. Ariana Need (1997), *"The Kindred Vote. Individual and Family Effects of Social Class and Religion on Electoral Change in the Netherlands, 1956 1994."* ICS-dissertation, Nijmegen

47. Jim Allen (1997), *"Sector Composition and the Effect of Education on Wages: an International Comparison."* Amsterdam: Thesis Publishers

48. Jack Hutten (1998), *"Workload and Provision of Care in General Practice. An Empirical Study of the Relation Between Workload of Dutch General Practitioners and the Content and Quality of their Care."* ICS-dissertation, Utrecht

49. Per Kropp (1998), *"Berufserfolg im Transformationsprozeß. Eine theoretisch empirische Studie über die Gewinner und Verlierer der Wende in Ostdeutschland."* ICS-dissertation, Utrecht

50. Maarten Wolbers (1998), *"Diploma-inflatie en verdringing op de arbeidsmarkt. Een studie naar ontwikkelingen in de opbrengsten van diploma's in Nederland."* ICS-dissertation, Nijmegen

51. Wilma Smeenk (1998), *"Opportunity and Marriage. The Impact of Individual Resources and Marriage Market Structure on First Marriage Timing and Partner Choice in the Netherlands."* ICS-dissertation, Nijmegen

52. Marinus Spreen (1999), *"Sampling Personal Network Structures: Statistical Inference in Ego-Graphs."* ICS-dissertation, Groningen

53. Vincent Buskens (1999), *"Social Networks and Trust."* ICS-dissertation, Utrecht

54. Susanne Rijken (1999), *"Educational Expansion and Status Attainment. A Cross-National and Over-Time Comparison."* ICS-dissertation, Utrecht

55. Mérove Gijsberts (1999), *"The Legitimation of Inequality in State-Socialist and Market Societies, 1987-1996."* ICS-dissertation, Utrecht

56. Gerhard van de Bunt (1999), *"Friends by Choice. An Actor-Oriented Statistical Network Model for Friendship Networks Through Time."* ICS-dissertation, Groningen

57. Robert Thomson (1999), *"The Party Mandate: Election Pledges and Government Actions in the Netherlands, 1986 1998."* Amsterdam: Thela Thesis

58. Corine Baarda (1999), *"Politieke besluiten en boeren beslissingen. Het draagvlak van het mestbeleid tot 2000."* ICS-dissertation, Groningen

59. Rafael Wittek (1999), *"Interdependence and Informal Control in Organizations."* ICS-dissertation, Groningen

60. Diane Payne (1999), *"Policy Making in the European Union: an Analysis of the Impact of the Reform of the Structural Funds in Ireland."* ICS-dissertation, Groningen

61. René Veenstra (1999), *"Leerlingen Klassen Scholen. Prestaties en vorderingen van leerlingen in het voortgezet onderwijs."* Amsterdam, Thela Thesis

62. Marjolein Achterkamp (1999), *"Influence Strategies in Collective Decision Making. A Comparison of Two Models."* ICS-dissertation, Groningen

63. Peter Mühlau (2000), *"The Governance of the Employment Relation. A Relational Signaling Perspective."* ICS-dissertation, Groningen

64. Agnes Akkerman (2000), *"Verdeelde vakbeweging en stakingen. Concurrentie om leden."* ICS-dissertation, Groningen

65. Sandra van Thiel (2000), *"Quangocratization: Trends, Causes and Consequences."* ICS-dissertation, Utrecht

66. Sylvia Peacock-Korupp (2000), *"Mothers and the Process of Social Stratification."* ICS-dissertation, Utrecht

67. Rudi Turksema (2000), *"Supply of Day Care."* ICS-dissertation, Utrecht

68. Bernard Nijstad (2000), *"How the Group Affects the Mind: Effects of Communication in Idea Generating Groups."* ICS-dissertation, Utrecht

69. Inge de Wolf (2000), *"Opleidingsspecialisatie en arbeidsmarktsucces van sociale wetenschappers."* ICS-dissertation, Utrecht

70. Jan Kratzer (2001), *"Communication and Performance: An Empirical Study in Innovation Teams."* ICS-dissertation, Groningen

71. Madelon Kroneman (2001), *"Healthcare Systems and Hospital Bed Use."* ICS/NIVEL-dissertation, Utrecht

72. Herman van de Werfhorst (2001), *"Field of Study and Social Inequality. Four Types of Educational Resources in the Process of Stratification in the Netherlands."* ICS-dissertation, Nijmegen

73. Tamás Bartus (2001), *"Social Capital and Earnings Inequalities. The Role of Informal Job Search in Hungary."* ICS-dissertation Groningen

74. Hester Moerbeek (2001), *"Friends and Foes in the Occupational Career. The Influence of Sweet and Sour Social Capital on the Labour Market."* ICS-dissertation, Nijmegen

75. Marcel van Assen (2001), *"Essays on Actor Perspectives in Exchange Networks and Social Dilemmas."* ICS-dissertation, Groningen

76. Inge Sieben (2001), *"Sibling Similarities and Social Stratification. The Impact of Family Background across Countries and Cohorts."* ICS-dissertation, Nijmegen

77. Alinda van Bruggen (2001), *"Individual Production of Social Well Being. An Exploratory Study."* ICS-dissertation, Groningen

