# How to design trust on market platforms?*

Ben Greiner, Wirtschaftsuniversität Wien & University of New South Wales

Timm Teubner, Technische Universität Berlin

Christof Weinhardt, Karlsruher Institut für Technologie

Abstract:
Reputation systems are core components of online platform markets, incentivizing trust and trustworthiness. Small market design details may have large implications on the extent of trade and platform revenues. We selectively review our own and others' research on the design of reputation mechanisms for platforms highlighting particular issues in the design of feedback systems which may affect feedback behavior and thus the informativeness and effectiveness of feedback information. We also discuss new research on conflict resolution mechanisms for platforms and the arguments for allowing reputation mobility. Finally, we conclude with implications for platform management and users.

* Ben Greiner thanks his co-authors Gary Bolton and Axel Ockenfels for the long-term collaboration on these topics. Timm Teubner and Christof Weinhardt thank their co-authors Marc Adam and Florian Hawlitschek for their cooperation in the research cited herein. We note that a similar review of our work, but targeted towards a law audience and in German language, appeared as Greiner, Teubner, and Weinhardt (2018).

## 1. Introduction: The economic principles of market platforms

Platforms are intermediaries that bring together groups of agents with the goal to enable economic and social exchange. Examples for such intermediaries are Facebook, Twitter and other social network sites, dating websites, or Visa, MasterCard and PayPal as payment platforms. One particular type of platform are online markets, in which the platform takes over the role of a market organizer that does not only bring together buyers and sellers, but also facilitates price discovery and other transaction processes. Examples for this type of platform are the e-commerce sites eBay, Airbnb, BlaBlaCar, Amazon "Marketplace", Apple's "App Store" (where developers find customers for their software), and Google's "Ad Platform" (where websites and advertisers find each other). The financial market valuations of some of these market platforms already exceed the value of some of the biggest "old commerce" companies.[1]

Reputation systems are part of the core functionality of these market platforms. In a reputation system, information about the behavior of market participants is collected and distributed to potential future transaction partners in aggregated or disaggregated form. Since the potential transaction partners may and should condition their decision about an engagement with a market participant on the historical information available, the system creates strong incentives to behave trustworthy in the market. In this way, reputation systems influence the (expected) costs and risks of participating in a market platform, but also the (expected) intrinsic benefit. They thus also affect the attractiveness and competitive advantage of a platform. Additionally, reputation histories enable a "lock-in" of users if they are not or only in a limited way transferable to other platforms (see Section 4). Last but not least, the design of reputation systems interacts with issues of liability and quality control on platforms. On Amazon, for example, the platform closely controls payment streams and product quality, while on eBay quality and trustworthiness of the seller predominantly rely on the information collected and incentives set by the reputation system.

[1] Uber went public in May 2019 at a valuation of US$ 82 billion. eBay's market cap was US$ 33 billion in early 2019. And Airbnb was valued US$ 31 billion at its last funding in March 2017 (Forbes 2018). A recent report by the European Commission estimates that 191 million citizens across the EU-28 countries have actively engaged in peer-to-peer platforms between 2015 and 2016. Including selling and buying of goods (e.g., eBay), sharing and renting of goods (e.g., Peerby), accommodation (e.g., Airbnb), rides (e.g., Uber, BlaBlaCar), and crowd work (e.g., TaskRabbit), annual expenditures on these platforms are estimated at EUR 27.9 billion (EU 2017).

In many ways, eBay has been an industry pioneer in the design (and prior experimental and empirical testing) of reputation systems. A large part of research into reputation systems has been based on eBay data (partly scraped from the Internet, partly provided by eBay itself), and competitors as well as collaborating companies have learned from eBay's experiences and followed its lead. Many research results from eBay proved to replicate on other platforms.

The reputation system on eBay was introduced in 1996. Over the course of time, the system underwent a number of major redesigns. At the beginning of 2007 it basically worked as follows. After each transaction, both transaction partners could submit a positive, neutral, or negative feedback (or no feedback at all) on the other market participant. Additionally, the traders were able to submit a short verbal comment on the transaction, the information content of which however was usually quite limited (e.g., "A++++ Top!!!!"). Feedback could be given within a time window of about 90 days after the transaction. Any feedback submitted on the platform was made publicly available right away. Feedback received by a market participant was aggregated to a "feedback score", which represented the number of received positive feedbacks (from unique transaction partners) minus the number of received negative feedbacks (from unique transaction partners). The score considered both feedback received as a buyer and received as a seller. A second prominent aggregate feedback indicator was the "percentage positive", that is, the relative share of positive feedback in the feedback score (neutral feedback was ignored).

A different variant of reputation system commonly found on market platforms is based on ratings on a scale of one to five "stars". Here, typically the average received star rating or even the whole distribution of received ratings is displayed to future transaction partners. Platforms using such a system include Amazon, Uber, virtually all popular C2C platforms (Teubner and Dann, 2018), and since 2007 also eBay with its so-called "detailed seller ratings".

