European Economic Review 160 (2023) 104586

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/eer

To insure or not to insure? Promoting trust and cooperation with insurance advice in markets

Ben Grodeck[a,*], Franziska Tausch[b], Chengsi Wang[c], Erte Xiao[c]

[a] Max Planck Institute for Evolutionary Biology, Germany
[b] Independent Consultancy, Cologne, Germany
[c] Department of Economics, Monash University, Australia

ARTICLE INFO

JEL codes:
C91
D9
D47
D82
L86

Keywords:
Asymmetric information
Insurance
Trust
Cooperation
Communication
Experimental economics

ABSTRACT

We design and test a novel insurance advice mechanism aimed at promoting trust and cooperation in markets with asymmetric information. In a buyer-seller game with third-party insurance, sellers have the option to advise buyers on whether to purchase insurance against the potential losses from the opportunistic behavior of strategic sellers. We hypothesize that advising not to purchase insurance introduces a psychological cost for defection. We develop a theoretical model that selects a pooling equilibrium where both cooperative and strategic sellers advise buyers not to purchase insurance. Once this advice has been given, strategic sellers choose not to defect if the associated psychological costs are sufficiently large. Data from a controlled laboratory experiment shows that the insurance advice mechanism significantly increases market efficiency, with buyers being more likely to purchase from sellers and sellers being more likely to cooperate. Furthermore, we find that the insurance advice mechanism is more effective when sellers can observe buyers' insurance purchase decisions.

1. Introduction

Asymmetric information is ubiquitous in economic transactions. Consumers are often unable to verify sellers' credibility before purchasing a product. In this case, if the consumer does not trust the seller, they may refrain from the transaction entirely. Many mechanisms have been designed to solve this asymmetric information problem and facilitate efficient transactions. Among them, insurance (or warranties) is a common practice. In particular, the buyer can purchase insurance or warranties from a third-party provider to protect his purchase. For example, online markets such as eBay offer the option of purchasing warranties from the third-party provider Squaretrade. Buyers can purchase the warranty either at the time of buying the product on eBay or directly from Squaretrade's website after purchasing the product (Steiner, 2012). The insurance provided by a third party offers additional or extended coverage to the existing manufacturer's warranty or protection when consumers purchase products, especially second-hand ones, for which the manufacturer's warranty is not honored. However, such insurance is often costly and comes with exclusions and limitations. If buyers are unwilling to pay the cost of insurance or are discouraged by complicated exclusion clauses, the inclusion of insurance may not generate more transactions. Moreover, insurance provided by a third party may not change the incentives for sellers or manufacturers to cooperate and thus have a limited impact on improving consumers' willingness to trade.

In this paper, we propose and test a novel insurance advice mechanism aimed at promoting trust and cooperation in markets with

* Corresponding author.
E-mail address: grodeck@evolbio.mpg.de (B. Grodeck).

https://doi.org/10.1016/j.euroecorev.2023.104586
Received 4 March 2022; Received in revised form 6 September 2023; Accepted 11 September 2023
Available online 22 September 2023

asymmetric information. The key to our proposed mechanism is to allow the seller (the party who has more information) to advise the buyer whether he should purchase third-party insurance.[1] In general, insurance often addresses two different types of risks: (1) risks about the seller's cooperative type, such as her intention to deliver the product on time or her intention to produce a high-quality product as advertised; and (2) natural risks that are out of her control, such as bad weather that causes the delay of the shipment. As our focus is on the asymmetric information problem, the proposed advice mechanism is related to insurance against the first type of risk. Our mechanism builds on markets where an insurance option is already in place, such that the marginal cost of introducing the advice option is negligible. As a first step, we test the mechanism built on insurance provided by a third party instead of the seller. This feature avoids potential confounds due to the additional profit incentives that sellers may have to sell insurance.

We hypothesize that advising not to purchase insurance introduces a psychological cost for defection. First, giving advice may lead the seller to feel more accountable for the buyer's payoffs, as she now plays a more active role in the buyer's decision (Tetlock, 1985; Lerner and Tetlock, 1994, 1999). If the seller is subject to omission bias (Ritov and Baron, 1992), she may judge defection—after advising the buyer not to protect himself from the risk—as morally worse than when she does not exert any influence on the buyer's decision. Second, if the advice of not purchasing insurance is taken as a statement that the seller will cooperate, subsequent defection may render the advice a lie and inflict psychological costs due to lying aversion (Cressey, 1986; Gneezy et al., 2013; Abeler et al., 2014, 2019). Third, if advising not to purchase insurance increases the buyer's expectation that the seller will cooperate, the seller may be averse to disappointing the buyer (Charness and Dufwenberg, 2006; Battigalli and Dufwenberg, 2007; Balafoutas and Sutter, 2017; Cartwright, 2019). Finally, these above-mentioned psychological mechanisms may mitigate the moral wiggle room that sellers can exploit to pursue self-interest (Dana et al., 2007; Bénabou and Tirole, 2016). For example, without insurance advice, the seller can plausibly justify her defection through reasoning, such as "It was the buyer's decision, I am not responsible for them choosing that." Such self-serving reasoning is no longer easy after the seller takes a more active role in the buyer's decision, or has to lie to defect, and thereby opportunistic behavior becomes more psychologically costly.[2]

We consider a theoretical model that contains both cooperative sellers, who always ship the product, and strategic sellers, who ships the product only if doing so maximizes their utility. Our model predicts that a pooling equilibrium with both types of sellers advising not to purchase insurance is likely to emerge, provided psychological costs are sufficiently large. At equilibrium, the buyer follows the advice and purchases the product without insurance, and the seller subsequently ships the product. As a result, the insurance advice mechanism achieves more efficient trades.

We conduct a controlled laboratory experiment to examine the effectiveness of the mechanism empirically. In particular, we address two main research questions. Does the insurance advice mechanism increase the number of buyers who enter transactions with sellers? Are sellers more likely to cooperate with buyers under the insurance advice mechanism?

Although in online marketplaces such as Amazon and eBay, buyers' insurance purchase decisions can be easily made observable to sellers, we take into account the fact that sellers do not always observe the buyers' insurance purchase decisions when the insurance is provided by a third party. Theoretically, the insurance advice mechanism can help build trust and improve efficiency even if the seller does not observe the buyer's actual insurance purchase decision. This is the case because the seller has a psychological cost associated with not shipping the product and becomes more likely to ship the product when she anticipates that the buyer may buy the product without insurance. In turn, knowing that the seller may ship the product, the buyer is also willing to follow the insurance advice with some non-zero probability. The improvement, however, is not as effective compared to when sellers perfectly observe the insurance purchase decisions. We empirically test whether the effectiveness of the mechanism varies based on the observability of the buyer's insurance purchase decision.

The experiment consisted of three treatments. The control treatment was a buyer-seller game with insurance. In the game, the buyer decided whether to purchase a product, and the seller decided whether to ship the product upon receiving the payment. If the buyer purchased the product, he could also purchase insurance against the risk that the seller might not ship the product after receiving the payment. To test the insurance advice mechanism, we designed two treatments: insurance advice (IA) treatment and insurance advice with hidden information (IA_HI) treatment. In both treatments, we introduced an insurance advice mechanism in which the seller had to advise the buyer whether to purchase the insurance before starting the buyer-seller game. Upon receiving the advice, the buyer decided whether to buy the product and, if so, whether to purchase the insurance. In the IA treatment, if the buyer purchased the product, the seller was informed of the buyer's insurance purchase decision before deciding whether to ship the product. In the IA_HI treatment, the seller never learned about the buyer's insurance purchase decision. This is the only difference between the two treatments. We used shipping as a simple way to introduce defections in the game. If the proposed mechanism works in this setting, it should also effectively reduce other types of defections, such as selling faulty products.

Our findings are consistent with our hypotheses. In the IA treatment, sellers advise not to purchase insurance 81 % of the time. Compared with the control treatment, the rate of product purchases increases by approximately 31 % in the IA treatment. Whereas buyers purchased the product 74 % of the time when sellers advised not to purchase insurance, they purchased it only 35 % of the time when the advice was to purchase insurance. The number of sellers who shipped the product also increases by almost 40 %. These changes result in an increase in market efficiency. The mechanism remains effective in the IA_HI treatment. About 71 % of sellers advise not to purchase insurance. Compared to the control, the product purchase rate increases by 33 %, and the shipping rate increases by approximately 22 % in the IA_HI treatment, increasing market efficiency. However, as our theoretical model predicts,

[1] For simplicity, we will use "she" to refer to the seller and "he" to refer to the buyer.

[2] On the other hand, advising the buyer to purchase the insurance could reduce the seller's accountability for the buyer's losses and increase the moral wiggle room for excusing her opportunistic behaviour.

compared to the IA treatment, the IA_HI treatment is less efficient because buyers were less likely to follow the advice of no insurance, and sellers were slightly less likely to ship the product.

This paper contributes to two strands of the literature. One is research on market design aimed at solving market failures due to information asymmetry. Several innovative solutions have been proposed and tested, including the widely studied reputation mechanism (for a review, see Chen et al. 2021). Previous studies have examined how to improve the reliability of reputation mechanisms that are subject to problems such as missing information (Resnick and Zeckhauser, 2002; Bolton et al., 2004; Dellarocas and Wood, 2008; Cabral and Hortacsu, 2010; Li and Xiao, 2014; Bolton et al., 2018; Bolton et al., 2019) and manipulating reviews (Mayzlin et al., 2014).

The simple mechanism we propose complements this literature by pointing out a new direction for solutions. For example, on eBay, buyers are offered extended warranties via Squaretrade or xcover.com. To implement the insurance advice mechanism, eBay could allow sellers to recommend to the buyer whether he should purchase the extended warranty. The advice mechanism can be especially beneficial for new sellers before they are able to establish a positive reputation via a feedback mechanism. The mechanism can also work in offline markets, such as the used car market. For instance, the original car owner may suggest whether the potential buyer should purchase an extended warranty from a third party.