78. Marcel Coenders (2001), *"Nationalistic Attitudes and Ethnic Exclusionism in a Comparative Perspective: An Empirical Study of Attitudes Toward the Country and Ethnic Immigrants in 22 Countries."* ICS-dissertation, Nijmegen

79. Marcel Lubbers (2001), *"Exclusionistic Electorates. Extreme Right Wing Voting in Western Europe."* ICS-dissertation, Nijmegen

80. Uwe Matzat (2001), *"Social Networks and Cooperation in Electronic Communities. A theoretical-empirical Analysis of Academic Communication and Internet Discussion Groups."* ICS-dissertation, Groningen

81. Jacques Janssen (2002), *"Do Opposites Attract Divorce? Dimensions of Mixed Marriage and the Risk of Divorce in the Netherlands."* ICS-dissertation, Nijmegen

82. Miranda Jansen (2002), *"Waardenoriëntaties en partnerrelaties. Een panelstudie naar wederzijdse invloeden."* ICS-dissertation, Utrecht

83. Anne-Rigt Poortman (2002), *"Socioeconomic Causes and Consequences of Divorce."* ICS-dissertation, Utrecht

84. Alexander Gattig (2002), *"Intertemporal Decision Making."* ICS-dissertation, Groningen

85. Gerrit Rooks (2000), *"Contract en Conflict: Strategisch Management van Inkooptransacties."* ICS-dissertation, Utrecht

86. Károly Takács (2002), *"Social Networks and Intergroup Conflict."* ICS-dissertation, Groningen

87. Thomas Gautschi (2002), *"Trust and Exchange, Effects of Temporal Embeddedness and Network Embeddedness on Providing and Dividing a Surplus."* ICS-dissertation, Utrecht

88. Hilde Bras (2002), *"Zeeuwse meiden. Dienen in de levensloop van vrouwen, ca. 1850 – 1950."* Aksant Academic Publishers, Amsterdam

89. Merijn Rengers (2002), *"Economic Lives of Artists. Studies into Careers and the Labour Market in the Cultural Sector."* ICS-dissertation, Utrecht

90. Annelies Kassenberg (2002), *"Wat scholieren bindt. Sociale gemeenschap in scholen."* ICS-dissertation, Groningen

91. Marc Verboord (2003), *"Moet de meester dalen of de leerling klimmen? De invloed van literatuuronderwijs en ouders op het lezen van boeken tussen 1975 en 2000."* ICS-dissertation, Utrecht

92. Marcel van Egmond (2003), *"Rain Falls on All of Us (but Some Manage to Get More Wet than Others): Political Context and Electoral Participation."* ICS-dissertation, Nijmegen

93. Justine Horgan (2003), *"High Performance Human Resource Management in Ireland and the Netherlands: Adoption and Effectiveness."* ICS-dissertation, Groningen

94. Corine Hoeben (2003), *"LETS' Be a Community. Community in Local Exchange Trading Systems."* ICS-dissertation, Groningen

95. Christian Steglich (2003), *"The Framing of Decision Situations. Automatic Goal Selection and Rational Goal Pursuit."* ICS-dissertation, Groningen

96. Johan van Wilsem (2003), *"Crime and Context. The Impact of Individual, Neighborhood, City and Country Characteristics on Victimization."* ICS-dissertation, Nijmegen

97. Christiaan Monden (2003), *"Education, Inequality and Health. The Impact of Partners and Life Course."* ICS-dissertation, Nijmegen

98. Evelyn Hello (2003), *"Educational Attainment and Ethnic Attitudes. How to Explain their Relationship."* ICS-dissertation, Nijmegen

99. Marnix Croes en Peter Tammes (2004). *Gif laten wij niet voortbestaan. Een onderzoek naar de overlevingskansen van joden in de Nederlandse gemeenten, 1940-1945.* Aksant Academic Publishers, Amsterdam.

100. Ineke Nagel (2004), *"Cultuurdeelname in de levensloop."* ICS-dissertation, Utrecht

101. Marieke van der Wal (2004), *"Competencies to Participate in Life. Measurement and the Impact of School."* ICS-dissertation, Groningen

102. Vivian Meertens (2004), *"Depressive Symptoms in the General Population: a Multifactorial Social Approach."* ICS-dissertation, Nijmegen

103. Hanneke Schuurmans (2004), *"Promoting Well-Being in Frail Elderly People. Theory and Intervention."* ICS-dissertation, Groningen

104. Javier Arregui Moreno (2004), *"Negotiation in Legislative Decision-Making in the European Union."* ICS-dissertation, Groningen

105. Tamar Fischer (2004), *"Parental Divorce, Conflict and Resources. The Effects on Children's Behaviour Problems, Socioeconomic Attainment, and Transitions in the Demographic Career."* ICS-dissertation, Nijmegen

106. René Bekkers (2004), *"Giving and Volunteering in the Netherlands: Sociological and Psychological Perspectives."* ICS-dissertation, Utrecht

107. Renée van der Hulst (2004), *"Gender Differences in Workplace Authority: An Empirical Study on Social Networks."* ICS-dissertation, Groningen

108. Rita Smaniotto (2004), *"'You Scratch My Back and I Scratch Yours' Versus 'Love Thy Neighbour'. Two Proximate Mechanisms of Reciprocal Altruism."* ICS-dissertation, Groningen

109. Maurice Gesthuizen (2004), *"The Life-Course of the Low-Educated in the Netherlands: Social and Economic Risks."* ICS-dissertation, Nijmegen