Empirical studies show that eBay's feedback system works pretty well (even though not perfectly). In 2007, it contained more than 6 billion individual feedbacks, and 4 million new feedbacks arrived every single day. About 70% of the market participants gave feedback after a transaction, buyers and sellers about equally often. Many academic studies found a positive relationship between a seller's feedback indicators ("percentage positive", feedback score) and the likelihood of a sale as well as the sale price (e.g. Bajari and Hortaçsu (2004), Cabral

and Hortaçsu (2010), Eaton (2007), Houser and Wooders (2005), Jin and Kato (2006), Livingston (2005), Lucking-Reiley, Bryan, Prasad, and Reeves (2007), Resnick, Zeckhauser, Swanson, and Lockwood (2006); see Dellarocas 2004 for a comprehensive early overview). Similar evidence exists for Airbnb and other platforms, suggesting that a good reputation on a market platform has a measurable economic impact. Statistical analyses of prices on Airbnb show that hosts with better reputation scores are able to charge higher prices for their apartments (Teubner, Hawlitschek, and Dann 2017). At the same time, surveys with Airbnb users show that they are more willing to book with hosts with very positive ratings as compared to hosts with less positive ratings.

Additional evidence comes from laboratory experiments. These studies show that trust and cooperation between sellers and buyers may be very fragile and usually breaks down in an environment where no information about past behavior is available. The existence of a reputation system, on the other hand, stabilizes market cooperation on a robust level. For example, Bolton, Katok and Ockenfels (2004) model a buyer-seller-game, in which the buyer first decides about sending his money or not, and after receiving the money the seller decides whether to send the product (in the promised quality) or not. While trade yields efficiency gains, such that both transaction partners may be better off, the seller has incentives not to send the product (in the promised quality) after receiving the money, such that the buyer has no reason to send money to the seller in the first place. In their laboratory experiment, participants play this game for 30 rounds, with a new transaction partner in each round. While in a baseline ("stranger") setting the market participants did not receive any information about their transaction partner's behavior in previous rounds, in a "reputation" treatment the traders were informed about their current partner's complete decision history in previous rounds. Bolton, Katok and Ockenfels (2004) observe that while the frequency of efficient transactions quickly converges to zero when no reputation system is in place, the number of efficient transactions stabilizes at about 50% of maximum efficiency gains when information about past behavior, that is a reputation system, is provided.

An important caveat of reputation systems in the real world is that reputation information is not exogenously and objectively produced, but is endogenously and subjectively provided to the platform by the transaction partners involved. In addition, it is not easily verifiable, as the

observation of detailed payment streams and product qualities comes with significant costs. Thus, reputation information might not be fully informative, and sometimes even biased.

The important question for economic theory and market design practice arising in face of this challenge is how an existing reputation system can be optimized such that it provides incentives to submit honest, unbiased, and reliable information, which in turn creates incentives for market participants to behave trustworthy and realize transaction efficiency gains. These issues should be of core interest for platform providers, in order to increase positive network effects and make the platform attractive to more users. In the remainder of this paper, we will discuss a number of these issues based on our own research findings as well as on those of other researchers in this field.

## 2. Reputation system design: Reciprocity and information

In 2007, the reputation system on eBay worked well, as also evidenced by the platform's success in the e-commerce sector. But it did not work perfectly, such that there was room for improvement. Bolton, Greiner, and Ockenfels (2013) document strongly reciprocal behavior in mutual feedback giving on eBay, which significantly affected the informativeness of feedback in the system and the trustworthiness of market participants. As mentioned above, in eBay's feedback system any feedback was immediately published on the platform. For example, if a buyer submitted positive feedback about the seller, then the seller was immediately informed about that feedback. In 99.9 percent of the cases in which the seller responded with a feedback, that feedback was positive as well. In the opposite case, however, when the buyer moved first and submitted a negative feedback, then the feedback with which the seller responded was negative in more than 90 percent of the cases. This establishes strong evidence for reciprocity in feedback giving. This reciprocal behavior may be socially or strategically motivated.

The existence of positive and negative reciprocity has differential effects on the incentives to submit feedback in the first place. A positive feedback for the transaction partner does not only reward their cooperation in the transaction, but also stipulates the trading partner to reply with a positive feedback of her own. A negative feedback punishes uncooperative behavior (a negative feedback has an empirically validated negative effect on future sale

prices of a trader), but may result in retaliation. Thus, reciprocity in feedback interactions increases the incentive to submit positive feedback and lowers the incentive to submit negative feedback. The eventual result is too much positive feedback in the reputation system with very little informational content. As a matter of fact, Bolton, Greiner, and Ockenfels (2013) report a share of 97-99% of positive feedback on eBay (with some small variation across categories and countries). The consequence may be lowered trust (both into sellers as well as into the market platform itself) and less trade.