Although our main interest is to propose and test the effectiveness of the insurance advice mechanism, it is interesting to consider how the advice mechanism relates to other types of communication (in particular promises) that have been shown to be effective in promoting cooperation (Ellingsen and Johannesson, 2004; Binmore, 2006; Charness and Dufwenberg, 2006; Bicchieri and Lev-On, 2007; Vanberg, 2008; Sánchez-Pagés and Vorsatz, 2007; Erat and Gneezy, 2012; Battigalli et al., 2013; López-Pérez and Spiegelman, 2013). Sally (1995) conducted a meta-analysis and found that communication was the most effective factor in promoting cooperation in prisoner's dilemma experiments. In particular, the analysis shows that elicited promises have a strong additional effect on cooperation. Recently, there has been emerging experimental research on the relationship between promises and cooperation. For example, researchers have shown that people follow their promise to either avoid lying costs (Ellingsen and Johannesson, 2004; Vanberg, 2008; Serra-Garcia et al., 2013) or as a result of guilt aversion (Charness and Dufwenberg, 2006; Battigalli et al., 2013). In our experiment, when the seller advises not to purchase insurance, the buyer may interpret this message as an implicit promise to ship the product.

It is worth noting that the literature on communication and promises highlights that compared to unconstructed free-form communication, restricted-form communication, such as binary messages, is often much less or not effective at all (Bracht and Feltovich, 2009). For example, a pre-formulated promise, such as "I promise to cooperate," may have a smaller effect than free-form communication or no effect at all on increasing cooperative behavior, specifically in strategic environments (Lundquist et al., 2009; Charness and Dufwenberg, 2010; Belot et al., 2012; Chen and Zhang, 2021; Brandts et al., 2019). The advice in our setup is given in a binary form. Sellers only choose between "Advise the buyer to purchase the insurance" or "Advise the buyer not to purchase the insurance." Thus, if the buyer perceives the advice as an implicit promise, it is, at best, a weak and indirect bare promise. The significant effect of a plain insurance advice message suggests that there may be some fundamental differences between advice and bare promise and that additional psychological channels may play a role.

Furthermore, according to the philosophy of language (Searle, 1975), the advice mechanism theoretically differs from bare promises. When a seller utters a bare promise to cooperate, she 'commits' to cooperate (without explicitly asking the buyer to take a certain action). This promise can be seen as a commissive speech act, defined as a speech act that the speaker intends to commit (Searle, 1975).[3] By contrast, when advising the buyer not to purchase insurance, the seller persuades the buyer to perform a specific action (e.g., not to take action to protect against potential losses). For Searle (1975), advice falls under the category of directive speech acts, defined as communication persuading another party to act in a particular manner. If a directive speech act is viewed as playing a more active role in the buyer's outcomes than a commissive speech act, omission bias and accountability theory predict that the advice mechanism can promote cooperation even when a bare promise is ineffective. Further, although a bare promise to ship the product may encourage more buyers to purchase the product and more sellers to cooperate with buyers, it is unclear how it affects insurance purchase decisions. One advantage of the insurance advice mechanism is that, in addition to promoting more transactions, it increases buyers' welfare by saving the cost of purchasing insurance—as observed in the comparisons between the IA treatment and the control. This could also be the difference between advising not to buy the insurance and other common advertising strategies used to persuade buyers to purchase a product. We discuss this in more detail at the end of the paper.

2. Theoretical framework and hypotheses

Consider a bilateral transaction between a buyer and a seller. We compare sellers' and buyers' decisions in three possible scenarios: a scenario without the insurance advice mechanism (control), a scenario with the insurance advice mechanism (IA), and a scenario with the insurance advice mechanism where sellers cannot observe buyers' insurance purchase decisions (IA_HI). Below, we present a theoretical framework for deriving predictions for sellers' and buyers' decisions in each scenario. We start with the scenario without the insurance advice mechanism (control).

The buyer demands one unit of the product that the seller produces and attaches a value ($v > 0$) to it. The seller attaches zero value to the product and can produce it at zero cost. We assume that both parties are risk-neutral. Following the previous literature (Bolton et al., 2004; Lafky, 2014; Li and Xiao, 2014), the product's price is set as fixed in the experiment. Specifically, the product price is

[3] Reassuring messages that aren't promises, such as "I plan to invest", as seen in Bracht and Feltovich, 2009 would also be considered commissive speech acts, as the statement is about the speaker's intended action.

exogenously given by $p \in (0,v)$. The fixed price excludes the possibility that prices could be used as a signal for seller type and allows us to provide clean evidence for the effect of the insurance advice.[4]

If the buyer purchases the product, the seller can ship the product at a cost of $d \in (0,p)$. Following the standard approach of modelling seller reputation in an asymmetric information environment with both moral hazard and adverse selection problems (Bar-Isaac and Tadelis, 2008), we assume that there are two types of sellers: a good type (type-*g*) who always ships the product and a strategic type (type-*s*) who maximizes her own utility, including potentially a psychological cost, which we explain below. Only the seller knows her type. The buyer does not know the seller's type, but he does know that the probability of encountering a type-g seller is $g \in (0,1)$, and that the probability of encountering a type-s seller is $1-g$.

Along with purchasing the product, the buyer has the option to buy insurance at price *w*, which allows the buyer to recoup *p* in case the product is not shipped.[5] We assume that the insurance is not too expensive, that is, $w \leq p\left(1-\frac{p}{v}\right)$, such that at least some buyers will buy the insurance.[6]

The buyer's purchase decision relies on the prior belief about the seller's type, *g*. The type-g seller always ships the product, whereas the type-s seller never ships the product. Given our assumption that $w \leq p\left(1-\frac{p}{v}\right)$, it is straightforward to show that the buyer's optimal decision is as follows:

$$\begin{cases} \textit{purchase the product without insurance}, \ \textit{if} \ g \geq 1-\dfrac{w}{p} \\ \textit{purchase the product with insurance}, \ \textit{if} \ \dfrac{w}{v-p} \leq g < 1-\dfrac{w}{p} \\ \textit{do not purchase the product}, \ \textit{if} \ g < \dfrac{w}{v-p} \end{cases} \tag{1}$$

That is, the buyer: purchases the product without insurance when *g* is relatively high; purchases the product with insurance if *g* is at some intermediate level; and does not purchase the product if *g* is very low.

Now consider the scenario in which the seller can advise the buyer whether to buy the insurance before making any purchase decision (IA). We denote the advice $a \in \{Y, N\}$, where *Y* means "buy the insurance" and *N* means "do not buy the insurance." In contrast to the control, in which the seller cannot influence the buyer's decisions, the seller's advice can change the buyer's expectation of the likelihood of receiving the product. As a result, the seller may become more accountable for the buyer's payoffs. Although the advice would not affect a type-g seller who always ships the product, we hypothesize a type-s seller will incur a psychological cost ($\alpha > 0$) for not shipping the product if (i) she advises *N* and (ii) the buyer purchases the product without insurance. The buyer does not know the exact value of α. However, for tractability, we assume he knows whether α is above or below *d*.

The seller may experience a psychological cost, even if the buyer purchases the insurance after she advises *N*. We assume the cost will be higher if the buyer follows the advice and does not purchase the insurance than purchase the insurance. In this sense, α can be understood as the incremental psychological cost between the two cases. Note that for simplicity, we also assume a type-s seller does not incur any psychological cost associated with the buyers' insurance purchase decisions made without the sellers' influence. That is, we assume that type-s seller's shipping decisions are not affected by buyers' insurance decisions in the control. This is because such a cost, if any, should be the same as in the IA. This simplicity allows us to focus on the effect of the advice mechanism.

The timing of the decision-making is summarized as follows. First, the seller advises whether to buy insurance. Next, the buyer receives the advice and decides whether to purchase the product and, if so, whether to buy the insurance. The seller observes the buyer's product purchase and insurance purchase decisions and decides whether to ship the product if the buyer purchases the product. The equilibrium concept is the weak perfect Bayesian equilibrium (WPBE). The formal propositions and proofs of our theoretical analysis can be found in Appendix A.

We note that the advice mechanism is not effective when $d > \alpha$ as the type-s seller would rather incur the psychological cost than the shipping cost. The mechanism can only be effective with $d \leq \alpha$. The analysis below will focus on the case where $d \leq \alpha$.

It is straightforward to show that there is no separating equilibrium. If the type-s seller's separating-equilibrium advice is *Y*, the buyer's optimal choice is not to purchase the product, as he anticipates that the type-s seller will not ship the product. Thus, the type-s

[4] In the field, sellers may use price to signal their type and affect buyers' purchase decisions. In particular, high prices can be the efficient means of signaling high quality because a loss in sales could hurt low cost, low-quality sellers more than high cost, high-quality sellers (Bagwell and Riordan, 1991; Bagwell, 1992). Assuming that a reasonable fraction of consumers can identify the true product quality before making purchases, if a low-quality seller pretends to be a high-quality seller by setting a high price, she will lose these well-informed buyers. That is, the loss in sales due to the high price will hurt low-cost, low-quality sellers more than high-cost, high-quality sellers, and high prices can then be used to effectively signal quality.

[5] We assume the seller does not receive commission from selling third-party insurances. If she does, the advice *N* should serve as an even stronger signal of being cooperative in the shipping stage. In reality, the buyer may also infer the distribution of the seller's type *g* from *w*. For example, the buyer may hold the belief of a low *g* in a market with high *w*. This will make him less likely to purchase the product, as we show below. However, it will not affect our hypotheses regarding the effects of the insurance advice mechanism.

[6] Without this assumption, buyers will never buy insurance. This is also the case when the insurance advice mechanism is introduced. We keep this assumption to make sure insurance is not redundant. In theory, if, instead, $w > p\left(1-\frac{p}{v}\right)$, buyers will never buy insurance because they receive a negative payoff. However, in this case, the advice mechanism is still beneficial in that it promotes more buyers to purchase the product and more sellers to ship the product. We discuss this in Appendix A4.

seller will advise *N* instead. If the type-s seller's separating-equilibrium advice is *N*, she will again be better off by instead advising *Y*, in which case she will get the full amount of payment p by not delivering, without incurring any psychological cost.

Next, we demonstrate that there exist two types of pooling equilibria when $d \le \alpha$ in the IA. First, there always exists a pooling equilibrium in which both types of sellers advise *N* (hereafter *N*-pooling equilibrium). Given that the type-g seller advises *N* in equilibrium, the type-s seller is better off pooling with them and advising *N*. As $d \le \alpha$, the type-s seller's optimal choice after advising *N* is to ship the product and incur the shipping cost. Thus, buyers purchase the product without insurance upon receiving advice *N*, and both types of sellers advise *N* and subsequently deliver the product.