110. Carlijne Philips (2005), *"Vakantiegemeenschappen. Kwalitatief en Kwantitatief Onderzoek naar Gelegenheid en Refreshergemeenschap tijdens de Vakantie."* ICS-dissertation, Groningen

111. Esther de Ruijter (2005), *"Household Outsourcing."* ICS-dissertation, Utrecht

112. Frank van Tubergen (2005), *"The Integration of Immigrants in Cross-National Perspective: Origin, Destination, and Community Effects."* ICS-dissertation, Utrecht

113. Ferry Koster (2005), *"For the Time Being. Accounting for Inconclusive Findings Concerning the Effects of Temporary Employment Relationships on Solidary Behavior of Employees."* ICS-dissertation, Groningen

114. Carolien Klein Haarhuis (2005), *"Promoting Anti-Corruption Reforms. Evaluating the Implementation of a World Bank Anti-Corruption Program in Seven African Countries (1999-2001)."* ICS-dissertation, Utrecht

115. Martin van der Gaag (2005), *"Measurement of Individual Social Capital."* ICS-dissertation, Groningen

116. Johan Hansen (2005), *"Shaping Careers of Men and Women in Organizational Contexts."* ICS-dissertation, Utrecht

117. Davide Barrera (2005), *"Trust in Embedded Settings."* ICS-dissertation, Utrecht

118. Mattijs Lambooij (2005), *"Promoting Cooperation. Studies into the Effects of Long-Term and Short-Term Rewards on Cooperation of Employees."* ICS-dissertation, Utrecht

119. Lotte Vermeij (2006), *"What's Cooking? Cultural Boundaries among Dutch Teenagers of Different Ethnic Origins in the Context of School."* ICS-dissertation, Utrecht

120. Mathilde Strating (2006), *"Facing the Challenge of Rheumatoid Arthritis. A 13-year Prospective Study among Patients and Cross-Sectional Study among Their Partners."* ICS-dissertation, Groningen

121. Jannes de Vries (2006), *"Measurement Error in Family Background Variables: The Bias in the Intergenerational Transmission of Status, Cultural Consumption, Party Preference, and Religiosity."* ICS-dissertation, Nijmegen

122. Stefan Thau (2006), *"Workplace Deviance: Four Studies on Employee Motives and Self-Regulation."* ICS-dissertation, Groningen

123. Mirjam Plantinga (2006), *"Employee Motivation and Employee Performance in Child Care. The effects of the Introduction of Market Forces on Employees in the Dutch Child-Care Sector."* ICS-dissertation, Groningen

124. Helga de Valk (2006), *"Pathways into Adulthood. A Comparative Study on Family Life Transitions among Migrant and Dutch Youth."* ICS-dissertation, Utrecht

125. Henrike Elzen (2006), *"Self-Management for Chronically Ill Older People."* ICS-Dissertation, Groningen

126. Ayse Güveli (2007), *"New Social Classes within the Service Class in the Netherlands and Britain. Adjusting the EGP Class Schema for the Technocrats and the Social and Cultural Specialists."* ICS-dissertation, Nijmegen

127. Willem-Jan Verhoeven (2007), *"Income Attainment in Post-Communist Societies."* ICS-dissertation, Utrecht

128. Marieke Voorpostel (2007), *"Sibling support: The Exchange of Help among Brothers and Sisters in the Netherlands."* ICS-dissertation, Utrecht

129. Jacob Dijkstra (2007), *"The Effects of Externalities on Partner Choice and Payoffs in Exchange Networks."* ICS-dissertation, Groningen

130. Patricia van Echtelt (2007), *"Time-Greedy Employment Relationships: Four Studies on the Time Claims of Post-Fordist Work."* ICS-dissertation, Groningen

131. Sonja Vogt (2007), *"Heterogeneity in Social Dilemmas: The Case of Social Support."* ICS-dissertation, Utrecht

132. Michael Schweinberger (2007), *"Statistical Methods for Studying the Evolution of Networks and Behavior."* ICS-dissertation, Groningen

133. István Henrik Back (2007), *"Commitment and Evolution: Connecting Emotion and Reason in Long-term Relationships."* ICS-dissertation, Groningen

134. Ruben van Gaalen (2007), *"Solidarity and Ambivalence in Parent-Child Relationships."* ICS-dissertation, Utrecht

135. Jan Reitsma (2007), *"Religiosity and Solidarity – Dimensions and Relationships Disentangled and Tested."* ICS-dissertation, Nijmegen

136. Jan Kornelis Dijkstra (2007), *"Status and Affection among (Pre)Adolescents and Their Relation with Antisocial and Prosocial Behavior."* ICS-dissertation, Groningen

137. Wouter van Gils (2007), *"Full-time Working Couples in the Netherlands. Causes and Consequences."* ICS-dissertation, Nijmegen

138. Djamila Schans (2007), *"Ethnic Diversity in Intergenerational Solidarity."* ICS-dissertation, Utrecht

139. Ruud van der Meulen (2007), *"Brug over Woelig Water: Lidmaatschap van Sportverenigingen, Vriendschappen, Kennissenkringen en Veralgemeend Vertrouwen."* ICS-dissertation, Nijmegen

140. Andrea Knecht (2008), *"Friendship Selection and Friends' Influence. Dynamics of Networks and Actor Attributes in Early Adolescence."* ICS-dissertation, Utrecht