Bolton, Greiner, and Ockenfels (2013) discuss and evaluate two possible solutions to this problem. The first approach is to make feedback "blind", that is, feedback submitted by a transaction partner is only published on the platform when the other transaction partner has submitted her feedback as well. This design would prevent any sequential reciprocity in feedback giving. However, it also comes with some caveats. A blind system would also prevent positive reciprocity, which serves as an important motivation for participating in the feedback system, in particular for reporting positive experiences. There would also still be room for "simultaneous" reciprocity, when a trader who *expects* negative feedback retaliates with a negative feedback to that expectation. A blind feedback system will need a well-defined time window (e.g. three weeks) within which feedback can be submitted and after which submitted feedback is published even when the transaction partner failed to submit her own feedback. This, however, creates room for strategic delay of feedback. A criminal seller would have three weeks time to cheat buyers on the platform before any of her behavior will be reflected in the platform's reputation system.

Empirical as well as experimental data support these theoretical considerations. The Brazilian eBay subsidy MercadoLivre uses the same auction and market mechanisms as eBay, but implements (based on historical reasons) its own blind reputation system with a three-week window to submit feedback. On MercadoLivre one could observe much more negative feedback and a significantly reduced correlation between mutual feedback, i.e. much less reciprocity in feedback giving. However, in such field data many details of the transaction (e.g. product quality) are not observable, and differences in feedback behavior may also root in cultural differences between Latin America and Western countries. This makes a strong case for laboratory experiments, where reputation system rules can be varied while holding all other market parameters as well as the trader population constant.

Bolton, Greiner, and Ockenfels (2013) corroborate the field results with causal evidence from a laboratory experiment. In their data, blind feedback yields more negative feedback with a reduced buyer-seller feedback correlation as compared to open feedback, and blindness of feedback eventually results in higher prices and higher product quality. Since reciprocity is excluded, also less feedback is given under the blind than under the open system.

A second solution approach was to make feedback one-sided, and only allowing the buyer to give feedback on the seller, but not vice versa. The argument behind this approach is that moral hazard opportunities (that is, room for misbehavior) are asymmetrically distributed between buyers and sellers, since buyers usually pay the product first and only after receiving payment the seller sends the product. Amazon's 5-star-rating is an example for a one-sided system. However, Amazon also handles all payments for transactions made on its platform, and manages complaints as well as returns. As such, Amazon has a close grip on most moral hazard opportunities of buyers, and mainly uses the reputation system to discipline sellers and incentivize their product qualities.

A one-sided feedback system on eBay would exclude feedback retaliation by definition. However, it would also exclude positive reciprocity, and thus lower incentives to participate in the reputation system. An open question for eBay is whether buyers indeed have no significant moral hazard opportunities. Buyers may not pay the auction price after they have won (which incurs costs for the seller to relist the item), may withdraw their credit card payment after receiving the product, or may submit unsubstantiated complaints about the product. A one-sided feedback system also offers opportunities for blackmailing of sellers, since it takes away a means to defend themselves. Based on these considerations, in 2007 eBay considered a one-sided system only *in addition* to the existing open two-sided system, as an additional, one-sided possibility for the buyer to rate the seller, anonymously and thus not subject to retaliatory actions. As a scale eBay used the 5-star-rating well known from Amazon and other platforms In the laboratory experiments of Bolton, Greiner, and Ockenfels (2013), this hybrid feedback system design performed remarkably well. Both prices as well as product quality (that is, both trust and trustworthiness) were significantly increased, even more than with a blind system. The hybrid system yielded more informativeness in the collected feedback ratings, but in contrast to the blind system this did not come at the cost of a reduction in feedback system participation.

From March to May 2007, eBay tested the new extended feedback system under the name of “Feedback 2.0 / detailed seller ratings” in a number of smaller and medium-sized eBay markets such as Australia, France, and the United Kingdom. After the successful testing phase, the new system was introduced world-wide, including the largest eBay platforms in Germany and the United States. Bolton, Greiner, and Ockenfels (2013) collected field data after the introduction of the new system and show that the new system indeed leads to less reciprocity in feedback giving and more information content in feedback. Nevertheless, the majority of feedback collected on eBay remained positive.

Also Airbnb revised its reputation system over the years. As Fradkin, Grewal, and Holtz (2017) report and evaluate, inspired by the research reported above in summer 2014 Airbnb’s system was migrated from an open system to a blind system, such that since then users have to submit their feedback without knowing the feedback giving by their transaction partner about themselves. This change made the reputation system more informative. However, negative experiences are still underreported (see also Zervas, Proserpio, and Byers 2015; Dann, Teubner, and Weinhardt 2019). Users on Airbnb seem not to fully use the whole breadth of the 5-point rating scale, and most users’ aggregated ratings are located in the range between 4.5 and 5.0 stars (Teubner, Hawlitschek, and Dann 2017). While there is little empirical support for the explanation that there are simply no negative experiences with hosts on Airbnb, this observation may root in the direct social interaction between sellers and buyers on the platform Airbnb – after all, in about 40% of all Airbnb offers host and guest spend some time under the same roof.