There may exist another pooling equilibrium with both types of sellers advising *Y* (hereafter *Y*-pooling equilibrium). The Y-pooling equilibrium exists only when $g \ge \frac{w}{v-p}$. When $g \ge \frac{w}{v-p}$, the type-s sellers do not want to deviate to advising *N* instead since advising *N* and revealing their type leads to a lower profit equal to $p - d$. The *Y*-pooling equilibrium does not exist if $g < \frac{w}{v-p}$. This is because buyers will not purchase the product after receiving advice *Y* and the type-s seller is better off deviating to advising *N*. Upon observing the off-path advice *N*, buyers will buy the product because they know that the seller will ship the product irrespective of her type given $d \le \alpha$.

The above analysis suggests that both types of pooling equilibria may exist when $g \ge \frac{w}{v-p}$. The type-s seller prefers the *Y*-pooling equilibrium to the *N*-pooling equilibrium as in the *Y*-pooling equilibrium buyers still buy the product, but the type-s seller can choose not to ship without incurring any psychological cost. However, the *Y*-pooling equilibrium can be ruled out by applying forward induction.[7] The basic intuition is that, upon observing the off-path advice *N*, the buyer should believe that the seller is of type-g, as this seller is more likely to choose *N* than the type-s seller given only the latter incurs a psychological cost. Given that the *Y*-pooling equilibria may not always exist and can be ruled out by forward induction whenever it exists, we select the *N*-pooling equilibrium for the scenario with the advice mechanism. We can test whether indeed the N-pooling equilibrium is more likely to be played than the Y-pooling equilibrium in our experiment.

Suppose the *N*-pooling equilibrium is selected. Compared to the control, the advice mechanism always promotes efficient trades, as long as $d \le \alpha$ for some sellers. Specifically, when $g < \frac{w}{v-p}$, buyers purchase the product only when the advice mechanism is introduced. When $g \ge \frac{w}{v-p}$, buyers purchase the product regardless of whether the advice mechanism is introduced. For the sellers, while, by definition, type-g sellers always deliver the product, type-s sellers do so only in the scenario with the insurance advice mechanism. That is, the values from transactions are only realized in the scenario with the advice mechanism. In addition, with the insurance advice mechanism, buyers will not buy the insurance and thus save on the insurance cost.

Now we consider the case when the market introduces the insurance advice mechanism but the seller does not observe the buyer's insurance purchase decision (IA_HI). Intuitively, the equilibrium outcome remains the same as in the control when $d > \alpha$. The more interesting case arises when $d \le \alpha$. First, recall that if the proportion of type-g sellers is sufficiently high such that $g \ge 1 - \frac{w}{p}$, buyers will buy the product without insurance in the control. Thus, even though the seller is not informed of the buyer's insurance purchase decision, she anticipates that the buyer will not purchase the insurance. As a result, there exists an *N*-pooling equilibrium: both types of sellers advise *N*, all buyers purchase the product without insurance, and both types of sellers ship the product. The *Y*-pooling equilibrium can be again eliminated using the same forward induction argument as in the IA.

Now consider the case when $g < 1 - \frac{w}{p}$. In this case, there can be multiple equilibria in the subgame after sellers advise *N*: (i) the buyer does not buy the product and the type-s seller does not ship the product upon receiving orders, (ii) the buyer buys the product without insurance and the type-s seller ships the product upon receiving orders, and (iii) the buyer mixes between buying the product with and without insurance, while the type-s seller mixes between shipping and not shipping the product. Specifically, in (iii), the buyer will buy the product without insurance with a probability of $\frac{d}{\alpha}$, while the type-s seller will ship the product with a probability of $1 - \frac{w}{p(1-g)}$. This mixed-strategy equilibrium exists given our assumption that the insurance price is relatively low, i.e., $w \le p\left(1 - \frac{p}{v}\right)$. We select the mixed-strategy equilibrium that, unlike other equilibria, does not require extreme forms of participants' (mis-)coordination when the insurance purchase decisions are unobservable.[8] In this mixed-strategy equilibrium, the proportion of each strategy realization is endogenously determined and the equilibrium outcome can be used to make comparisons between the control and IA. As in the IA scenario, we will focus on the equilibrium in which both types of sellers advise *N*. If all sellers advise *Y*, there will be no mixed-strategy equilibrium in the subsequent subgame as the type-s seller will never ship the product.

By comparing the equilibrium in all three scenarios, we derive the following hypotheses. Note that our theoretical results do not qualitatively depend on the assumption that there exist type-g sellers who always ship the product. The insurance advice mechanism works by forcing type-s sellers to commit to shipping the product, which does not depend on the existence of type-g sellers. We discuss the details in Appendix A4.

Hypothesis 1. In the scenarios with the insurance advice mechanism, sellers will advise *N*.

As long as α is sufficiently large ($\alpha \ge d$) for some sellers:

[7] The details of the equilibrium refinement using forward induction can be found in the Appendix A2.

[8] If sellers are optimistic and believe that buyers will follow the advice *N*, there exists an equilibrium such that when $d \le \alpha$ all sellers advise *N*, buyers purchase the product without purchasing insurances, and all sellers ship the product. Similarly, if sellers are pessimistic and believe that buyers will not follow the advice *N*, there exists another equilibrium such that when $d \le \alpha$ all participants behave as in the scenario without advice mechanism despite all sellers advising *N*. These equilibria are less convincing as they rely on sellers holding extreme beliefs (See Appendix A for more details).

Hypothesis 2. Both the IA and the IA_HI increase the frequency of buyers purchasing the product compared to the control:

$$\text{freq}(\text{BuyProd}|\text{control}) < \text{freq}(\text{BuyProd}|\text{IA\_HI}) = \text{freq}(\text{BuyProd}|\text{IA}).$$

Hypothesis 3. The frequency of buyers purchasing the product without insurance (BuyProd & NoIns) is highest in the IA and lowest in the control:

$$\text{freq}(\text{BuyProd\&NoIns}|\text{control}) < \text{freq}(\text{BuyProd\&NoIns}|\text{IA\_HI}) < \text{freq}(\text{BuyProd\&NoIns}|\text{IA}).$$

Hypothesis 4. The frequency of sellers shipping the product is highest in the IA and lowest in the control:

$$\text{freq}(\text{Ship}|\text{control}) < \text{freq}(\text{Ship}|\text{IA\_HI}) < \text{freq}(\text{Ship}|\text{IA}).$$

These hypotheses imply that the insurance advice mechanism can improve market efficiency. We can measure market efficiency by the frequency of efficient trades. In our setup, there are two types of efficient trades. One is characterized by the number of buyers purchasing the product (regardless of their insurance purchase decisions) and the number of sellers shipping the product. The other is characterized by the number of buyers purchasing the product without insurance and the number of sellers shipping the product. When considering the total welfare of buyers and sellers, this latter definition of efficiency is a Pareto improvement over the first definition, as the buyer's earnings increase while the seller's earnings do not change. [9] To differentiate these two definitions of efficiency, we refer to the first type as *standard efficient trades* and to the second type as *optimal efficient trades*. According to Hypotheses 2, 3, and 4, we expect the frequency of efficient trades to be highest in the IA and lowest in the control. Consequently, the insurance advice mechanism is most effective in the IA, especially when considering the proportion of optimal efficient trades.

Hypothesis 5. The frequency of standard and optimal efficient trades freq(Eff):

$$\text{freq}(\text{Eff}|\text{control}) < \text{freq}(\text{Eff}|\text{IA\_HI}) < \text{freq}(\text{Eff}|\text{IA}).$$

3. Experiment

3.1. Experimental design

To test the hypotheses, we design an experiment with three treatments, corresponding to the three scenarios described in the theoretical analysis. At the beginning of each treatment, subjects were randomly assigned to the role of either buyer or seller. Following Li and Xiao (2014), each treatment consisted of 10 rounds. Repeating the game allowed us to obtain a larger number of observations and provided participants with opportunities to learn to converge to equilibrium. Both buyers and sellers received a full history of their decisions, which was provided and updated at the end of every round (see Appendix B for screenshots of the decision-making stage). To minimize the potential reputation effect, we randomly matched each buyer with a seller at the beginning of each round. At the end of the experiment, one round was randomly selected as the payment round, such that the earnings outcome in one round was unlikely to have any income effect on the decisions in later rounds. The instructions are provided in Appendix C.

In the control treatment (illustrated in Fig. 1), at the beginning of each round, buyers and sellers were endowed with 35 points, and the buyer could choose to purchase a product with insurance, purchase a product without insurance, or not purchase the product. The product value for the buyer $v = 40$ points. The price fo the product $p = 25$ points. Following the previous literature (Li and Xiao, 2014; Lafky, 2014), we set the price of the product to be fixed to exclude the possibility that sellers could use the price to signal their intention to cooperate, which would complicate the study of the advice mechanism. If the buyer decided not to purchase the product, the round ended, and each participant's earnings remained at 35 points. If the buyer decided to purchase the product (with or without insurance), the seller received the payment of 25 points from the buyer and then decided whether to ship the product.[10] Shipping cost $d = 10$ points. Thus, if the seller shipped the product, her earnings for that round were 50 points, and if she did not ship the product, her earnings were 60 points.

The insurance cost $w = 8$ points. If the buyer purchased the product, but the seller did not ship it, the insurance would cover the loss of the 25 points that the buyer had paid to the seller. Once the buyer decided to purchase the insurance, he would pay the cost of 8 points, regardless of whether the seller shipped the product. All these factors were common knowledge.

Our design of the parameter values was guided by the theoretical analysis in Section 2. First, the parameter values satisfied the assumption $w \leq p\left(1 - \frac{p}{v}\right)$, such that at least some buyers would potentially buy the insurance. Second, the parameter values ensured that in the control treatment the buyer's decision (whether that be to purchase the product without insurance, purchase the product

[9] Here we only consider the total welfare of buyers and sellers since the insurance is exogenously provided by the experimenter. In a natural setting, if the insurance market is close to perfectly competitive or insurers serve a large number of markets, a reduced number of insurance purchases in one market only leads to a negligible loss for the insurance company.

[10] Although we used the shipping context (also see Bolton et al. 2004, Li and Xiao 2014), the nature of the decision making, however, can also extend to other settings such as the choice of the quality of the products or the speed of shipping.