141. Ingrid Doorten (2008), *"The Division of Unpaid Work in the Household: A Stubborn Pattern?"* ICS-dissertation, Utrecht

142. Stijn Ruiter (2008), *"Association in Context and Association as Context: Causes and Consequences of Voluntary Association Involvement."* ICS-dissertation, Nijmegen

143. Janneke Joly (2008), *"People on Our Minds: When Humanized Contexts Activate Social Norms."* ICS-dissertation, Groningen

144. Margreet Frieling (2008), *"'Joint production' als motor voor actief burgerschap in de buurt."* ICS-dissertation, Groningen

145. Ellen Verbakel (2008), *"The Partner as Resource or Restriction? Labour Market Careers of Husbands and Wives and the Consequences for Inequality Between Couples."* ICS-dissertation, Nijmegen

146. Gijs van Houten (2008), *"Beleidsuitvoering in gelaagde stelsels. De doorwerking van aanbevelingen van de Stichting van de Arbeid in het CAO-overleg."* ICS-dissertation, Utrecht

147. Eva Jaspers (2008), *"Intolerance over Time. Macro and Micro Level Questions on Attitudes Towards Euthanasia, Homosexuality and Ethnic Minorities."* ICS-dissertation, Nijmegen

148. Gijs Weijters (2008), *"Youth Delinquency in Dutch Cities and Schools: A Multilevel Approach."* ICS-dissertation, Nijmegen

149. Jessica Nooij (2009), *"The Self in Social Rejection."* ICS-dissertation, Groningen

150. Gerald Mollenhorst (2009), *"Networks in Contexts. How Meeting Opportunities Affect Personal Relationships."* ICS-dissertation, Utrecht

151. Tom van der Meer (2009), *"States of Freely Associating Citizens: Comparative Studies into the Impact of State Institutions on Social, Civic and Political Participation."* ICS-dissertation, Nijmegen

152. Manuela Vieth (2009), *"Commitments and Reciprocity in Trust Situations. Experimental Studies on Obligation, Indignation, and Self-Consistency."* ICS-dissertation, Utrecht

153. Rense Corten (2009), *"Co-evolution of Social Networks and Behavior in Social Dilemmas: Theoretical and Empirical Perspectives."* ICS-dissertation, Utrecht

154. Arieke Rijken (2009), *"Happy Families, High Fertility? Childbearing Choices in the Context of Family and Partner Relationships."* ICS-dissertation, Utrecht

155. Jochem Tolsma (2009), *"Ethnic Hostility among Ethnic Majority and Minority Groups in the Netherlands. An Investigation into the Impact of Social Mobility Experiences, the Local Living Environment and Educational Attainment on Ethnic Hostility."* ICS-dissertation, Nijmegen

156. Freek Bucx (2009), *"Linked Lives: Young Adults' Life Course and Relations With Parents."* ICS-dissertation, Utrecht

157. Philip Wotschack (2009), *"Household Governance and Time Allocation. Four studies on the combination of work and care."* ICS-dissertation, Groningen

158. Nienke Moor (2009), *"Explaining Worldwide Religious Diversity. The Relationship between Subsistence Technologies and Ideas about the Unknown in Pre-industrial and (Post-)industrial Societies."* ICS-dissertation, Nijmegen

159. Lieke ten Brummelhuis (2009), *"Family Matters at Work. Depleting and Enriching Effects of Employees' Family lives on Work Outcomes."* ICS-dissertation, Utrecht

160. Renske Keizer (2010), *"Remaining Childless. Causes and Consequences from a Life Course Perspective."* ICS-dissertation, Utrecht

161. Miranda Sentse (2010), *"Bridging Contexts: The interplay between Family, Child, and Peers in Explaining Problem Behavior in Early Adolescence."* ICS-dissertation, Groningen

162. Nicole Tieben (2010), *"Transitions, Tracks and Transformations. Social Inequality in Transitions into, through and out of Secondary Education in the Netherlands for Cohorts Born Between 1914 and 1985."* ICS-dissertation, Nijmegen

163. Birgit Pauksztat (2010), *"Speaking up in Organizations: Four Studies on Employee Voice."* ICS-dissertation, Groningen

164. Richard Zijdeman (2010), *"Status Attainment in the Netherlands, 1811-1941. Spatial and Temporal Variation Before and During Industrialization."* ICS-dissertation, Utrecht

165. Rianne Kloosterman (2010), *"Social Background and Children's Educational Careers. The Primary and Secondary Effects of Social Background over Transitions and over Time in the Netherlands."* ICS-dissertation, Nijmegen

166. Olav Aarts (2010), *"Religious Diversity and Religious Involvement. A Study of Religious Markets in Western Societies at the End of the Twentieth Century."* ICS-dissertation, Nijmegen

167. Stephanie Wiesmann (2010), *"24/7 Negotiation in Couples Transition to Parenthood."* ICS-dissertation, Utrecht

168. Borja Martinovic (2010), *"Interethnic Contacts: A Dynamic Analysis of Interaction Between Immigrants and Natives in Western Countries."* ICS-dissertation, Utrecht

169. Anne Roeters (2010), *"Family Life Under Pressure? Parents' Paid Work and the Quantity and Quality of Parent-Child and Family Time."* ICS-dissertation, Utrecht