In 2008, eBay went one step further in the redesign of their reputation system. Another change in the feedback rules stated that sellers were still able to submit positive or neutral feedback, but had no possibility anymore to submit negative feedback about the buyer. The idea behind this rule change was indeed to continue to allow for positive reciprocity, but to prevent negative reciprocity and feedback retaliation. However, from a game-theoretical perspective, this change should lead to different incentives for feedback timing and transaction behavior. Firstly, the new system prevents feedback retaliation. Sellers who employ feedback strategically will now only have the channel of positive reciprocity. Contrary to the situation with negative reciprocity where the seller should wait for the buyer to submit feedback first, they should now submit their positive feedback as early as possible, in order

to trigger a positive reaction on the buyer side. Secondly, buyers should be more satisfied after the rule change, since the incentives for sellers to behave in a trustworthy manner are stronger. But the system prevents the punishment of misbehaving buyers. The reduced incentives to behave trustworthy should result on less cooperative behavior of buyers, possibly ranging from lower payment reliability over less feedback system participation to blackmail using the threat of negative feedback. Sellers may be much less happy under the new system.

Field data on the new system provide evidence for some of these theoretical predictions. After the 2008 rule change in the reputation system that prevented sellers from giving negative feedback, sellers started to submit their feedback much earlier, and buyers were less and less likely to participate in the reputation system (see Bolton, Greiner and Ockenfels, in progress). Hui, Saeedi and Sundaresan (2017) as well as Klein, Lambertz and Stahl (2016) find evidence for better average product quality and higher satisfaction of buyers after the system change. Reliable data on satisfaction of sellers and misbehavior of buyers, however, is hard to obtain. There exists only some anecdotal evidence that the number of complaints from sellers about buyers increased after the change, and that in particular blackmailing of sellers with the threat of negative feedback is not uncommon. Using laboratory experiments, Greiner and Ockenfels (in progress) show that on platforms with two-sided moral hazard (that is, where both buyers and sellers have incentives to misbehave), a one-sided reputation system performs significantly worse than a two-sided system. These results are in line with the general insight from reputation system research that the appropriate design of a feedback system depends on the particularities of the underlying market.

## 3. Conflict resolution and escalation

Notwithstanding the steady development of new technologies to track payment and product streams between trading partners, there will always be details of a transaction that cannot be perfectly observed and controlled by the platform, such that it needs to rely on user reports and feedback. For example, eBay cannot directly verify whether or not a delivered product arrived in the previously advertised quality. There are additional legal reasons why platforms do not want to get involved in disputes over details of a transaction, since they may see their position as pure intermediaries, only facilitating match and price discovery between

buyer and seller but not being party to the trade contract. For this reason, many platforms offer decentralized conflict resolution mechanisms, in particular the institution of feedback withdrawal. The idea is to allow transaction partners, after they have established the existence of a conflict (usually manifested in negative feedback), to make good on their transaction behavior in a decentralized way, through repair, refund, or other reparatory actions. Depending on the exact rules on the specific platform, the feedback can then be withdrawn or revised, either unilaterally or only after mutual agreement of both transaction partners.

In their theoretical, empirical, and experimental study of the issue, Bolton, Greiner, and Ockenfels (2018) point out that the possibility to withdraw feedback and the details of the withdrawal mechanism may in turn affect feedback behavior as well as incentives to cooperate in the transaction. In the worst case, the existence of a conflict resolution mechanism may completely backfire, such that it does not contribute to a resolution of the conflict, but to its escalation. This is best illustrated with advice given to eBay sellers on the webpage ecommerce-guide.com. The author Frank Fortunato writes there:

> *"'Mutual Feedback Withdrawal' is the easiest and surest way to remove a negative from your rating. […] I recommend calling them after leaving the other party a negative feedback in response. It gives you leverage in further negotiations, and may be your only chance to do so because once you enter the Mutual Feedback Withdrawal process, eBay will not allow you to leave feedback for the transaction."*

The statement implies that make-good effort is only required when a seller failed to secure sufficient negotiation power for the feedback withdrawal phase by giving negative feedback on the buyer. Submitting negative feedback is thus strategically advisable, even when the seller does not yet know the buyer's feedback but suspects some potential dissatisfaction on the buyer's side. However, when feedback is submitted (and withdrawn) purely strategically, then the reputation system will lose its bite, and will not provide incentives that discipline trading behavior. If it is possible for a seller to avoid any make-good and reverse received negative feedback by submitting a negative feedback herself, then she does not have any incentives to behave in a trustworthy manner in the transaction in the first place.