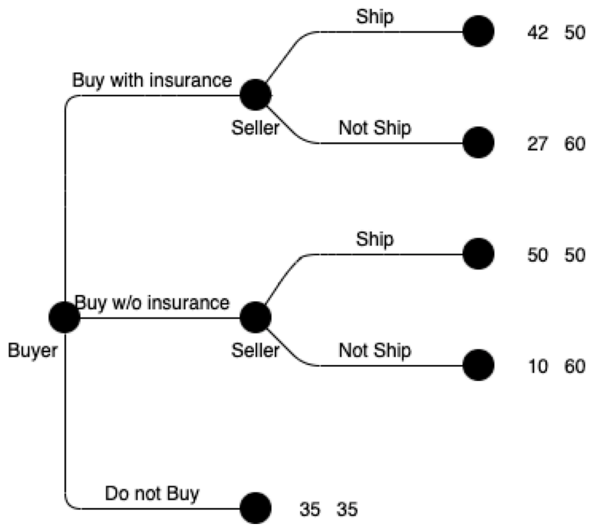

Fig. 1. Buyer-seller game with insurance (control treatment).

with insurance, or not purchase the product) differed depending on their prior belief about the seller's type, g.

The IA treatment—the timing of the game is described in Fig. D1 in Appendix D—is the same as the control treatment, except that we added a stage before the buyer made their product and insurance purchase decisions. At this stage, the seller had to advise the buyer whether to purchase the insurance. The buyer then made his decision after he observed the seller's advice. The seller was informed of the buyer's insurance purchase decision before she made the shipping decision. All this was common knowledge. The rest of the game was the same as the control treatment.

The IA_HI treatment has the same structure as the IA treatment, except that the seller was not informed of the buyer's insurance purchase decision throughout the experiment.

3.2. Experimental procedure

The experiment was conducted at the Monash Laboratory for Experimental Economics (MonLEE) using z-tree (Fischbacher, 2007). The experimenter read the instructions aloud, and the subjects completed a comprehension quiz (see Appendix E) to ensure that they understood the task and the payoffs associated with each decision.

We ran 24 sessions in total—8 sessions per treatment—and we recruited, on average, 14 subjects in each session. Each session lasted less than one hour. Subjects were randomly assigned the role of either buyer or seller and maintained this role for the entirety of the experiment. In each round, a buyer was randomly and anonymously rematched with a seller. At the end of the experiment, one round was randomly selected as the payment round. In total, we recruited 332 subjects: 108 for the control treatment, 116 for the IA treatment, and 108 for the IA_HI treatment. Each subject was paid $4 AUD for participating, adding to the earnings from the games. The exchange rate was 1 point = $0.4 AUD. Subjects were paid privately, earning about $20 AUD on average.

4. Results

We first report what advice sellers gave buyers in the two advice treatments. Then, we compare buyers' purchase decisions. Next, we examine how insurance advice affects sellers' shipping decisions and whether shipping frequency differs between the treatments. Finally, we report the treatment effect on market efficiency. When reporting the findings, we always first report descriptive data followed by results from regression analysis to provide statistical tests for the comparisons. Since in each session we anonymously rematch participants after each round, the regression analysis allows us to control for the potential session effects, by clustering standard errors at the session level (Fréchette, 2012). We also analyze the dynamics of behavior as participants may not play the equilibrium at the beginning of the game, but learn to converge to the equilibrium overtime.

Table 1 summarizes the main descriptive results for each treatment. In total, there are 108 subjects in the control, 116 subjects in the IA treatment, and 108 subjects in the IA_HI treatment.

4.1. Insurance advice

Supporting Hypothesis 1, over the 10 rounds, we observe a large proportion of sellers advised N in both the IA (81.2 %) and the IA_HI treatment (71.1 %). Fig. 2 plots the proportion of sellers who advised N in each round. In both treatments, the frequency of advising N is relatively lower in the first round and increases over time.

To provide statistical evidence for the treatment differences, we analyze sellers' insurance advice decisions over time using a

Table 1
Descriptive summary of decisions.

Treatment	Sellers advised N (%)	Buyers purchased product (%)	Buyers purchased product without insurance (%)	Sellers shipped product (%)	Standard efficient trades (%)	Optimal efficient trades (%)
Control	-	51.1	20.2	43.9	22.4	9.1
IA	81.2	66.9	48.3	61.4	42.1	33.3
Advice *N*	-	74.2	59.4	64.2	49.9	41.0
Advice *Y* *	-	35.0	2.7	25.0	8.3	0
IA_HI	71.1	67.8	35.6	53.5	37.8	22.8
Advice *N* *	-	73.3	43.6	58.0	46.9	29.9
Advice *Y**	-	52.0	13.6	26.2	15.4	5.1

Note: There are in total 54 buyers/sellers in the control; 58 buyers/sellers in IA and 54 buyers/sellers in IA_HI.
* The number of subjects in these three cases are smaller than the whole sample because some sellers never advised *N* or never advised *Y*, and some buyers never received advice *N* or never received advice *Y*. Specifically, in the IA treatment, 27 sellers never advised *Y* and 9 buyers never received advice *Y*, leaving 49 buyers and 31 sellers in the Advice *Y* condition. In the IA_HI treatment, 4 sellers never advised *N* and 20 sellers never advised *Y*, leaving 54 buyers and 50 sellers in the Advice *N* condition, and 54 buyers and 34 sellers in the Advice *Y* condition.

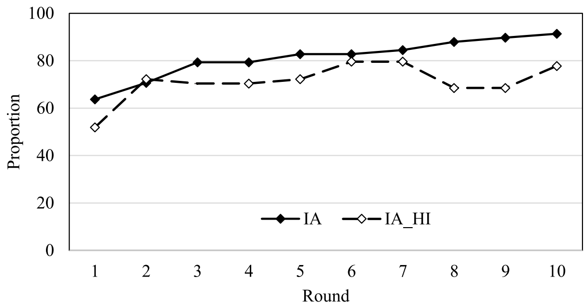

Fig. 2. Proportion of sellers who advised *N*.

random-effects linear probability.[11] We report the results in Table 2 below. The dependent variable is whether the seller advised *N* or *Y* in each round. Regressions (1) and (2) compare the IA treatment to the IA_HI treatment. The independent variable in Regression (1) includes only the IA treatment dummy variable. We find the coefficient of IA (β_1) is significantly positive, meaning that sellers in the IA treatment are significantly more likely to advise *N* than those in the IA_HI treatment. In Regression (2), we add the independent variables Round and IA*Round. We find that as the rounds progressed, sellers were significantly more likely to advise *N* in both the IA treatment ($\beta_2 + \beta_{3,}\ p = 0.001$) and in the IA _HI treatment($\beta_{2,}\ p = 0.032$). We report in the next section that a buyer was more likely to purchase the product when he received advice *N* compared to when he received advice *Y*. The increasing rate of advising *N* suggests that sellers gained experience and learned to advise buyers *N* over time.

Result 1. **:** *In both the IA and IA_HI treatments, the majority of sellers advised N. The frequency of advising N increases over time in both treatments.*

4.2. Purchase decision

Supporting Hypothesis 2, the purchase rate is higher in both the IA (66.9 %) and IA_HI (67.8 %) treatments compared to the control treatment (51.1 %). Fig. 3 plots the product purchase decision conditional on the advice that buyers received. As buyers did not receive any advice in the control treatment, we use the dotted line to mark the average product purchase rate in the control. Compared with the control, buyers in the two advice treatments were more likely to purchase the product when the sellers advised *N*. (IA: 74.2 %; IA_HI: 73.2 %; Control: 51.1 %). By contrast, when buyers received advice *Y* in the IA treatment, the purchase rate was lower than in the

[11] For all the regressions reported in this paper, we cluster standard errors at the session level. We also test the robustness of the results by (1) including control variables that include gender, major, and how well the individual understood the experimental instructions. All the main results are robust when including these controls (see Appendix F). We also conduct probit regressions, and all results are robust (see Appendix H).

Table 2
Random individual effects LPM regression analysis of insurance advice decisions.

Independent variables	Dependent variable: Advice $N_{i,t}$ =1, if seller *i* advised *N* in round *t* =0, *o.w.* (1)	(2)
β_1: IA	0.101**	0.032
	(0.040)	(0.084)
β_2: Round		0.014**
		(0.007)
β_3: IA_Round		0.013
		(0.011)
Constant	0.711***	0.633***
	(0.023)	(0.046)
H0: $\beta_2 + \beta_3$=0		p = 0.001
N	1,120	1,120

Note: IA_HI is the constant. Robust standard errors clustered at the session level are reported in parentheses.
*** $p < 0.01$,
** $p < 0.05$, * $p < 0.1$

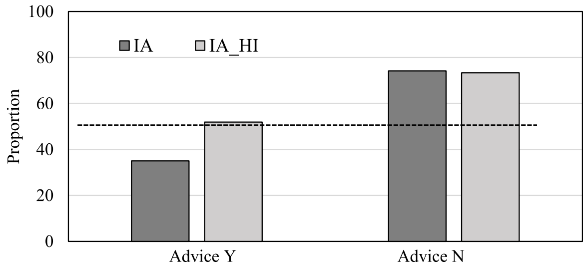

Fig. 3. Product purchase rate conditional on advice. Note: The dotted line marks the purchase rate (51.1 %) in the control treatment.

control (IA: 35.0 %; Control: 51.1 %). In the IA_HI treatment, the purchase rate when advice is Y is about the same as that in the control (51.9 %). These results suggest that the increase in the product purchase rate in the IA and IA_HI treatments is mainly driven by buyers who received advice *N*. When the seller advised *N*, the buyer's expectations of the seller shipping the product increased, and consequently, were more likely to purchase the product.

It is interesting to observe that upon receiving advice *Y*, 35 % of buyers in the IA and 51.9 % of buyers in the IA_HI treatment purchased the product when in theory, they should not. However, the sample size is small; thus, caution must be observed when drawing inferences from these observations. Nevertheless, we make the following two notes. In the IA treatment, the proportion of buyers who purchased the product after receiving advice *Y* is relatively higher in the earlier rounds: 47.2 % in the first five rounds and 13.5 % in the last five rounds, with no buyers choosing to purchase the product in the final two rounds. This result suggests that buyers learned not to purchase the product when they received advice *Y* as the rounds progressed. In the IA_HI treatment, however, we still observe that 20 % of buyers purchased the product in the last round. The three buyers who purchased the product in the IA_HI treatment when the seller advised *Y* appeared to be more cautious as all of them also decided to purchase insurance. By contrast, in the control treatment, 5 out of 17 buyers who bought the product in the last round did not purchase the insurance.