170. Jelle Sijtsema (2010), *"Adolescent Aggressive Behavior: Status and Stimulation Goals in Relation to the Peer Context."* ICS-dissertation, Groningen

171. Kees Keizer (2010), *"The Spreading of Disorder."* ICS-dissertation, Groningen

172. Michael Mäs (2010), *"The Diversity Puzzle. Explaining Clustering and Polarization of Opinions."* ICS-dissertation, Groningen

173. Marie-Louise Damen (2010), *"Cultuurdeelname en CKV. Studies naar effecten van kunsteducatie op de cultuurdeelname van leerlingen tijdens en na het voortgezet onderwijs."* ICS-dissertation, Utrecht

174. Marieke van de Rakt (2011), *"Two generations of Crime: The Intergenerational Transmission of Convictions over the Life Course."* ICS-dissertation, Nijmegen

175. Willem Huijnk (2011), *"Family Life and Ethnic Attitudes. The Role of the Family for Attitudes Towards Intermarriage and Acculturation Among Minority and Majority Groups."* ICS-dissertation, Utrecht

176. Tim Huijts (2011), *"Social Ties and Health in Europe. Individual Associations, Cross-National Variations, and Contextual Explanations."* ICS-dissertation, Nijmegen

177. Wouter Steenbeek (2011), *"Social and Physical Disorder. How Community, Business Presence and Entrepreneurs Influence Disorder in Dutch Neighborhoods."* ICS-dissertation, Utrecht

178. Miranda Vervoort (2011), *"Living Together Apart? Ethnic Concentration in the Neighborhood and Ethnic Minorities' Social Contacts and Language Practices."* ICS-dissertation, Utrecht

179. Agnieszka Kanas (2011), *"The Economic Performance of Immigrants. The Role of Human and Social Capital."* ICS-dissertation, Utrecht

180. Lea Ellwardt (2011), *"Gossip in Organizations. A Social Network Study."* ICS-dissertation, Groningen

181. Annemarije Oosterwaal (2011), *"The Gap between Decision and Implementation. Decision making, Delegation and Compliance in Governmental and Organizational Settings."* ICS-dissertation, Utrecht

182. Natascha Notten (2011), *"Parents and the Media. Causes and Consequences of Parental Media Socialization."* ICS-dissertation, Nijmegen

183. Tobias Stark (2011), *"Integration in Schools. A Process Perspective on Students' Interethnic Attitudes and Interpersonal Relationships."* ICS-dissertation, Groningen

184. Giedo Jansen (2011), *"Social Cleavages and Political Choices. Large-scale Comparisons of Social Class, Religion and Voting Behavior in Western Democracies."* ICS-dissertation, Nijmegen

185. Ruud van der Horst (2011), *"Network Effects on Treatment Results in a Closed Forensic Psychiatric Setting."* ICS-dissertation, Groningen

186. Mark Levels (2011), *"Abortion Laws in European Countries between 1960 and 2010. Legislative Developments and Their Consequences for Women's Reproductive Decision-making."* ICS-dissertation, Nijmegen

187. Marieke van Londen (2012), *"Exclusion of ethnic minorities in the Netherlands. The effects of individual and situational characteristics on opposition to ethnic policy and ethnically mixed neighbourhoods."* ICS-dissertation, Nijmegen

188. Sigrid Mohnen (2012), *"Neighborhood context and health: How neighborhood social capital affects individual health."* ICS-dissertation, Utrecht

189. Asya Zhelyazkova (2012), *"Compliance under Controversy: Analysis of the Transposition of European Directives and their Provisions."* ICS-dissertation, Utrecht

190. Valeska Korff (2012), *"Between Cause and Control: Management in a Humanitarian Organization."* ICS-dissertation, Groningen

191. Maike Gieling (2012), *"Dealing with Diversity: adolescents' support for civil liberties and immigrant rights."* ICS-dissertation, Utrecht

192. Katya Ivanova (2012), *"From Parents to Partners: The Impact of Family on Romantic Relationships in Adolescence and Emerging Adulthood."* ICS-dissertation, Groningen

193. Jelmer Schalk (2012), *"The Performance of Public Corporate Actors: Essays on Effects of Institutional and Network Embeddedness in Supranational, National, and Local Collaborative Contexts."* ICS-dissertation, Utrecht

194. Alona Labun (2012), *"Social Networks and Informal Power in Organizations."* ICS-dissertation, Groningen

195. Michal Bojanowski (2012), *"Essays on Social Network Formation in Heterogeneous Populations: Models, Methods, and Empirical Analyses."* ICS-dissertation, Utrecht

196. Anca Minescu (2012), *"Relative Group Position and Intergroup Attitudes in Russia."* ICS-dissertation, Utrecht

197. Marieke van Schellen (2012), *"Marriage and crime over the life course. The criminal careers of convicts and their spouses."* ICS-dissertation, Utrecht

198. Mieke Maliepaard (2012), *"Religious Trends and Social Integration: Muslim Minorities in the Netherlands."* ICS-dissertation, Utrecht

199. Fransje Smits (2012), *"Turks and Moroccans in the Low Countries around the year 2000: determinants of religiosity, trend in religiosity and determinants of the trend."* ICS-dissertation, Nijmegen

200. Roderick Sluiter (2012), *"The Diffusion of Morality Policies among Western European Countries between 1960 and 2010. A Comparison of Temporal and Spatial Diffusion Patterns of Six Morality and Eleven Non-morality Policies."* ICS-dissertation, Nijmegen