Bolton, Greiner, and Ockenfels (2018) formalize this argument in a game-theoretical model and verify its implications in field data from eBay as well as in a laboratory experiment. On

eBay they observe that the initiation of a feedback withdrawal process is often accompanied by negative feedback reciprocity. When a seller responds to a buyer's negative feedback with a negative feedback by herself, then the likelihood that she will also initiate a feedback withdrawal is 39 percent, while this likelihood is only 16 percent when the seller did not retaliate. Another observation in the empirical data, however, is that the success rate of initiating a feedback withdrawal (i.e., whether feedback was eventually withdrawn or not) is *lower* when it was accompanied by a negative feedback as opposed to no such retaliatory feedback. Had the seller responded with a negative feedback, then the probability that the negative feedback is eventually withdrawn after initiation of the withdrawal process equals 54 percent. When the seller did not retaliate, this success rate is 74 percent. This observation contradicts the theoretical argument made above, that responding with a negative feedback gives you more strategic leverage in the feedback withdrawal process. It may be that some buyers react aversely when a seller does not only deliver low product quality but then also gives unsubstantiated negative feedback and wants her feedback to be withdrawn. Unfortunately, it is difficult to verify such explanations based on eBay field data, since we cannot observe the original product quality or in how far the seller made good after receiving negative feedback. Laboratory experiments can deliver complementary evidence that shed light on these questions.

In their laboratory experiment, Bolton, Greiner, and Ockenfels (2018) model the interaction between buyer and seller as a multi-stage process, which abstracts away from many particularities in the field. First, two agents decide whether they want to trade with each other or not. Second, the buyer decides about whether to send the payment or not, while simultaneously the seller determines the (initial) quality of the product to be sent. In the third stage, after observing the transaction result, both transaction partners simultaneously give feedback, which can be positive or negative. After being informed about the feedback they received, in the fourth stage the buyer has another chance to facilitate payment (if he hasn't done so yet) while the seller can improve upon the initial quality of the product. Experimental treatments only differ in the fifth and last stage of the interaction. In one condition, the transaction partners can mutually withdraw their feedback (FBW), while in the other condition, such feedback withdrawal is not feasible (noFBW).

Data from these laboratory experiments provide a clear picture. The mere existence of a feedback withdrawal stage invites strategic gaming in the previous stages, a behavior that manifests itself in several observations. Negative feedback is much more common in the FBW condition, and less dependent on the actual trading behavior of the transaction partner. Make-good is less often observed when the respective trader has given a negative feedback herself (such that she has bargaining power in the subsequent withdrawal stage). In the FBW treatment, agreement to withdrawing feedback is more conditioned on whether a trader has received a negative feedback herself rather than on whether the other has actually made good. In total, these strategic reactions to the FBW stage yield a strong bias in the feedback information collected in the reputation system. In the end, 50 percent of traders who did not cooperate end up with a positive feedback. This bias results in less trust and less trustworthiness in the transaction stage. In other words, the existence of a feedback resolution stage does not reduce conflicts, but actually escalates them. Another observation in the experiment, however, is that there also exist traders who, independently of whether they received a negative feedback themselves, give negative feedback when the other trading partner does not cooperate, and condition withdrawal only on proper make-good behavior. These "altruistic punishers" thus provide a public good of (costly) disciplining non-cooperators in the market, and their existence weakens the negative strategic effects of the conflict resolution mechanism on market and feedback behavior.

In a follow-up study, Bolton, Breuer, Greiner, and Ockenfels (2019) propose a new, improved design of a conflict resolution mechanism and test it experimentally against the previously described institutional setups. They term their design the "feedback revision option" (FBRO), and it features only a very minor change compared to the FBW rule described above. While under FBW both trading partners agree to mutually withdraw their feedback, under FBRO both trading partners agree to *allow* each other to unilaterally withdraw their feedback. That is, after both agreed, each trader can still independently (and simultaneously, i.e. unobserved) decide whether she would like to go through with the feedback withdrawal or not. In theory, this breaks the link between having given negative feedback and maintaining negotiation power in the withdrawal stage, since the withdrawal of one's own negative feedback cannot be conditioned on the withdrawal of the other's negative feedback. As a consequence, the negative strategic effects of a feedback withdrawal option should not be

existent with such a (slightly tweaked) design. The experimental data squarely support these theoretical predictions. As a matter of fact, the FBRO design does not only prevent the negative effects of FBW, it has even positive effects on trade and efficiency, because it allows to coordinate mutual expectations of traders in the transaction.

## 4. Reputation portability between platforms

The existence of separate reputation systems on many different platforms such as Amazon, eBay, or Airbnb raises the question of opportunities and perils of reputation portability, that is, the transfer of reputation from one context or platform to another (Teubner, Hawlitschek, and Adam 2019). Considering the effects of reputation *within* a specific marketplace and the increasing "platformization" of our world, a permeability of online reputation *across* platform borders appears quite intriguing. Already back in the 1990s, when eBay and Amazon started their reputation systems, Amazon tried to allow its users to import ratings from eBay (Resnick et al. 2000). However, eBay was not appreciative of this and claimed its users' ratings as a proprietary resource, accompanied by legal threats. Amazon eventually discontinued the reputation import service (Dellarocas, Dini, and Spagnolo 2009). In the next years, the idea of reputation portability did not receive much attention until it was picked up again recently with the wide-spread proliferation of sharing platforms (Tadelis 2016; Puschmann and Alt 2016). For such consumer-to-consumer platforms, which facilitate the sale or rental of goods and services between private individuals, reputation portability is particularly relevant. Typical examples include Airbnb and Homestay for apartments and rooms, BlaBlaCar and Uber for car sharing and cab services, Turo (private mutual car rentals), or TaskRabbit for a variety of home services. The non-professional transaction partners on these platforms do not have the facilities to build up and rely on a brand name that is known across platforms, but must built and maintain a reputation on their own, in order to be successful as a seller or even be allowed to book, buy, or rent as a consumer.