Next, we compare the dynamics of the product purchase decisions in each treatment. Fig. 4(a) plots the proportion of product purchases over 10 rounds, while 4(b) and 4(c) plot the same proportion when the advice was *N* or *Y*, respectively. Although we observe a rapid decay in the purchase proportion in the control treatment, the decay is relatively slower in the IA and IA_HI treatments. Upon separating the cases based on sellers' insurance advice, we find that the decay is slower only when the sellers advised *N*.

To provide statistical evidence for the comparisons reported above, we analyze the buyer's product purchase decisions (including when the advice was *Y*) using a random-effects linear probability model. We report the results in Table 3 below. The dependent variable is whether the buyer purchased the product in each round. Regressions (1)–(3) compare the IA treatment to the control. The independent variable in Regression (1) only includes the IA treatment dummy variable. We find the coefficient of IA (β_1) is significantly positive, meaning that buyers in the IA treatment are significantly more likely to purchase the product than those in the control. In Regression (2), the independent variables include the treatment dummy, round, and the interaction between the treatment and round. The coefficient of "Round" is negative and statistically significant, indicating that the product purchase rate decays significantly over time in the control treatment. There is also significant decay in the IA treatment ($\beta_3 + \beta_4$, $p = 0.026$). However, consistent with the observation in Fig. 2, the IA treatment slows down the decay as compared to the control, the coefficient of the interaction variable IA*Round (β_4) is positive and marginally significant ($p = 0.093$). Regression (3) further includes the dummy variable Advice *N* to control for advice and the interaction variable Advice *N**Round. We find that the reduced decay rate in the IA treatment is mainly driven by sellers advising *N*. To see this, note that IA*Round is not statistically significant (and the direction is negative), indicating

(a) Proportion of buyers who purchased the product per round (total)

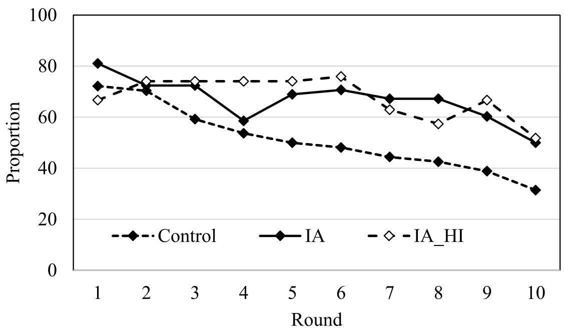

(b) Proportion of buyers who purchased the product per round (Advice *N*)

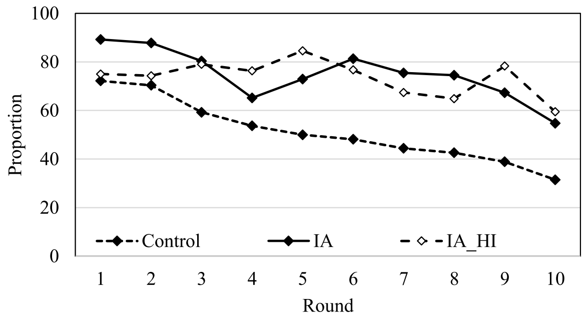

(c) Proportion of buyers who purchased the product per round (Advice *Y*)

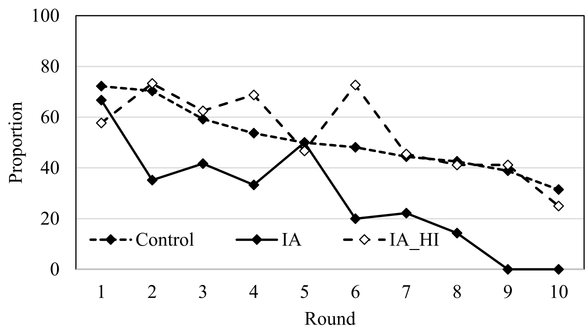

Fig. 4. Proportion of buyers who purchased the product per round.

that advising Y did not slow down the rate of decay. By contrast, Advice *N* and the interaction variable Advice *N**Round are both positive and statistically significant.

Regressions (4)–(6) provide similar analyses of the IA_HI treatment compared to the control. The coefficient of IA_HI (β_2) is significantly positive in Regression (4), meaning that buyers in the IA_HI treatment are significantly more likely to purchase the product than those in the control. Regression (5) shows that although there is still a significant rate of decay in IA_HI ($\beta_3 + \beta_5$, $p =$ 0.012), it is slower in the IA_HI treatment compared to the control (β_5 is significantly positive). Similar to the findings of the IA treatment, when we control for the advice in Regression (6), we find that the reduction of the decay rate in the IA_HI treatment is significant only when sellers advised *N* (β_5 is negative and not significant; β_7 is significantly positive).

These results are consistent with our theoretical analysis, in which, for both the IA and IA_HI treatments, advice *N* leads to a higher product purchase rate compared to the control treatment.

We next tested Hypothesis 3 by comparing the proportion of buyers that purchased the product without insurance across treatments. Assuming that sellers shipped the product, buyers achieved the highest earnings in this scenario. Supporting Hypothesis 3, the

Table 3
Random individual effects LPM regression analysis of product purchase decisions.

Independent variables	Dependent variable: $Buy_{j,t}$ =1, if buyer j purchased the product in round t = 0, o.w. (1) IA and Control	(2) IA and Control	(3) IA and Control	(4) IA_HI and Control	(5) IA_HI and Control	(6) IA_HI and Control
β_1: IA	0.158**	0.045	-0.130			
	(0.070)	(0.081)	(.019)			
β_2: IA_HI				0.167*	0.033	0.0111
				(0.089)	(0.089)	(0.087)
β_3: Round		-0.042***	-0.042***		-0.042***	-0.042***
		(0.007)	(0.007)		(0.007)	(0.007)
β_4: IA*Round		0.020*	-0.018			
		(0.012)	(0.149)			
β_5: IA_HI*Round					0.024**	-0.001
					(0.010)	(0.014)
β_6: Advice N			0.289***			0.042
			(0.079)			(0.058)
β_7: Advice N *Round			0.033***			0.033**
			(0.009)			(0.017)
Constant	0.511***	0.744***	0.744***	0.511***	0.744***	0.744***
	(0.065)	(0.068)	(0.068)	(0.065)	(0.068)	(0.068)
H0: $\beta_3 + \beta_4$=0		$P = 0.026$	$P < 0.001$			
H0: $\beta_3 + \beta_5$=0					$P = 0.012$	$P < 0.001$
N	1120	1120	1120	1080	1080	1080

Note: Advice $N = 1$ if the seller advised N; = 0, o.w. Robust standard errors clustered at the session level are reported in the parentheses,
*** $p < 0.01$
** $p < 0.05$
* $p < 0.1$.

proportion of buyers that purchased the product without insurance is highest in the IA treatment and lowest in the control treatment. The order is significant (IA: 48.3 %; IA_HI: 35.6 %; Control: 20.2 %; Jonckheere–Terpstra test, $p < 0.011$).[12] We report the dynamics of this proportion over the 10 rounds in each treatment in Fig. G1 in Appendix G. As shown in Fig. G1, the order is very similar over the 10 rounds. This result suggests that one benefit of the IA mechanism is that buyers saved expenses on insurance (without increasing the risk of losing payment when encountering strategic sellers, as reported below).

Fig. 5 further shows the proportion of buyers who purchased the product without insurance, conditional on the advice received. Since there was no advice opportunity in the control treatment, we mark the average proportion of product purchases without insurance by the dotted line. As shown in Fig. 5, compared with the control treatment, the proportion of buyers who purchased the product without insurance after receiving advice N in both the IA and the IA_HI treatments more than doubled (IA: 59.4 % IA_HI: 43.6 %; Control: 20.2 %). To test the treatment differences, we conducted a random-effects linear probability regression analysis of the insurance purchase decisions for IA and IA_HI, respectively. In each regression, we include only the treatment dummy variable (IA or IA_HI) and use the control as the baseline. In both regressions, the coefficient of the treatment dummy is significantly positive (IA: 0.39, $p = 0.001$; IA_HI: 0.24, $p < 0.001$). By contrast, when receiving advice Y, fewer buyers purchased the product without insurance in the IA treatment (3.7 %) and the IA_HI treatment (13.6 %) compared to the control treatment (20.2 %,).[13] These results show that the lower insurance purchase rates in the IA and the IA_HI treatments reported above are due to buyers receiving advice N.

Result 2. **:** *Buyers were more likely to purchase the product in the IA and IA_HI treatments than in the control treatment. This increase in the product purchase rate is mainly driven by advice N.*

Result 3. **:** *Buyers were more likely to purchase the product without insurance in the two advice treatments than in the control treatment. The increase is mainly driven by the effect of advising N.*

4.3. Shipping decision

Supporting Hypothesis 4, the overall shipping rate is highest in the IA and lowest in the control treatment (IA: 61.4 %; IA_HI: 53.5 %; Control: 43.8 %, Jonckheere–Terpstra test, $p = 0.003$). The increase in the shipping rate is mainly driven by advising N (IA (advice N): 64.2 %; IA_HI (advice N): 58.0 %; Control: 43.9 %). By contrast, when sellers advised Y, the shipping rate in the two advice treatments is lower than in the control (IA (advice Y): 25.0 %; IA_HI (advice Y): 26.2 %; Control: 43.9 %). Fig. 6 plots the shipping rate

[12] For all Jonckheere-Terpstra tests, we calculate the average at the session level and use each session as an independent observation ($n = 8$ for each treatment respectively, $n = 24$ in total).
[13] We also tested the treatment differences using a similar regression analysis as above. We find the coefficient of the IA dummy is significantly negative (IA: -0.17, $p < 0.001$) and the coefficient of the IA_HI dummy is negative, but not significant (IA_HI: -0.06, $p = 0.218$).