201. Nicoletta Balbo (2012), *"Family, Friends and Fertility."* ICS-dissertation, Groningen

202. Anke Munniksma (2013), *"Crossing ethnic boundaries: Parental resistance to and consequences of adolescents' cross-ethnic peer relations"* ICS-dissertation, Groningen

203. Anja-Kristin Abendroth (2013), *"Working Women in Europe. How the Country, Workplace, and Family Context Matter."* ICS-dissertation, Utrecht

204. Katia Begall (2013), *"Occupational Hazard? The Relationship between Working Conditions and Fertility."* ICS-dissertation, Groningen

205. Hidde Bekhuis (2013), *"The Popularity of Domestic Cultural Products: Cross-national Differences and the Relation to Globalization."* ICS-dissertation, Utrecht

206. Lieselotte Blommaert (2013), *"Are Joris and Renske more employable than Rashid and Samira? A study on the prevalence and sources of ethnic discrimination in recruitment in the Netherlands using experimental and survey data."* ICS-dissertation, Utrecht

207. Wiebke Schulz (2013), *"Careers of Men and Women in the 19th and 20th Centuries."* ICS-dissertation, Utrecht

208. Ozan Aksoy (2013), *"Essays on Social Preferences and Beliefs in Non-embedded Social Dilemmas."* ICS-dissertation, Utrecht

209. Dominik Morbitzer (2013), *"Limited Farsightedness in Network Formation."* ICS-dissertation, Utrecht

210. Thomas de Vroome (2013), *"Earning Your Place: The Relation Between Immigrants' Economic and Psychological Integration in the Netherlands."* ICS-dissertation, Utrecht

211. Marloes de Lange (2013), *"Causes and Consequences of Employment Flexibility Among Young People. Recent Developments in the Netherlands and Europe."* ICS-dissertation, Nijmegen

212. Roza Meuleman (2014), *"Consuming the Nation. Domestic cultural consumption: its stratification and relation with nationalist attitudes."* ICS-dissertation, Utrecht

213. Esther Havekes (2014), *"Putting interethnic attitudes in context. The relationship between neighbourhood characteristics, interethnic attitudes and residential behaviour."* ICS-dissertation, Utrecht

214. Zoltán Lippényi (2014), *"Transitions toward an open society? Intergenerational occupational mobility in Hungary in the 19th and 20th centuries."* ICS-dissertation, Utrecht

215. Anouk Smeekes (2014), *"The presence of the past: Historical rooting of national identity and current group dynamics."* ICS-dissertation, Utrecht

216. Michael Savelkoul (2014), *"Ethnic Diversity and Social capital. Testing Underlying Explanations Derived from Conflict and Contact Theories in Europe and the United States."* ICS-dissertation, Nijmegen

217. Martijn Hogerbrugge (2014), *"Misfortune and family: How negative events, family ties, and lives are linked."* ICS-dissertation, Utrecht

218. Gina-Felicia Potarca (2014), *"Modern Love. Comparative Insights in Online Dating Preferences and Assortative Mating."* ICS dissertation, Groningen

219. Mariska van der Horst (2014), *"Gender, Aspirations, and Achievements: Relating work and family aspirations to occupational outcomes."* ICS-dissertation, Utrecht

220. Gijs Huitsing (2014), *"A Social network perspective on bullying"* ICS dissertation, Groningen

221. Thomas Kowalewski (2015), *"Personal growth in organizational contexts."* ICS-dissertation, Groningen

222. Manuel Muñoz-Herrera (2015), *"The impact of individual differences on network relations: Social exclusion and inequality in productive exchange and coordination games."* ICS-dissertation, Groningen

223. Tim Immerzeel (2015), *"Voting for a change. The democratic lure of populist radical right parties in voting behavior."* ICS-dissertation, Utrecht

224. Fernando Nieto Morales (2015), *"The control imperative: Studies on reorganization in the public and private sectors."* ICS-dissertation, Groningen

225. Jellie Sierksma (2015), *"Bounded helping: How morality and intergroup relations shape children's reasoning about helping."* ICS-dissertation, Utrecht

226. Tinka Veldhuis (2015), *"Captivated by fear. An evaluation of terrorism detention policy."* ICS-dissertation, Groningen

227. Miranda Visser (2015), *"Loyality in humanity.Turnover among expatriate humanitarian aid workers."* ICS-dissertation, Groningen

228. Sarah Westphal (2015), *"Are the kids alright? Essays on postdivorce residence arrangements and children's well-being."* ICS-dissertation, Utrecht

229. Britta Rüschoff (2015), *"Peers in careers: Peer relationships in the transition from school to work."* ICS dissertation, Groningen

230. Nynke van Miltenburg (2015), *"Cooperation under peer sanctioning institutions: Collective decisions, noise, and endogenous implementation."* ICS-dissertation, Utrecht

231. Antonie Knigge (2015), *"Sources of sibling similarity. Status attainment in the Netherlands during modernization."* ICS-dissertation, Utrecht

232. Sanne Smith (2015), *"Ethnic segregation in friendship networks. Studies of its determinants in English, German, Dutch, and Swedish school classes."* ICS-dissertation, Utrecht