Typically, users have to create new profiles and build up a new history and reputation on each platform they want to transact on. This raises the question whether it should be possible to transfer reputational information across platforms, and if so, to which extent. A reputation transfer may also address the notorious cold-start problem (Wessel, Thies, and Benlian 2017) typical for platforms, two-sided markets, and other environments governed by network

effects. New hosts on Airbnb, for example, would be able to refer to their ratings on eBay, and thus benefit from their long-standing and impeccable history as reputable sellers. Otherwise, they would represent an unknown quantity on Airbnb, reducing their (and the platform's) chances for profitable transactions.

Despite the increasing importance of consumer-to-consumer platforms and calls for research combining "economic thinking and analyses with human-computer interaction" (Tadelis 2016, p. 338) and on "how consumers [might] connect different identities on different sharing platforms towards a cross-platform identity management" (Puschmann and Alt 2016, p. 98), research into this matter is scarce. Recent studies have assessed the effectiveness of portable reputation from a user's psychological perspective. Using a scenario-based experiment, Otto et al. (2018) presented users within a ride sharing scenario with imported ratings from Airbnb, finding that that imported star ratings facilitated trust in prospective service providers. Zloteanu et al. (2018) explored the impact of "trust and reputation information" (TRI) on users' intentions to book on a hypothetical accommodation platform. They find a positive effect of TRI on users' perception of a host's trustworthiness and increased booking intention. As one form of TRI, the authors explored "online market reputation," that is, ratings imported from other platforms. Their data, however, made it difficult to assess which type of TRI users value most, though it seemed that imported reputation played a role for creating trust.

Teubner, Adam, and Hawlitschek (in progress) study how trusting potential consumers are when service providers' (i.e., seller, host, driver, etc.) transfer reputation between platforms. In particular they study situations in which the provider has not yet built a platform-specific reputation, but information about the reputation of this provider on other platforms is made available. This setting is compared to situations in which a) the provider has already established a reputation on this specific platform, and b) the provider has not yet established a reputation and no other information is available. The empirically emerging behavioral patterns are consistent across a number of different real-world platforms and combinations of these. Providers with transferred reputation from another platform are perceived as significantly more reputable and trustworthy than providers without any reputational history. However, transferred reputation scores are less effective (for establishing trust and purchase intentions) than reputation built up on the very platform, and the effectiveness of transferred reputation hinges on the fit between source and target context.

Thus, even when created in other contexts, reputation can be seen as part of a provider's economic capital to attract consumers. The fact that individual reputation histories are distributed across isolated platforms hinders the establishment of trust between platform users, reduces the number of realized transactions, and thus lowers economic efficiency in the larger e-commerce sector. Platforms may be subject to a prisoner's dilemma situation, in that they would benefit when other platforms release their reputation data, but have no incentives to allow to transfer their own data since that would also release users who are currently locked in (Krämer 2018).

Hence, from a regulatory perspective, reputation portability represents an intriguing subject where the EU may even force platforms to enable the export and import of reputation. In fact, the European Commission views reputation portability as an important means to solve issues of personal data sovereignty, platform lock-ins, and competition between platforms (European Commission 2017, p. 93). In this regard, scholars have recently referred to Article 20 of the EU's General Data Protection Regulation. The Article compels platform operators (referred to as "controllers") to grant their users (referred to as "data subjects") the right to data portability. Specifically, a user "shall have the right to receive the personal data concerning him or her, which he or she has provided" to the platform (EU 2018). It is unlikely, however, that this provision applies to reputational information such as rating scores and text reviews, for at least two reasons. First, the article explicitly refers to information which users have provided *themselves* (i.e., comments, posts, likes, photos, etc.). Reputation data, however, is provided by *others*, namely by prior transaction partners. Second, even if Airbnb and other platforms would enable the export or download of a user's reputation technically, this data would lose its credibility and hence trust-building potential as soon as it leaves the platform, as it could become subject to post-hoc manipulation by the user. Given that most reputation systems' effectiveness is based on the fact that they are not easy to manipulate, such portable reputation would either have to be sealed through a digital signature or transferred between platforms directly. Both are neither warranted by EU's regulation nor is it likely that the major platforms are willing to implement such measures any time soon.

However, there exist several services which congregate reputation profiles from different platforms and provide them in aggregated form (e.g., Deemly.co and Traity.com). Users need to provide their access credentials (username and password) to the different platforms to the

aggregator, which requires a large amount of trust into the institution and thus represents a significant roadblock to broad participation in these systems.