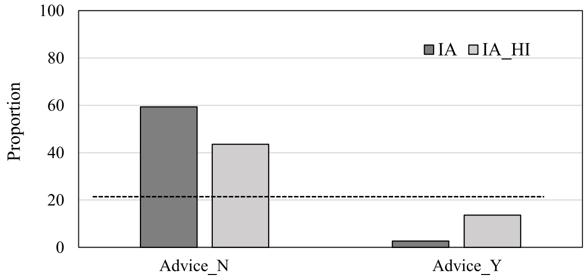

Fig. 5. Proportion of buyers who purchased the product without insurance conditional on advice. Note: The dotted line marks the product purchase rate (20.2 %) in the control treatment (no advice was given in the control).

over the 10 rounds. Fig. 6(a) reports the overall shipping rate. We observe that the overall shipping rate is highest in the IA treatment in 7 out of the 10 rounds. By separating the cases into advice *N* (Fig. 6(b)) and advice *Y* (Fig. 6(c)), we again find that the higher shipping rate in the IA treatment is mainly driven by advising *N* throughout the experiment. These results are consistent with our theoretical framework, in which the seller incurs a psychological cost for not shipping the product after she advised *N*.

In our theoretical framework, we assume that the psychological cost is higher when the seller observes that the buyer follows her advice of *N* than when the buyer does not follow her advice. This assumption would predict that, in the IA treatment, sellers who advise *N* would be more likely to ship the product after observing the buyer who purchased the product without insurance than buyers who purchased the product with insurance. To check this, Fig. 7 reports the average shipping rates of sellers who advised *N* in the IA treatment when buyers purchased the product with and without insurance. For comparison, the dotted lines mark the average shipping rate when buyers purchased the product with and without insurance in the control treatment. We also include data from the IA_HI treatment. As sellers never knew the buyer's insurance purchase decision in the IA_HI treatment, we do not expect to see any correlation between the sellers' shipping decisions and the buyers' insurance purchase decision. Our data are consistent with the prediction. The average shipping rate in the IA treatment is higher when the buyer did not purchase the insurance (67.8 %) than when the buyer purchased insurance (56.2 %). The difference is smaller in the IA_HI treatment and in the control.

To provide statistical evidence for the above comparisons, we conduct a regression analysis of the sellers' shipping decisions using a random effects linear probability model. For each of the advice treatments, we start with a regression that only includes the treatment dummy (IA or IA_HI) (Regressions 1 and 4). Results are reported in Table 4. the results from Regression (1) show that, in the IA treatment, sellers are significantly more likely to ship than those in the control (β_1 is statistically significant). However, in Regression (4) while we observe the direction is higher in IA_HI as compared to the control: β_2 is positive but not significant.[14] Next, we add another independent variable "Advice *N*" = 1 if the seller advised *N* (Regression 2 and 5). Results from the two regressions show that the insurance advice mechanism of advising N is what drives the increase in shipping rates (β_6 is statistically significant in both treatments).

To test the effect of buyers' insurance purchase decisions, in Regressions (3) and (6) we add the independent variable "Noinsure"=1 if the buyer purchased the product without insurance and the interaction of the Noinsure variable with the dummy treatment variable. As our hypothesis on the psychological cost is only for sellers who advised *N*, we do not include data from the advice treatments when the advice was *Y*. [15] Regression (6) shows that the shipping rate is not affected by the insurance purchase decision in the control treatment (β_3 is not statistically significant), or in the IA_HI treatment ($\beta_3 + \beta_4$, $p = 0.204$). In contrast, results from Regression (3) suggest that, in the IA treatment, if the buyer purchased the product without insurance after receiving advice N, sellers are significantly more likely to ship the product ($\beta_3 + \beta_4$, $p<0.001$).

In summary, our data suggest that advice *N* provided by the sellers causes an increase in the shipping rate. Furthermore, for sellers, the buyers' insurance purchase decisions did matter when they advised buyers *N*, as long as they could observe whether the advice was followed.

We also explore individual differences in shipping behavior. For each seller, we calculate the frequency of shipping the product when the paired buyer decided to purchase the product. The distribution of the shipping rate in each treatment is shown in Appendix G (Fig. G2). In all treatments, the two most common behavior profiles are to never or always ship. Table 5 summarizes the proportion of sellers who "always shipped" and "never shipped" in each treatment. In the two advice treatments, we report the proportions for the

[14] We also ran a similar regression model using only data from IA and IA_HI with IA_HI being the baseline, we find that the coefficient of the IA dummy is positive but not significant (IA: 0.08, $p = 0.232$)

[15] Similar regression analysis using only data from the advice treatments when the advice was *Y* shows that in the IA treatment sellers were significantly less likely to ship the product when buyers did not follow the advice *Y* and did not purchase insurance than the control ($\beta_3 + \beta_4 =$ -0.146, $p = 0.028$). There is no significant difference in shipping rates between the control and the IA_HI treatment when the buyer purchased the product without insurance ($\beta_3 + \beta_4 = 0.100$, $p = 0.573$).

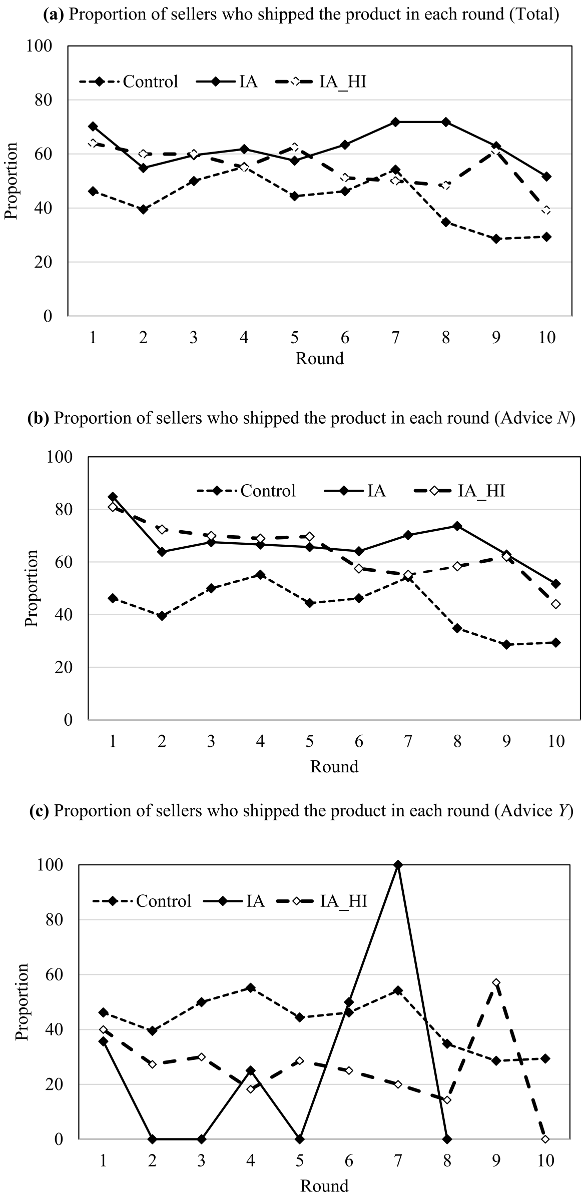

Fig. 6. Proportion of sellers who shipped the product in each round. Note: In rounds 9 and 10 of the IA treatment, there were no observation in which the seller advised *Y* and the buyer purchased the product.

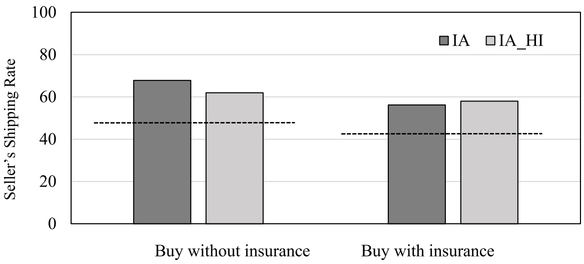

Fig. 7. Shipping rates when sellers advised *N*. Note: The dotted line marks the shipping rate in the control treatment when the buyer did not purchase an insurance (48.4 %, on the left) and when the buyer purchased insurance (42.0 %, on the right).

Table 4
Random individual effects LPM regression analysis of shipping decisions.

Independent variables	Dependent variable: $Ship_{i,t}$ =1, if the seller *i* shipped the product in round *t*; =0, o.w.					
	(1) IA and Control	(2) IA and Control	(3) IA (advice *N*) and Control	(4) IA_HI and Control	(5) IA_HI and Control	(6) IA_HI (advice *N*) and Control
β_1: IA	0.179*** (0.053)	-0.062 (0.073)	0.096 (0.066)			
β_2: IA_HI				0.101 (0.0063)	-0.059 (0.073)	0.159** (0.078)
β_3: Noinsure			0.046 (0.058)			0.046 (0.058)
β_4: Noinsure*IA			0.121* (0.064)			
β_5: Noinsure*IA_HI						-0.018 (0.063)
β_6: Advice *N*		0.273*** (0.061)			0.216*** (0.074)	
Constant	0.436*** (0.036)	0.436*** (0.036)	0.419*** (0.049)	0.437*** (0.036)	0.437*** (0.049)	0.419*** (0.036)
H0: $\beta_3 + \beta_4 = 0$			$P < 0.001$			
H0: $\beta_3 + \beta_5 = 0$						$p = 0.204$
N	664	664	625	642	642	558

Note: Noinsure = 1 if the buyer did not purchase insurance; = 0, o.w. Advice *N* = 1 if the seller advised *N*; = 0, o.w. Robust standard errors clustered at the session level are reported in parentheses.
*** $p < 0.01$
** $p < 0.05$
* $p < 0.1$

Table 5
Frequency of sellers who either always or never shipped the product.