233. Patrick Präg (2015), *"Social stratification and health. Four essays on social determinants of health and wellbeing."* ICS dissertation, Groningen

234. Wike Been (2015), *"European top managers' support for work-life arrangements"* ICS-dissertation, Utrecht

235. André Grow (2016), *"Status differentiation: New insights from agent-based modeling and social network analysis"* ICS-dissertation, Groningen

236. Jesper Rözer (2016), *"Family and personal networks. How a partner and children affect social relationships"* ICS-dissertation, Utrecht

237. Kim Pattiselanno (2016), *"At your own risk: The importance of group dynamics and peer processes in adolescent peer groups for adolescents' involvement in risk behaviors"* ICS-dissertation, Groningen

238. Vincenz Frey (2016), *"Network formation and trust"* ICS-dissertation, Utrecht

239. Rozemarijn van der Ploeg (2016), *"Be a buddy, not a bully? Four studies on social and emotional processes related to bullying, defending, and victimization"* ICS-dissertation, Groningen

240. Tali Spiegel (2016), *"Identity, career trajectories and wellbeing: A closer look at individuals with degenerative eye conditions"* ICS- dissertation, Groningen

241. Felix Christian Tropf (2016), *"Social Science Genetics and Fertility"* ICS-dissertation, Groningen

242. Sara Geven (2016), *"Adolescent problem behavior in school: the role of peer networks"* ICS-dissertation, Utrecht

243. Josja Rokven (2016), *"The victimization-offending relationship from a longitudinal perspective"* ICS-dissertation, Nijmegen

244. Maja Djundeva (2016), *"Healthy ageing in context: Family welfare state and the life course"* ICS-dissertation, Groningen

245. Mark Visser (2017), *"Inequality between older workers and older couples in the Netherlands. A dynamic life course perspective on educational and social class differences in the late career"* ICS-dissertation, Nijmegen

246. Beau Oldenburg (2017), *"Bullying in schools: The role of teachers and classmates"* ICS-dissertation, Groningen

247. Tatang Muttaqin (2017), *"The education divide in Indonesia: Four essays on determinants of unequal access to and quality of education"* ICS-dissertation, Groningen

248. Margriet van Hek (2017), *"Gender inequality in educational attainment and reading performance. A contextual approach."* ICS-dissertation, Nijmegen

249. Melissa Verhoef-van Dorp (2017), *"Work schedules, childcare and well-being. Essays on the associations between modern-day job characteristics, childcare arrangements and the well-being of parents and children"* ICS-dissertation, Utrecht

250. Timo Septer (2017), *"Goal priorities, cognition and conflict: Analyses of cognitive maps concerning organizational change"* ICS-dissertation, Groningen

251. Bas Hofstra (2017), *"Online Social Networks: Essays on Membership, Privacy, and Structure"* ICS-dissertation, Utrecht

252. Yassine Khoudja (2018), *"Women's labor market participation across ethnic groups: The role of household conditions, gender role attitudes, and religiosity in different national contexts"* ICS-dissertation, Utrecht

253. Joran Laméris (2018), *"Living together in diversity. Whether, why and where ethnic diversity affects social cohesion"* ICS-dissertation, Nijmegen

254. Maaike van der Vleuten (2018), *"Gendered Choices. Fields of study of adolescents in the Netherlands"* ICS-dissertation, Utrecht

255. Mala Sondang Silitonga (2018), *"Corruption in Indonesia: The impact of institutional change, norms, and networks"* ICS-dissertation, Groningen

256. Manja Coopmans (2018), *"Rituals of the past in the context of the present. The role of Remembrance Day and Liberation Day in Dutch society"* ICS-dissertation, Utrecht

257. Paul Hindriks (2018), *"The Struggle for Power: Attitudes towards the political participation of ethnic minorities"* ICS-dissertation, Utrecht

258. Nynke Niezink (2018), *"Modeling the dynamics of networks and continuous behavior"* ICS-dissertation, Groningen

259. Simon de Bruijn (2018), *"Reaching agreement after divorce and separation. Essays on the effectiveness of parenting plans and divorce mediation"* ICS-dissertation, Utrecht

260. Susanne van 't Hoff-de Goede (2018), *"While you were locked up. An empirical study on the characteristics, social surroundings and wellbeing of partners of prisoners in The Netherlands"* ICS-dissertation, Utrecht

261. Loes van Rijsewijk (2018), *"Antecedents and Consequences of Helping among Adolescents"* ICS-dissertation, Groningen

262. Mariola Gremmen (2018), *"Social network processes and academic functioning. The role of peers in students' school well-being, academic engagement, and academic achievement"* ICS-dissertation, Groningen

263. Jeanette Renema (2018), *"Immigrants' support for welfare spending. The causes and consequences of welfare usage and welfare knowledgeability"* ICS-dissertation, Nijmegen

264. Suwatin Miharti (2018), *"Community health centers in Indonesia in the era of decentralization. The impact of structure, staff composition and management on health outcomes"* ICS-dissertation, Groningen

265. Chaïm la Roi (2019), *"Stigma and stress: Studies on attitudes towards sexual minority orientations and the association between sexual orientation and mental health"* ICS-dissertation, Groningen

266. Jelle Lössbroek (2019), *"Turning grey into gold. Employer-employee interplay in an ageing workforce"* ICS-dissertation, Utrecht

267. Nikki van Gerwen (2019), *"Employee cooperation through training. A multi-method approach"* ICS-dissertation, Utrecht