## 5. Conclusions

We summarize the findings from the literature on reputation systems, including both our own and the work of other researchers in the field, in a number of suggestions and implications relevant to managers and professionals who interact with online markets, be it as buyers, sellers, or platform providers themselves.

**Trust and trustworthiness on platforms are supported by reputation and conflict resolution systems.** However, small details in a platform's specific institutional rules can have large impacts on individual and aggregate market behavior. The optimal design depends on the institution-dependent moral hazard opportunities of different user groups on the platform, and the value these user groups contribute to the platform.

**Our research strongly suggests that a one-size-fits-all solution does not exist for reputation systems.** The rules necessary to reach a market equilibrium that features high trust and trustworthiness as well as economic efficiency will depend on the specific economic conditions on the platform.

**One major focus for market designers and sellers should lie on creating incentives for users to provide feedback.** Strong participation in the feedback system remedies the issues of social and strategic biases, selective reporting, as well as fake positive and negative reviews. Encouragement of feedback giving may include future rebates, reminders, sequencing from general to more detailed feedback (as used on Booking.com, for example), or even enforcement before another transaction is possible (as implemented on Uber).

**Open feedback system with endogenous feedback timing will further reciprocity in feedback giving.** Reciprocity has negative consequences for the informativeness of feedback. While purely one-sided or blind/simultaneous feedback systems may resolve the issue of reciprocity, they may also result in lower participation in the feedback system. Hybrid systems may be able to strike the right balance.

**Conflict resolution systems need to be carefully designed, since they may escalate rather than dampen the likelihood of conflicts.** If users anticipate the possibility of withdrawing feedback and can strategically influence their subsequent negotiation power, the feedback system may be rendered less effective. Proper withdrawal rules that mitigate incentives for strategic gaming can improve coordination between trading partners and thus conflict resolution.

**Feedback and conflict resolution rules may also affect endogenous, social disciplining mechanisms.** In particular, there exist "altruistic punishers" among traders who provide a public good to their market platform, and whose intrinsic incentives may be crowded out when market rules are changed.

**Allowing the transfer of user reputation has the potential to increase trust and trustworthiness between market sides across platform boundaries.** Many users are active on multiple platforms, on each for a different reason. Even though many technical as well as legal details with respect to transferring feedback and reputation from one platform to another are still debated and subject to research, it appears that reputation mobility can increase trading volume and market efficiency in electronic commerce. For (new) market platform providers, this implies that they should seriously consider to allow their users to import (or credibly refer to) reputational data from other platforms, thus giving them a kickstart in terms of perceived trustworthiness, and lower entry barriers.

**And finally: Well-functioning reputation systems are characterized by a fine balance between mutual punishment and counter-punishment opportunities.** We strongly recommend proper theoretical analysis, empirical evaluation, and experimental testing before rules governing feedback and conflict resolution on market platforms are introduced or changed. Otherwise, market providers face the risk of biasing information, setting negative incentives, or crowding out intrinsic motivation. Academic researchers can be of tremendous help here to practitioners, since any experience with real-world market design problems also furthers our knowledgebase of the functioning of reputation systems.

## Conflict of interest statement

On behalf of all authors, the corresponding author states that there is no conflict of interest.

**Literature**


Bajari, P. and Hortaçsu, A. (2004), Economic Insights from Internet Auctions, *Journal of Economic Literature* 42 (2), 457-486.

Bolton, G., Breuer, K., Greiner, B. and Ockenfels, A., (2019). The design of feedback revision rules – An experimental study. *Working Paper*.

Bolton, G., Greiner, B. and Ockenfels, A., (2013). Engineering trust: Reciprocity in the production of reputation information. *Management Science* 59(2), 265–285.

Bolton, G., Greiner, B., and Ockenfels, A., (2018). Dispute Resolution or Escalation? The Strategic Gaming of Feedback Withdrawal Options in Online Markets. *Management Science* 64(9), 2018, 4009–4031.

Bolton, G., Greiner, B., and Ockenfels, A., (in progress). One-sided feedback with two-sided moral hazard – An empirical study.

Bolton, G., Katok, E., and Ockenfels, A., (2004). How effective are online reputation mechanisms? An experimental study. *Management Science* 50(11), 1587–1602.

Cabral, L. and Hortaçsu, A. (2010), The Dynamics of Seller Reputation: Evidence from eBay, *Journal of Industrial Economics* 58, 54-78.

Dann, D., Teubner, T., and Weinhardt, C. (2019). Poster child and guinea pig—Insights from a structured literature review on Airbnb. *International Journal of Contemporary Hospitality Management 31*(1), pp. 427-473.

Dellarocas, C., (2004). Building trust on-line: The design of robust reputation mechanisms for online trading communities. In: *Social and Economic Transformation in the Digital Era* (Eds. Doukidis, Mylonopoulos, Pouloudi). Idea Group Publishing, Hershey, PA, S. 95–113.

Dellarocas, C., Dini, F., and Spagnolo, G., (2009). Designing reputation mechanisms, in *Handbook of Procurement*, pp. 446-482.