Treatment (# of sellers)	Always Ship (%)	vs. Control (p-value)	Never Ship (%)	vs. Control (p-value)
Control (54)	18.5	-	33.3	-
IA (58)	32.8	0.053	17.2	0.035
Advice_*N* (58)	41.4	0.002	19.0	0.073
Advice_*Y* (20)	10.0	0.417	60.0	0.034
IA_HI (54)	24.0	0.460	24.1	0.292
Advice_*N* (47)	34.0	0.049	27.7	0.567
Advice_*Y* (25)	4.0	0.001	48.0	0.167

Note: For Advice_*N* (Advice_*Y*), we only considered the seller's shipping decisions when the advice was *N* (*Y*) and a buyer purchased the product. The p-value was based on an LPM regression analysis with standard errors clustered at the session level, where the independent variable is either a treatment dummy, or an advice dummy.[19]

case when sellers advised *N* and *Y*, respectively.[16]

Our theoretical framework predicts that a type-s seller who never ship in the control treatment will advise *N* and subsequently always ship in the IA treatment, provided that the psychological cost of not shipping is sufficiently large. In the IA_HI treatment, there is a mixed strategy equilibrium where type-s sellers advise *N*, but only some will always ship the product. Thus, compared to the control treatment, we expect to see a higher proportion of "always ship" and a lower proportion of "never ship" when sellers advise *N* in the IA treatment. The effect of advice, albeit positive, is weaker in the IA_HI treatment. We thus compare the proportion of "always ship" and "never ship" when sellers advise *N* in the two advice treatments with the control treatment.

Supporting the theoretical framework, we find that in IA and IA_HI, when the advice was *N*, the proportion of "always ship" is highest in the IA treatment and lowest in the control, and the order is statistically significant (41.4 % >34.0 % >18.5 %, Jonckheere–Terpstra test, $p = 0.003$). Similarly, the proportion of "never ship" is highest in the control and lowest in the IA treatment, and

[16] Calculating the two types using the pooled data in the two advice treatments can be misleading. Take an extreme case as an example. Suppose all sellers advise *N* in 8 rounds and advise *Y* in 2 rounds. Also suppose all sellers always ship when advising *N* and never ship when advising *Y*. In this case, when calculating the overall proportion of "always ship" and "never ship", we will get 0% of both "always ship" and "never ship" in the advice treatments.

the order is also statistically significant (33.3 >27.7 % >19.0 %, Jonckheere–Terpstra test, $p = 0.033$).[17][18]

Result 4. **:** *Sellers were more likely to ship the product in the two advice treatments than in the control. The increase in the shipping rate is driven by those who advised N.*

Result 5. **:** *In the IA treatment, sellers who advised N were more likely to ship the product when the buyer followed the advice than when he did not follow the advice.*

4.4. Market efficiency

Fig. 8 plots the proportion of standard efficient trades and optimal efficient trades in each treatment. As defined in Section 2, standard efficient trades occur when the product is bought and shipped, and optimal efficient trades occur when the product is bought without insurance and shipped. Fig. 8 supports Hypothesis 5 in that the frequency of standard efficient trades is highest in the IA and lowest in the control treatment (IA: 42.1 %; IA_HI: 37.8 %; Control: 22.4 %, Jonckheere–Terpstra test, $p = 0.002$). and we observe the same pattern for optimal efficient trades (IA: 33.3 %; IA_HI: 22.8 %; Control: 9.1 %, Jonckheere–Terpstra test, $p < 0.001$).

We also conduct a random-effects linear probability regression analysis of efficient trades to test the effect of IA and IA_HI, respectively. In each regression, we include only the treatment dummy variable (IA or IA_HI) and use the control as the baseline. We find the coefficient for each treatment dummy is significantly positive (standard efficient trades: IA: 0.20, $p = 0.001$; IA_HA: 0.15, $p = 0.0029$; optimal efficient trades: IA: 0.24, $p < 0.001$; 0.14, $p = 0.012$), indicating that the increase in the efficient trades in both treatments is statistically significant. In addition, we compare the efficient trades between IA and IA_HI treatments using a similar regression analysis where we include only the IA dummy and use IA_HI as the baseline. We find the coefficient of IA is positive for both standard efficient trades (0.04, $p = 0.531$) and optimal efficient trades (0.10, $p = 0.093$) although only the latter is marginally significant.

As a result of the increase in the number of efficient trades, both buyers and sellers made higher profits in the two insurance advice treatments than in the control. The average earnings per round for the buyer increase by 9.1 % in the IA treatment (36.1 points) compared to the control (33.1 points). The average earnings per round for the seller increase by 4.4 % in the IA treatment (47.5 points) compared to the control (45.5 points). We observe similar results in the IA_HI treatment. Compared to the control, the average earnings per round for the buyers increase by 5.4 % (34.9 points), and the average earnings per round for the seller increase by 5.9 % (48.2) in the IA_HI treatment.

Result 6. **:** *The insurance advice mechanism significantly increases the frequency of efficient trades, especially when buyers' insurance purchase decisions were known to the sellers.*

5. Discussion and conclusion

We designed and tested a novel insurance advice mechanism aimed at promoting efficient trades in a market with asymmetric information. We show both theoretically and experimentally that under this mechanism, buyers purchased the product significantly more often, and sellers were also more likely to ship the product than in the control treatment. The insurance advice mechanism also has an indirect welfare effect on buyers by reducing the frequency of purchasing insurance. The comparison between the two advice treatments further shows that the mechanism is most effective when sellers could observe buyers' insurance purchase decisions. This finding suggests that online marketplaces may want to make buyers' insurance purchase decisions salient to sellers, alongside the introduction of the insurance advice mechanism.

Our study points out a new direction for designing market mechanisms to overcome asymmetric information problems. Existing instruments, such as reputation mechanisms, warranties, and insurance, designed to promote market efficiency, are often costly to sellers and/or buyers (Li and Xiao, 2014; Bolton et al., 2018; Bolton et al., 2019; Andreoni, 2018). Although some big companies or manufacturers can provide insurance or warranties, many small sellers that populate online marketplaces, such as eBay and Amazon, cannot do so. Sellers can sometimes use prices to signal their quality. However, this type of signaling can be quite costly: good sellers have to distort their prices to differentiate themselves from bad sellers (Bagwell and Riordan, 1991, Bagwell, 1992). Importantly, price signaling does not change the non-cooperative behavior of strategic sellers.

The insurance advice mechanism can be a low-cost complement to these existing instruments. A simple message on the seller's online page can do the job. With the rapid growth of the digital economy, more third-party insurance companies—such as Squaretrade

[17] Sellers behaved very differently when they took the off-equilibrium strategy of advising *Y*. The proportion of "always ship" is highest in the control treatment and lowest in the IA_HI (Control: 18.5%; IA: 10%; IA_HI: 4%). For the proportion of "never ship", it is highest in the IA and lowest in the control (IA: 60%; IA_HI: 48.6%; Control: 33.3%)

[18] In the IA treatment, we find that 11 out of 58 sellers never shipped when advising N. Among these 11 sellers, 5 only advised N. Among the other 6 sellers who sometimes advised N and sometimes advised Y, 5 of them never shipped and the other one shipped 50% of the time after advising Y. Thus, we may argue at most, 10 out of 58 (17%) sellers had $d > \alpha$. This indicates that there is a small proportion of sellers who had $d > \alpha$.

[19] When comparing between IA and IA_HI using similar regression analysis with IA_HI being the baseline, we find the sign of the coefficient of the IA dummy treatment is consistent with our hypothesis, although it is not significant: Always ship (IA: 0.07, p=0.338; when including only advice N: IA: 0.06, p=0.368); Never ship (IA: -0.09, p=0.326; when including only advice *Y:* IA: 0.12, p=0.368).

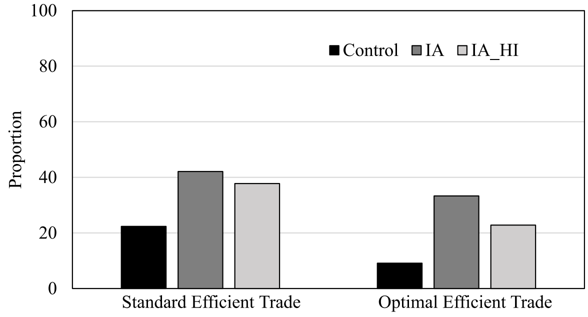

Fig. 8. Proportion of efficient trades by treatments.

and xcover. com—have emerged to protect consumers against risks not covered by the manufacturer's warranty. These insurance products provide a natural opportunity to introduce the advice mechanism with minimal changes to the current platform design. It is worth noting that in this paper, the benefits of the advice mechanism are measured using the outcomes in a market that already has third-party insurance as the reference point. Our experiment cannot speak for whether a market without third-party insurance can benefit from introducing "insurance" and "advice" mechanisms simultaneously. A mechanism that requires adding both insurance and advice is surely not as simple.

One may argue that even in a market with insurance, it might still be sufficient to simply advise the buyers to purchase the product. We make two remarks on this: First, the fact that sellers are selling their products online or in stores is probably equivalent to advising buyers to purchase the product. That is, the advice to buy the product is probably already embedded in the market and does not provide any new information (e.g., the quality of the product as argued in our paper). Advising whether consumers should purchase the insurance from a third party, in contrast, provides new information.

Second, in our framework, we assume that under the advice mechanism, sellers' psychological cost of not shipping depends on whether buyers follow the advice not to buy insurance. Consistent with this assumption, data from our experiment show that one advantage of the insurance advice mechanism is that, in addition to promoting more transactions, it increases buyers' welfare by saving the cost of purchasing insurance—as observed in the comparisons between the IA treatment and the control. If the seller's advice is simply to buy the product, buyers can follow the advice by purchasing the product with or without the insurance. Thus, we expect that advice to purchase the product will be less likely to help buyers save the cost of buying insurance. Future studies could test these hypotheses and compare the effects of the advice mechanism with other forms of communication, such as bare promises (both structured and free forms) or direct advice to buy products.

A potential barrier to implementing the advice mechanism is whether third-party insurance companies will agree to it. One concern may be that third-party insurance companies will be worse off if fewer people purchase insurance. However, it is unclear whether the introduction of the insurance mechanism necessarily makes the insurance company worse off. In the Advice_HI treatment, on average, more products were purchased with insurance (32.2 %) than in the control treatment (30.1 %). This means that if there is no observability, insurance advice may result in companies like Squaretrade being better off (or at least not worse off) because more buyers participate in the market. Finally, it is the online platform that would implement the mechanism, not the third-party insurance company. Assuming that online platforms such as eBay benefit from more buyers purchasing products, we should expect that there is a potential incentive to introduce the advice mechanism.