268. Paula Thijs (2019), *"Trends in cultural conservatism: The role of educational expansion, secularisation, and changing national contexts"* ICS-dissertation, Nijmegen

269. Renske Verweij (2019), *"Understanding childlessness: Unravelling the link with genes and the socio-environment"* ICS-dissertation, Groningen

270. Niels Blom (2019), *"Partner relationship quality under pressing work conditions. Longitudinal and cross-national investigation"* ICS-dissertation, Nijmegen

271. Müge Simsek (2019), *"The dynamics of religion among native and immigrant youth in Western Europe"* ICS-dissertation, Utrecht

272. Leonie van Breeschoten (2019), *"Combining a career and childcare: The use and usefulness of work-family policies in European organizations"* ICS-dissertation, Utrecht

273. Roos van der Zwan (2019), "*The political representation of ethnic minorities and their vote choice"* ICS-dissertation, Nijmegen

274. Ashwin Rambaran (2019), "*The classroom as context for bullying: A social network approach*" ICS-dissertation, Groningen

275. Dieko Bakker (2019), *"Cooperation and social control: Effects of preferences, institutions, and social structure"* ICS-dissertation, Groningen

276. Femke van der Werf (2019), "*Shadow of a rainbow? National and ethnic belonging in Mauritius*" ICS-dissertation, Utrecht

277. Robert Krause (2019), *"Multiple Imputation for Missing Network Data"* ICS-dissertation, Groningen

278. Take Sipma (2020), *"Economic insecurity and populist radical right voting"* ICS-dissertation, Nijmegen

279. Mathijs Kros (2020), *"The nature of negative contact: studies on interethnic relations in western societies"* ICS-dissertation, Utrecht

280. Lonneke van den Berg (2020), *"Time to leave: Individual and contextual explanations for the timing of leaving home"* ICS-dissertation, Amsterdam

281. Marianne Hooijsma (2020), "*Clashrooms: Interethnic peer relationships in schools"* ICS-dissertation, Groningen

282. Marina Tulin (2020*), "Blind Spots in Social Resource Theory: Essays on the Creation, Maintenance and Returns of Social Capital"* ICS-dissertation, Amsterdam

283. Tessa Kaufman (2020), "*Toward Tailored Interventions: Explaining, Assessing, and Preventing Persistent Victimization of Bullying"* ICS-dissertation, Groningen

284. Lex Thijssen (2020), *"Racial and Ethnic Discrimination in Western Labor Markets: Empirical Evidence from Field Experiments"* ICS-dissertation, Utrecht

285. Lukas Norbutas (2020), *"Trust on the Dark Web: An analysis of illegal online drug markets"* ICS-dissertation, Utrecht

286. Tomáš Diviák (2020), *"Criminal Networks: Actors, Mechanisms, and Structures"* ICS-dissertation, Groningen

287. Tery Setiawan (2020), "*Support for Interreligious Conflict in Indonesia*" ICS-dissertation, Nijmegen

288. Vera de Bel (2020), "*The ripple effect in family networks. Relational structures and well-being in divorced and non-divorced families*" ICS-dissertation, Groningen

289. Diego Palacios (2020), *"How Context and the Perception of Peers' Behaviors Shape Relationships in Adolescence: A Multiplex Social Network Perspective"* ICS-dissertation, Groningen

290. Saskia Glas (2020), *"Where are the Muslim Feminists?: Religiosity and Support for Gender Equality in the Arab region*" ICS-dissertation, Nijmegen

291. Tomas Turner-Zwinkels (2020), *"A New Macro-Micro Approach to the Study of Political Careers: Theoretical, Methodological and Empirical Challenges and Solutions"* ICS-dissertation, Groningen

292. Lotte Scheeren (2020), "*Not on the Same Track? Tracking Age and Gender Inequality in Education"* ICS-dissertation, Amsterdam

293. Joris Broere (2020), "*Essays on how social network structure affects asymmetric coordination and trust*" ICS-dissertation, Utrecht

294. Marcus Kristiansen (2021*), "Contact with Benefits: How Social Networks affect Benefit Receipt Dynamics in the Netherlands"* ICS-dissertation, Utrecht

295. Judith Kas (2021), "*Trust and reputation in the peer-to-peer platform economy*" ICS-dissertation, Utrecht

People are more likely to interact with another individual if they know that person has been trustworthy in the past. Reputation systems, by providing that information on past trustworthiness, thus enable trust. This dissertation studies three different mechanisms through which reputation systems affect behavior in the context of the platform economy.

The first issue that is addressed, is whether reputation systems help solve discrimination. While reputation information may on the one hand draw the attention of trustors away from the demographic characteristics of the trustee, reputation systems may reproduce existing differences in the probability to be trusted. Another question that is being answered is to what extent other types of reputation information affect the decision to place trust or not. Contrary to what previous research has focused on, reputation systems may not only convey information about the trustee's past trustworthiness, but also about their past trustfulness and past honesty. This dissertation provides an answer to these questions by using a combination of field data (generated by users of real platforms), laboratory experiments and online experiments.

Judith Kas obtained her Master's degree in Management, Economics and Consumer Studies at Wageningen University. She conducted the present study as part of her Ph.D. research at the Interuniversity Center for Social Science Theory and Methodology (ICS) and the department of Sociology of Utrecht University.