Easley, D., and Kleinberg, J. (2010). Network Effects, In *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*, Cambridge University Press, pp. 509–542.

Eaton, D. H. (2007), The Impact of Reputation Timing and Source on Auction Outcomes, The *B.E. Journal of Economic Analysis & Policy* 7(1), Article 33.

European Commission (2017). Exploratory study of consumer issues in online peer-to-peer platform markets. http://ec.europa.eu/newsroom/just/item-detail.cfm?&item_id=77704

European Union (2018). REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation).

Forbes (2018). As A Rare Profitable Unicorn, Airbnb Appears To Be Worth At Least $38 Billion. Accessible at https://www.forbes.com/sites/greatspeculations/2018/05/11/as-a-rare-profitable-unicorn-airbnb-appears-to-be-worth-at-least-38-billion/#3d77a6022741; Retrieved February 12th, 2019.

Fradkin, A., Grewal, E., and Holtz, D., (2017). The determinants of online review informativeness: Evidence from field experiments on Airbnb, *Working Paper*.

Greiner, B. and Ockenfels, A. (in progress). One-sided and two-sided feedback systems in markets with one-sided and two-sided moral hazard – An experiment.

Greiner, B., Teubner, T., and Weinhardt, C. (2018). Grundfragen der Plattformökonomie – Wie man Vertrauen designt. In *Plattformen – Geschäftsmodell und Verträge* (Eds. Uwe

Blaurock, Martin Schmidt-Kessel and Katharina Erler), Schriften der Ernst von Caemmerer Stiftung vol. 10, Nomos.
Houser, D. and Wooders, J. (2005), Reputation in Auctions: Theory and Evidence from eBay, *Journal of Economics and Management Strategy* 15 (2), 353-369.
Hui, X., Saeedi, M., and Neel Sundaresan, N. (2017). Adverse Selection or Moral Hazard: An Empirical Study. *Working Paper*.
Jin, G.Z. and Kato, A. (2006). Price, Quality and Reputation: Evidence from an Online Field Experiment. *RAND Journal of Economics* 37(4), 983-1005.
Klein, T. J., Lambertz, C., and Stahl, K., (2016). Market Transparency, Adverse Selection, and Moral Hazard, *Journal of Political Economy* 124(6): 1677–1713.
Krämer, J. (2018). Datenschutz 2.0 – ökonomische Auswirkungen von Datenportabilität im Zeitalter des Datenkapitalismus. In *T. Kretschmer et al. (Ed.), Datenkapitalismus – eine ökonomische Betrachtung* (Vol. 98, pp. 459–480).
Livingston, J. A. (2005), How Valuable is a Good Reputation? A Sample Selection Model of Internet Auctions, *Review of Economics and Statistics* 87 (3), 453-465.
Lucking-Reiley, D., Bryan, D., Prasad, N., and Reeves, D. (2007), Pennies from eBay: the Determinants of Price in Online Auctions*, Journal of Industrial Economics* 55(2), 223-233.
Otto, L., Angerer, P., and Zimmermann, S. (2018). Incorporating external trust signals on service sharing platforms. In *ECIS 2018 Proceedings*, Portsmouth, UK.
Puschmann, T., and Alt, R. (2016). Sharing economy. *Business & Information Systems Engineering*, 58(1), 93-99.
Resnick, P., R. Zeckhauser, E. Friedman, and K. Kuwabara (2000), Reputation Systems, *Communications of the ACM* 43 (12), 45-48.
Tadelis, S. (2016). Reputation and feedback systems in online platform markets. *Annual Review of Economics*, 8, 321-340.
Teubner, T., Adam. M. T. P., and Hawlitschek, F. (in progress). Unlocking online reputation: On the effectiveness of cross-platform signaling in the sharing economy.
Teubner, T., and Dann, D. (2018). How platforms build trust. *Working Paper.*
Teubner, T., and Hawlitschek, F. (2018). The economics of P2P online sharing, In: *The Sharing Economy: Possibilities, Challenges, and the way forward* (Eds. Albinsson, P. A., Perera, B. Y.): 129–156, Praeger Publishing.
Teubner, T., Hawlitschek, F., and Adam, M. T. P., (2019). Reputation Transfer*. Business & Information Systems Engineering* (in press).
Teubner, T., Hawlitschek, F., and Dann, D., (2017). Price determinants on Airbnb: How reputation pays off in the sharing economy. *Journal of Self-Governance and Management Economics* 5(4): 53–80.
Wessel, M., Thies, F., and Benlian, A. (2017), Competitive Positioning of Complementors on Digital Platforms: Evidence from the Sharing Economy, in *ICIS 2017 Proceedings*.
Zervas, G., Proserpio, D., and Byers, J., (2015). A first look at online reputation on Airbnb, where every stay is above average, *Working Paper*.
Zloteanu, M., Harvey, N., Tuckett, D., & Livan, G. (2018). Digital Identity: The effect of trust and reputation information on user judgement in the Sharing Economy. *PLOS ONE*, 13(12), pp. 1-18.