Future research could be valuable by examining factors that may impact the mechanism's effectiveness. A critical condition for the insurance advice mechanism to be effective is that the psychological cost of not shipping after advising no insurance is greater than the cost of shipping. In our experiment, the seller's shipping cost was relatively small.[20] It would be interesting to explore whether the psychological cost is sufficient to discourage non-cooperative behavior in markets where the cost of cooperation is much greater (e.g. when products are expensive). Previous studies on lying aversion have shown that although people are more likely to lie when they have more to gain, they are also less inclined to lie when the other person has more to lose (Gneezy, 2005; Lundquist et al., 2009). If the psychological costs increase with these gains from non-delivery, the advice mechanism may remain effective, even for relatively large-value transactions.

To avoid confounding effects in the experiment, we exogenously set a constant price for the product. However, in real-world transactions, sellers can endogenously set the price. Furthermore, this price itself may act as a signal to the buyer. It is interesting to consider whether and how sellers would use the insurance advice mechanism when they can also choose the product's price. The insurance advice mechanism can effectively convince buyers to purchase the product at a fixed price under certain conditions, as we have shown in this paper. Hence, the crucial question here is whether the necessary conditions are satisfied at the price that maximizes sellers' profit when buyers believe the sellers always ship. If the answer is yes, sellers may simply rely on the costless insurance advice mechanism to signal their intention of shipping the product, because any adjustment of the price will lead to a lower overall profit.

[20] Although the cost of shipping is only 10 points ($4 AUD), previous studies have shown that people are willing to incur a high cost to cooperation. For example, Abeler et al. (2019) find that people are willing to act honestly at a cost as high as $50 USD,

That is, the insurance advice mechanism can act as a substitute to price signals. The interaction between the price strategy and the insurance advice mechanism is more complicated when the necessary conditions are not satisfied. In certain circumstances, the seller may be able to adjust the price to satisfy the relevant conditions and restore the advice mechanism. In this case, price and insurance advice can be viewed as complements in persuading buyers. However, it is also possible that the price adjustment is too costly or ineffective in restoring those conditions. In this case, only the price may be used to persuade buyers, making it a substitute to the advice mechanism. The exact equilibrium interplay between price and insurance in determining buyers' beliefs is beyond the scope of this paper, and we leave it to future research.

In our experiment, there was no uncertainty associated with receiving compensation under the insurance policy. However, asymmetric information problems are also common in the insurance market. In particular, consumers are often unsure about the coverage. It would be valuable to examine the relationship between the effect of the advice mechanism and trust in the insurance policy. Although we forced the sellers to provide advice in our experiment, it may be more feasible to introduce a mechanism in which insurance advice is presented as an option or upon request from the buyers and sellers can choose to be silent. It would be fruitful to compare the effect of the advice mechanism when advising is mandatory as opposed to when it is optional.

Data availability

All replication files are available on the OSF link attached

Acknowledgments

The authors thank Arthur Campbell and Joshua Miller, as well as seminar and conference participants at the University of Queensland's BESC e-seminar, M-BEEs/M-BEPs 2019 conference, Virtual East Asia Experimental and Behavioral Economics Seminar, East China Normal University, and ANZWEE 2019 conference for valuable feedback and comments. The authors would also like to thank two anonymous referees and the editor for their valuable feedback, comments, and suggestions throughout the review process. Erte Xiao and Chengsi Wang gratefully acknowledge financial support from the Australian Research Council Discovery Project DP170101167 and DP210102015 respectively.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.euroecorev.2023.104586.

References

Abeler, J., Becker, A., Falk, A., 2014. Representative evidence on lying costs. J. Public Econ. 113, 96–104.
Abeler, J., Nosenzo, D., Raymond, C., 2019. Preferences for truth-telling. Econometrica 87 (4), 1115–1153.
Andreoni, J., 2018. Satisfaction guaranteed: when moral hazard meets moral preferences. Am. Econ. J. Microecon. 10 (4), 159–189.
Balafoutas, L., Sutter, M., 2017. On the nature of guilt aversion: insights from a new methodology in the dictator game. J. Behav. Exp. Financ. 13, 9–15.
Bagwell, K., 1992. Pricing to signal product line quality. J. Econ. Manag. Strategy 1 (1), 151–174.
Bagwell, K., Riordan, M.H., 1991. High and declining prices signal product quality. Am. Econ. Rev. 81 (1), 224–239.
Bar-Isaac, H., Tadelis, S., 2008. Seller reputation. Found. Trends Microecon. 4 (4), 273–351.
Battigalli, P., Charness, G., Dufwenberg, M., 2013. Deception: the role of guilt. J. Econ. Behav. Organ. 93, 227–232.
Battigalli, P., Dufwenberg, M., 2007. Guilt in games. Am. Econ. Rev. 97 (2), 170–176.
Belot, M., Bhaskar, V., Van De Ven, J., 2012. Can observers predict trustworthiness? Rev. Econ. Stat. 94 (1), 246–259.
Bénabou, R., Tirole, J., 2016. Mindful economics: The production, consumption, and value of beliefs. J. Econ. Perspect. 30 (3), 141–164.
Bicchieri, C., Lev-On, A., 2007. Computer-mediated communication and cooperation in social dilemmas: an experimental analysis. Politics Philos. Econ. 6 (2), 139–168.
Binmore, K., 2006. Why do people cooperate? Politics Philos. Econ. 5 (1), 81–96.
Bolton, G., Greiner, B., Ockenfels, A., 2018. Dispute resolution or escalation? The strategic gaming of feedback withdrawal options in online markets. Manag. Sci. 64 (9), 4009–4031.
Bolton, G.E., Katok, E., Ockenfels, A., 2004. How effective are electronic reputation mechanisms? An experimental investigation. Manag. Sci. 50 (11), 1587–1602.
Bolton, G.E., Kusterer, D.J., Mans, J., 2019. Inflated reputations: uncertainty, leniency, and moral wiggle room in trader feedback systems. Manag. Sci. 65 (11), 5371–5391.
Bracht, J., Feltovich, N., 2009. Whatever you say, your reputation precedes you: observation and cheap talk in the trust game. J. Public Econ. 93 (9-10), 1036–1044.
Brandts, J., Cooper, D.J., Rott, C., 2019. Communication in laboratory experiments. Handbook of Research Methods and Applications in Experimental Economics. Edward Elgar Publishing.
Cabral, L., Hortacsu, A., 2010. The dynamics of seller reputation: evidence from eBay. J. Ind. Econ. 58 (1), 54–78.
Cartwright, E., 2019. A survey of belief-based guilt aversion in trust and dictator games. J. Econ. Behav. Organ. 167, 430–444.
Charness, G., Dufwenberg, M., 2006. Promises and partnership. Econometrica 74 (6), 1579–1601.
Charness, G., Dufwenberg, M., 2010. Bare promises: an experiment. Econ. Lett. 107 (2), 281–283.
Chen, Y., Cramton, P., List, J.A., Ockenfels, A., 2021. Market design, human behavior, and management. Manag. Sci. 67 (9), 5317–5348.
Chen, Y., Zhang, Y., 2021. Do elicited promises affect people's trust?—Observations in the trust game experiment. J. Behav. Exp. Econ. 93, 101726.
Cressey, D.R., 1986. Why managers commit fraud. Aust. N. Z. J. Criminol. 19 (4), 195–209.
Dellarocas, C., Wood, C.A., 2008. The sound of silence in online feedback: estimating trading risks in the presence of reporting bias. Manag. Sci. 54 (3), 460–476.
Dana, J., Weber, R.A., Kuang, J.X., 2007. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. Econ. Theory 67–80.
Ellingsen, T., Johannesson, M., 2004. Promises, threats and fairness. Econ. J. 114 (495), 397–420.
Erat, S., Gneezy, U., 2012. White lies. Manag. Sci. 58 (4), 723–733.

Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.
Fréchette, G.R., 2012. Session-effects in the laboratory. Exp. Econ. 15, 1–14.
Gneezy, U., 2005. Deception: the role of consequences. Am. Econ. Rev. 95 (1), 384–394.
Gneezy, U., Rockenbach, B., Serra-Garcia, M., 2013. Measuring lying aversion. J. Econ. Behav. Organ. 93, 293–300.
Lafky, J., 2014. Why do people rate? Theory and evidence on online ratings. Games Econ. Behav. 87, 554–570.
Li, L., Xiao, E., 2014. Money talks: rebate mechanisms in reputation system design. Manag. Sci. 60 (8), 2054–2072.
Lerner, J.S, Tetlock, P.E, 1994. Accountability and social cognition. Encyclopedia of Human Behav. 1, 3098–3121.
Lerner, J.S., Tetlock, P.E., 1999. Accounting for the effects of accountability. Psychol. Bull. 125 (2), 255.
López-Pérez, R., Spiegelman, E., 2013. Why do people tell the truth? Experimental evidence for pure lie aversion. Exp. Econ. 16 (3), 233–247.
Lundquist, T., Ellingsen, T., Gribbe, E., Johannesson, M., 2009. The aversion to lying. J. Econ. Behav. Organ. 70 (1-2), 81–92.
Mayzlin, D., Dover, Y., Chevalier, J., 2014. Promotional reviews: an empirical investigation of online review manipulation. Am. Econ. Rev. 104 (8), 2421–2455.
Resnick, P., Zeckhauser, R., 2002. Trust among strangers in Internet transactions: empirical analysis of eBay's reputation system. Economics of the Internet and E-Commerce 11 (2), 23–25.
Ritov, I., Baron, J., 1992. Status-quo and omission biases. J. Risk Uncertain. 5 (1), 49–61.
Sally, D., 1995. Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. Ration. Soc. 7 (1), 58–92.
Sánchez-Pagés, S., Vorsatz, M., 2007. An experimental study of truth-telling in a sender–receiver game. Games Econ. Behav. 61 (1), 86–112.
Searle, J.R., 1975. A taxonomy of illocutionary acts. K. Gunderson (ed.). Language, Mind and Knowledge. University of Minnesota Press, Minneapolis, MN, pp. 344–369.
Serra-Garcia, M., Van Damme, E., Potters, J., 2013. Lying about what you know or about what you do? J. Eur. Econ. Assoc. 11 (5), 1204–1229.
Steiner, I.. eBay expands squaretrade warranties in home and garden. EcommerceBytes. https://www.ecommercebytes.com/cab/abn/y12/m08/i09/s04.
Tetlock, P.E., 1985. Accountability: a social check on the fundamental attribution error. Soc. Psychol. Q. 48 (3), 227–236.
Vanberg, C., 2008. Why do people keep their promises? An experimental test of two explanations 1. Econometrica 76 (6), 1467–1480.